

Richard Schussnig

Generalised Newtonian Fluids in Cardiovascular Fluid-Structure Interaction

MONOGRAPHIC SERIES TU GRAZ
STRUCTURAL ANALYSIS



Richard Schussnig

**Generalised Newtonian Fluids in
Cardiovascular Fluid-Structure Interaction**

Monographic Series TU Graz

Structural Analysis

Series Editors

G. Beer	Institute of Structural Analysis
T.-P. Fries	Institute of Structural Analysis

Monographic Series TU Graz

Structural Analysis

Richard Schussnig

Generalised Newtonian Fluids in Cardiovascular Fluid-Structure Interaction

This work is based on the dissertation "*Generalised Newtonian Fluids in Cardiovascular Fluid-Structure Interaction*", presented at Graz University of Technology, Institute of Structural Analysis in 2022.

Supervision / Assessment:

Thomas-Peter Fries (Graz University of Technology)

Marek Behr (RWTH Aachen University)

Cover photo	Gernot Beer
Cover layout	Verlag der Technischen Universität Graz
Print	Buchschmiede (Dataform Media GmbH, Vienna)

2023 Verlag der Technischen Universität Graz
www.tugraz-verlag.at

Print

ISBN 978-3-85125-943-8
Not for sale

E-Book

ISBN 978-3-85125-944-5
DOI 10.3217/978-3-85125-943-8



This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://creativecommons.org/licenses/by-nc-nd/4.0/deed.en>

This CC license does not apply to the cover, third party material (attributed to other sources) and content noted otherwise.

Contents

1	Introduction	1
1.1	Goals and outline of this work	2
1.2	Contributions and acknowledgements	3
2	Incompressible viscous flows: coupled approach	7
2.1	Governing equations	8
2.2	Residual-based stabilisation	14
2.3	PPE-based inf-sup stabilisation	16
2.3.1	Pressure Poisson equation	18
2.3.2	Momentum equation	20
2.3.3	Coupled variational formulation	21
2.4	Discretisation in space and time	22
2.5	Nonlinear solver and preconditioning	25
2.6	Timestepping schemes	32
2.6.1	Predictor-corrector approach	33
2.6.2	Fractional-step θ -scheme	34
2.7	Computational results	38
2.7.1	Kovaszny flow problem	39
2.7.2	Poiseuille flow problem in 3D	44
2.7.3	Lid-driven cavity flow	46
2.7.4	Manufactured solution in 2D	48
2.7.5	Carreau fluid channel flow in 2D	52
2.7.6	Time step selection in Carreau channel flow	56
2.7.7	Flow through an idealised aneurysm	63
2.8	Summary and conclusion	68
3	Incompressible viscous flows: split-step approach	71
3.1	A weak form for C^0 -interpolation	73
3.1.1	Pressure Poisson equation	74
3.1.2	Pressure Dirichlet condition	76
3.1.3	Momentum equation	76
3.2	Discretisation in time	77
3.3	Improving mass conservation	80
3.4	Solving the linear systems	84
3.5	Computational results	86
3.5.1	Manufactured solution in 2D	86
3.5.2	Lid-driven cavity flow	89
3.5.3	Flow through an idealised aneurysm	95
3.5.4	Computational performance study	98

3.6	Summary and conclusion	102
4	Incompressible viscous flows: extensions to biomedical applications	105
4.1	Flows in moving domains	105
4.1.1	BVS method in ALE form	109
4.1.2	ALE split-step scheme	113
4.2	Robin boundary conditions in the fluid problem	119
4.3	Stabilising dominant convection	124
4.4	Inflow over non-trivial inlets	125
4.5	Backflow stabilisation	127
4.6	Outflow conditions for truncated domains	130
4.7	Summary and conclusion	132
5	Structural dynamics	135
5.1	Governing equations	136
5.2	Tissue models accounting for fibre-reinforcement	142
5.3	Discretisation in time	144
5.4	Fully discrete, linearised problem	147
6	Structural dynamics: extensions to biomedical applications	155
6.1	Material orientation for cardiovascular settings	156
6.2	Accounting for external tissue support	162
6.3	Robin interface conditions for fluid–structure interaction	164
6.4	Prestressed reference configurations	166
6.5	Summary and conclusion	168
7	The coupled problem	171
7.1	Governing equations	176
7.2	Partitioned schemes based on a split-step method	183
7.3	Semi-implicit variants of the coupling scheme	190
7.4	Accelerating FSI schemes	193
7.5	Mesh motion in ALE methods	199
8	The coupled problem: computational results	207
8.1	Analytical solution: rectangular piston	207
8.2	Analytical solution: circular piston	212
8.3	Flow through a straight vessel	214
8.3.1	Pressure pulse	216
8.3.2	Pulsatile flow	224
8.4	ALE extension: beam in a channel	229
8.5	Abdominal aortic aneurysm	242
8.6	Iliac bifurcation	251
8.7	Aortic dissection	257
8.8	Human phonation	273
8.9	Summary and conclusion	279
9	Concluding remarks and future perspectives	283

1 Introduction

A seemingly infinite number of complex phenomena in science, engineering, medicine or nature itself share a fundamental ingredient: fluid mechanics. Ranging from the design of airfoils, jet engines, whole flying vehicles or parachutes, wind-turbines or hydro powerplants, dam or ship construction, avalanche or mudslide protection, medical device design or blood flowing through the circulatory network, ventilation or cooling systems of concert venues, tunnels or computing clusters all the way to large scale phenomena such as weather forecasts, convection in the earth's mantle, ocean currents or rivers flowing. Fluid dynamics have thus been investigated since first curious minds tried to grasp its underlying principles. During the past decades, mankind's curiosity, monetary interests or noble motives such as improving or developing strategies for green energy production led researchers from academia and industry to explore these endless planes. In an attempt to categorize, describe and ultimately predict fluid dynamics, mathematical modelling is indispensable. Naturally, the field of computational fluid dynamics (CFD) has ever since been an outstandingly active one. As mathematical and numerical concepts improved, together with continuously growing computing power, tackling larger and ever more complex problems was rendered feasible. Thus, a strong third pillar of science, being computational techniques alongside theory and experiments allows to analyse, understand and improve designs, overall processes or even develop fundamentally new concepts in science and engineering. Computational modelling is an immensely powerful tool, which (only) together with the other two established concepts can help us move forward in technology and science.

On the other side, many practical applications and phenomena involving fluid mechanics also have to incorporate effects from transient flow domains. Based on fluid flows in moving, deforming or even topology-changing geometries, accounting for the flow field's interaction with deformable structures leads to the coupled multiphysics problem of fluid-structure interaction (FSI). Such interactions need to be accounted for, whenever fluid forces acting on the structure influence the stress state and deformation in the structure, such that the flow domain is altered in turn, having a possibly drastic impact on the fluid forces again. Consequently, the engineering design process or any reliable model of such physical phenomena must consider the coupling of fluid and structure.

Great advances in the field during the past 40 years led to robust, efficient and accurate solution procedures, and interest grew in direction of applications apart from the classical engineering contexts already mentioned. The field of biomedicine is among the most promising fields in this regard, applying CFD and FSI to arterial blood flow, phonation or respiration. With modern numerical tools and computer hardware, increasingly rich and detailed models replicating reality in health and disease, so-called digital twins, are being developed and studied to gain new insights on underlying basic principles, risk factors, alternative treatment methods, to inspire medical device design and more. However,

existing methods need to be adapted to incorporate central modelling aspects or avoid possible performance bottlenecks in such applied scenarios. Including the required modelling aspects capturing effects of central relevance in practice, such as non-Newtonian rheology or outflow models relating volumetric flow rates to pressure, most often call for changes and modifications at the heart of the numerical schemes—an often overlooked or knowingly ignored detail, capable of diminishing or at least greatly reducing predictive power.

1.1 Goals and outline of this work

Major contributions of the present work are the inclusion of generalised Newtonian, or so-called quasi-Newtonian, fluid models in finite-element-based incompressible flow solvers. This key ingredient allows for an inhomogeneous viscosity to account for shear resistance depending on the shear rate, as is the case, e.g., for shear-thinning fluids such as blood or polymer melts. Such a modification of the rheological law, leading to a different representation of the viscous contribution in the fluid’s stress tensor, initiated the development of two formulations of finite-element-based incompressible flow solvers: (i) a stable or stabilised approach coupling velocity and pressure unknowns and (ii) a split-step approach, decoupling velocity and pressure based on a reformulation of the classical Navier–Stokes equations for incompressible flows via a pressure Poisson equation (PPE) with consistent boundary conditions. Both these fundamentally different concepts are used to develop state-of-the-art CFD solvers, including linearisation, decoupling, stabilisation, preconditioning and more practically relevant aspects. These topics are covered starting from Ch. 2, where the mathematical problem is presented in Sec. 2.1, followed by Sec. 2.2, introducing a first CFD solver based on a coupled velocity-pressure formulation. Ch. 3 then derives a related pressure-Poisson-based split-step approach and Ch. 4 extends both these formulations towards moving domains. Furthermore, modelling aspects and numerical techniques relevant for applications in the biomedical context are introduced. An Arbitrary Lagrangian–Eulerian formulation is combined with well established methods to counteract dominant convection, incorporate given volumetric flow data to prescribe velocity profiles on non-circular inlets, backflow stabilisation and three-element Windkessel models to determine appropriate pressure levels at terminal ends of the considered computational domain. Here, the main contribution lies in the careful intergration of these mentioned schemes into the PPE-based split-step scheme, neither impairing temporal stability nor accuracy.

Thereafter, Ch. 5 introduces displacement-based formulations for the structural solver including nearly incompressible hyperelastic continua with fiber reinforcement. Contributions of this work regarding the structural solver are an improved algorithm to generate fiber orientations in configurations with thin structural layers wetted from both sides as encountered, e.g., in aortic dissection and a pseudo-timestepping and continuation method to ease construction of prestress tensor fields to account for prestressed reference geometries.

The FSI problem is then finally introduced in Ch. 7, starting from the basic problem based on the previously considered subproblem formulations, sketching variants of the

pressure-Poisson-based, partitioned FSI scheme, semi-implicit coupling and algorithms combining Robin interface conditions with classical black-box acceleration schemes. To the best of the authors knowledge, this work is the first to *combine* semi-implicit coupling schemes with Robin interface conditions *and* the Interface Quasi-Newton Inverse Least-Squares method, allowing for added-mass stable, fast coupling of the structural displacement and fluid pressure, while the remaining subproblems are treated explicitly. Moreover, the investigated interplay of modelling aspects introduced in the fluid's and structure's subproblems with the coupling scheme and adaptive timestepping is a topic rarely commented on in literature and can thus be considered viable contributions to the field.

Some selected variants of the proposed schemes are then applied to problems in the cardiovascular context in Secs. 8.5–8.7, namely to blood flow in an idealised abdominal aortic aneurysm, an iliac bifurcation and in a dissected aorta. The tremendous improvements achieved compared to basic setups are highlighted in smaller problem settings, whereas a patient-specific geometry of an aortic dissection case shows the scheme's performance in a clinical context. Lastly, applicability in the aeroelastic regime is demonstrated by investigating human phonation in Sec. 8.8, which differs substantially in terms of present physical parameters.

This outline already indicates the scope of this work being on the development of non-Newtonian flow solvers, an extension towards FSI problems and application in patient-specific scenarios in the cardiovascular context. At the core of this work, there are combinations of physical models and their finite element formulation, but rigorous proofs thereof go beyond the scope of this work. Also, we do not aim at thoroughly introducing the fundamentals of fluid dynamics, continuum mechanics, temporal- and spatial discretisation or basic algorithms and concepts of linear algebra and numerics, as these topics are thoroughly covered in literature.

1.2 Contributions and acknowledgements

The work presented within this thesis is partially funded by the LEAD project “Mechanics, Modeling and Simulation of Aortic Dissection” (biomechaorta.tugraz.at) at Graz University of Technology and hence profited immensely from the interdisciplinary project group, ranging from computer graphics and vision over theoretical physics to biomechanical modelling and characterisation of healthy and diseased aortic tissue. Therefore, several contributions entering this thesis one way or the other are thus naturally contributed by other people than the author. The fluid mechanics contributions presented herein are closely related and interconnected with D.R.Q. Pacheco's work, where the PPE-based inf-sup stabilisation for incompressible flows is joint work published in [1] laying the foundation and [2] extending the scheme towards non-Newtonian fluids and further incorporating adaptive timestepping strategies, linearisation techniques and suitable preconditioner design. The second flow solver based on a split-step scheme is joint work as well, which lead to a first publication [3], such that a partial overlap of the basic formulations presented in Chs. 2–3 with D.R.Q. Pacheco's thesis [4] is natural. Several extensions of the basic scheme [3] towards practical applications in the haemodynamic

context presented in Ch. 4 and [5, 6] are further contributions by the author. Similarly, the application within an accelerated, semi-implicit FSI scheme is joint work presented in [5, 6] with large parts contributed by the author. As PPE-based splitting schemes are further developed at the Graz Center of Computational Engineering (gcce.tugraz.at), further extensions to multi-phase flow [7] or thrombus formation [8] are constantly improved and developed within the current research team.

Regarding the applications of the computational methods in the cardiovascular context, close collaboration with a focus on aortic dissection with G.A. Holzapfel, M. Rolf-Pissarczyk and K. Bäumlér lead to a series of publications in this context [9–12], which enter this thesis as input data for the abdominal aortic aneurysm, iliac bifurcation and aortic dissection cases presented in Secs. 8.5–8.7. In a similar manner, M. Kaltenbacher both inspired and guided the application of the FSI solver in an idealised human phonation setup in Sec. 8.8, which highlights the split-step scheme’s applicability in aeroelasticity.

It is also emphasised that *all* work of the author is carried out in close collaboration with T.P. Fries, who contributed to *all* of the work presented herein in various ways ranging from implementational details, guidance regarding model choice or mathematical formulation to general advice on presentation, project development and funding.

Main contributions of this work are thus flow solvers for generalised Newtonian fluids, namely, a stabilised coupled (monolithic) velocity-pressure formulation presented in Ch. 2, and a split-step scheme introduced in Ch. 3, both of which are based on a pressure Poisson equation with fully consistent boundary conditions. Moreover, practically relevant modelling and numerical aspects are included into the introduced CFD solvers in Ch. 4. After presenting the adopted structural dynamics formulation in Ch. 5, the following Ch. 6 demonstrates inclusion of prestressing strategies into the algorithm, suitable creation of local material orientations and how viscoelastic support can be considered. Ch. 7 then couples the split-step scheme and the displacement-based structural dynamics solver in a partitioned way, accounting for Robin coupling conditions and acceleration schemes, which in combination tremendously speed-up and robustify the fluid–structure interaction scheme. In Ch. 8, several variants of the coupling scheme are critically compared in examples ranging from academic setups to showcase expected mathematical properties up to realistic simulations in the cardiovascular setting such as flow through an iliac bifurcation or a patient-specific case of aortic dissection. Owing to this broad range of topics and their high complexity, corresponding literature surveys are given in the respective chapters for the sake of a better overview.

2 Incompressible viscous flows: coupled approach

Flows of complex fluids in engineering and medical applications are amongst the most challenging simulations today. Even with access to modern supercomputer systems and novel computational methods, practical applications can easily require thousands of CPU hours to obtain reliable solutions—meaning runtimes of days, weeks or even months depending on the problem at hand. In this context, so-called non-Newtonian fluids increase computational demand further, introducing more involved rheological laws and associated fields to the flow problem. To counteract the rapidly increasing computational complexity, generalised Newtonian (or quasi-Newtonian) fluids are often considered a well balanced middle-ground in terms of phenomenological accuracy and associated numerical effort. In such models, viscoelastic effects (see, e.g., [13–15]) are neglected, but a spatially varying viscosity allows reproducing flow characteristics of particular importance for example in polymeric flows [16] or pathologic vessel configurations or malformations such as aneurysms, stenoses or aortic dissections [13, 17, 18]. However, adapting the rheological law being at the very heart of any flow solver, subtle differences can have drastic consequences on the performance of the overall method. The developments regarding flow solvers introduced in the present work are the formulation of new schemes to incorporate generalised Newtonian rheological laws and adaptation of existing concepts available for incompressible flow of Newtonian fluids to account for inhomogeneous viscosity.

A fundamental ingredient for the methods presented in this chapter is the reformulation of a pressure Poisson equation (PPE) by Pacheco and Steinbach [19], that includes fully consistent boundary conditions, and allows for inhomogeneous viscosity and variable density. This work exploits the first two of these features, while the fluid density is assumed constant. This modified PPE leads to several possibilities: first, it can be used as part of a stabilisation when applying equal-order interpolation of velocity and pressure, initially introduced for stationary flows of Newtonian fluids in [1] and extended to a stabilised generalised Newtonian flow solver with semi-implicit, adaptive timestepping and an effective preconditioner in [2]. Second, the PPE can be used to recover fluid pressure from a given velocity field in a split-step framework, such that velocity and pressure fields may be separated. This second idea builds on the work of [20, 21] which decouples momentum and continuity equations in the Newtonian case. A non-trivial generalisation to quasi-Newtonian fluids is introduced in [3], while an extension via the Arbitrary Lagrangian–Eulerian (ALE) framework to account for a moving and deforming reference fluid domain is presented in [5]. Herein, we focus on the accomplishments from [2] and [5], but also introduce the necessary ingredients and motivations from [1] and [3], all of which build the basis for Chs. 2 and 3.

Before focusing on any of these two options being a coupled velocity-pressure formulation or a PPE-based split-step scheme, the joint basic problem description is introduced first in Sec. 2.1. Then, the remainder of this chapter focuses on the solution approach treating velocity and pressure in a monolithic, coupled form, discussed in Sec. 2.2. Afterwards, Ch. 3 considers a split-step approach decoupling the velocity and pressure fields. Closing the discussion of the fluid subproblem, Ch. 4 presents numerical techniques and modelling aspects extending applicability of these two frameworks in the context of haemodynamics. Therein, flows in moving and deforming domains, stabilisations of convective-dominant and re-entrant flow and in- and outflow conditions in the cardiovascular setting are discussed.

2.1 Governing equations

Let us consider a domain $\Omega \subset \mathbb{R}^d$ in $d = 2$ or 3 spatial dimensions with Lipschitz boundary $\Gamma := \partial\Omega$, which we further decompose into two non-overlapping Neumann and Dirichlet segments Γ_N and Γ_D , respectively, where $\Gamma_D \cup \Gamma_N = \Gamma$ and $\Gamma_D \cap \Gamma_N = \emptyset$. The balance of linear momentum and continuity equations for incompressible flow compose the standard Navier–Stokes system given as

$$\rho \partial_t \mathbf{u} + \rho (\nabla \mathbf{u}) \mathbf{u} - \nabla \cdot \boldsymbol{\sigma} = \mathbf{b} \quad \text{in } \Omega \times (0, T], \quad (2.1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \Omega \times (0, T], \quad (2.2)$$

with the sought fluid velocity \mathbf{u} and pressure p as unknowns, the fluid's density ρ and a volumetric force \mathbf{b} acting on the fluid in the time interval $I_t := (0, T]$ from $t = 0$ to a specified target time $t = T$. In Eqn. (2.1), the nonlinear term, $(\nabla \mathbf{u}) \mathbf{u}$, is written in convective form, which is obtained from the conservative form, $\nabla \cdot (\mathbf{u} \otimes \mathbf{u})$, using

$$\nabla \cdot (\mathbf{u} \otimes \mathbf{u}) = (\nabla \mathbf{u}) \mathbf{u} + \mathbf{u} \nabla \cdot \mathbf{u} = (\nabla \mathbf{u}) \mathbf{u}, \quad (2.3)$$

since $\nabla \cdot \mathbf{u} = 0$ due to the incompressibility constraint (2.2). The Cauchy stress $\boldsymbol{\sigma}$ can be split into the viscous stress tensor \mathbf{S} and an isotropic contribution in terms of the pressure as

$$\boldsymbol{\sigma} := -p \mathbf{I} + \mathbf{S}, \quad \text{where } \mathbf{S} := 2\mu \nabla^s \mathbf{u}. \quad (2.4)$$

Therein, \mathbf{I} is the $d \times d$ unit tensor, the symmetric velocity gradient $\nabla^s \mathbf{u}$ is defined as

$$\nabla^s \mathbf{u} := \frac{1}{2} (\nabla \mathbf{u} + \nabla^\top \mathbf{u}), \quad \text{with } \nabla^\top \mathbf{u} := (\nabla \mathbf{u})^\top \quad (2.5)$$

and the viscosity is denoted by μ . In the case of Newtonian fluids, $\mu = \text{const.}$ linearly relates strains to stresses, which for quasi-Newtonian fluids, however, is generalised using a nonlinear relation $\eta(\dot{\gamma}) : \mathbb{R}^+ \rightarrow \mathbb{R}^+ \setminus \{0\}$ capturing more complex rheological behaviour and effects such as bulk flow, shear-thickening or shear-thinning. These phenomena are incorporated by modelling the apparent fluid viscosity dependent on the shear rate $\dot{\gamma}$,

$$\mu := \eta(\dot{\gamma}), \quad (2.6)$$

where the shear rate is given by

$$\dot{\gamma} := \sqrt{1/2 \nabla^s \mathbf{u} : \nabla^s \mathbf{u}}. \quad (2.7)$$

The rheological models considered herein are of particular importance in the fields of biomedical engineering and medicine for the modelling of blood flow, but also in industrial applications such as flows of polymer melts. As we shall see in the following sections, generalised Newtonian fluid models are easily exchanged in all of the presented solution approaches, but a first more general rheological law is given by [13]

$$\eta(\dot{\gamma}) := \eta_\infty + (\eta_0 - \eta_\infty) [\kappa + (\lambda \dot{\gamma})^a]^{\frac{b-1}{a}}, \quad (2.8)$$

where η_∞ and η_0 denote lower and upper viscosity limits, respectively, and the remaining constants κ , λ , a and b are fitting parameters specifying the transition from η_∞ to η_0 . Eqn. (2.8) boils down to the Power law for $\kappa = \eta_\infty = 0$, but at the same time can also represent the Carreau ($\kappa = 1$ and $a = 2$) and Carreau-Yasuda ($\kappa = 1$) models of particular relevance in haemodynamic flows. Finally, when choosing $\eta_\infty = \eta_0$, we retrieve the standard Newtonian model with constant viscosity.

The standard Navier–Stokes system (2.1)–(2.2) is then supplemented by an initial condition for the velocity at $t = 0$,

$$\mathbf{u} = \mathbf{u}_0 \quad \text{in } \Omega, \quad (2.9)$$

with $\nabla \cdot \mathbf{u}_0 = 0$ for consistency and Dirichlet conditions prescribing the fluid’s velocity vector on the Dirichlet boundary segment Γ_D ,

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D. \quad (2.10)$$

Natural boundary conditions enforced weakly via boundary integrals stemming from integration by parts differ depending on the specific formulation chosen for the viscous stress tensor. Using the stress-divergence form, i.e.,

$$\nabla \cdot \boldsymbol{\sigma} = \nabla \cdot (-p\mathbf{I} + \mathbf{S}),$$

applying Gauss’ divergence theorem on the entire stress term allows prescribing real tractions in direction of the unit outward normal \mathbf{n} , since

$$\mathbf{t} = \boldsymbol{\sigma} \mathbf{n} = (-p\mathbf{I} + 2\mu \nabla^s \mathbf{u}) \mathbf{n} \quad \text{on } \Gamma_N, \quad (2.11)$$

which directly corresponds to the Cauchy stress present in the fluid. Hence, this formulation is preferred especially when considering interface-coupled problems such as fluid–structure interaction. Alternatively, we can rewrite the stress term to obtain the generalised Laplacian form,

$$\begin{aligned}
\nabla \cdot \boldsymbol{\sigma} &= -\nabla p + \nabla \cdot (2\mu \nabla^s \mathbf{u}) \\
&= -\nabla p + \mu [\nabla \cdot (\nabla \mathbf{u}) + \nabla (\nabla \cdot \mathbf{u})] + 2\nabla^s \mathbf{u} \nabla \mu \\
&= -\nabla p + \mu \Delta \mathbf{u} + 2\nabla^s \mathbf{u} \nabla \mu,
\end{aligned} \tag{2.12}$$

which can be used to prescribe boundary data in terms of so-called pseudo tractions

$$\tilde{\mathbf{t}} = (-p\mathbf{I} + \mu \nabla \mathbf{u}) \mathbf{n} \quad \text{on } \Gamma_N, \tag{2.13}$$

integrating only the velocity Laplacian and pressure gradient by parts. The reason for enforcing such an artificial boundary condition is that prescribing the real traction, e.g., as $\mathbf{t} = \mathbf{0}$ is not desired in scenarios, where the computational domain has to be truncated. Setting real tractions at the outlets leads to more fanned velocity vectors on the outlet and redirection of the velocity field towards the exterior, whereas enforcing pseudo tractions $\tilde{\mathbf{t}} = \mathbf{0}$ allows the fluid to exit the domain with little to no disturbance and additionally enforces a mean zero pressure on a planar outlet (see, e.g., [22–25] for a more in-depth discussion on outlet conditions). For generalised Newtonian fluids, however, pseudo tractions do not enforce mean pressures over planar outlets exactly, but *approximately* [26]. Also note here that the difference between the two variants is simply

$$\mathbf{t} = \tilde{\mathbf{t}} + \mu (\nabla^\top \mathbf{u}) \mathbf{n}, \tag{2.14}$$

which will be used at a later point to augment the final weak forms, switching from one traction condition to the other by simply adapting a single scalar parameter. However, this comes at the cost of considering an additional boundary integral depending on the formulation of the stress term, which might further influence the linear solver and obtained solution of the momentum balance equation. Thus we shall use both forms in the following, but stick to the most natural viscous stress term, meaning, the viscous stress term resulting in the natural boundary condition we wish to enforce.

For pure Dirichlet problems, i.e., when $\Gamma_D = \Gamma$ (in contrast to our requirements posed initially), two additional complications need to be considered: first, the Dirichlet datum needs to comply to

$$\int_{\Omega} \nabla \cdot \mathbf{u} \, d\Omega = \langle \nabla \cdot \mathbf{u}, 1 \rangle_{\Omega} = 0 = \langle \mathbf{u}, \mathbf{n} \rangle_{\Gamma_D} = \langle \mathbf{g}, \mathbf{n} \rangle_{\Gamma_D} = \int_{\Gamma_D} \mathbf{g} \cdot \mathbf{n} \, d\Gamma_D \tag{2.15}$$

for solvability, denoting by $\langle \cdot, \cdot \rangle_{\Omega}$ and $\langle \cdot, \cdot \rangle_{\Gamma_D}$ the standard L^2 scalar products in Ω and on Γ_D , respectively. Second, the pressure is in this setting only uniquely defined up to a constant. A suitable pressure scaling can be achieved requiring a mean pressure equal to zero in the whole domain or over the boundary via

$$\langle p, 1 \rangle_{\Omega}, \quad \text{or} \quad \langle p, 1 \rangle_{\Gamma_D}. \tag{2.16}$$

This augmentation is admissible in the context of pure flow problems and can easily be included. In fluid–structure interaction, however, the interface coupling conditions include fluid tractions, where the pressure scaling directly enters. Despite being relevant

in scenarios where the fluid is fully enclosed by the structure, we will not focus on pressure uniqueness in enclosed flow problems and hence refrain from a more detailed discussion at this point.

Before diving deeper into the two solution paradigms of coupled and split-step schemes, let us introduce an alternative to the classical Navier–Stokes system comprised of momentum balance and continuity equations (2.1)–(2.2) and initial and boundary conditions as given in Eqns. (2.9)–(2.11). This presentation follows Pacheco [4], is closely related to joint work presented in [1–3, 5] and further incorporates adaptations of [19–21]. The formulation shown here allows switching from the generalised Laplacian ($\chi = 0$) to the stress-divergence form ($\chi = 1$) using a single scalar parameter χ . The system (2.1)–(2.2) and (2.9)–(2.11) is rewritten as

$$\rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2 \nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b} \quad \text{in } \Omega \times (0, T], \quad (2.17)$$

$$-\Delta p + \nabla \cdot [2 \nabla^s \mathbf{u} \nabla \mu - \rho (\nabla \mathbf{u}) \mathbf{u}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = -\nabla \cdot \mathbf{b} \quad \text{in } \Omega \times [0, T], \quad (2.18)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D \times (0, T], \quad (2.19)$$

$$\left(-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \right) \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N \times (0, T], \quad (2.20)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0, \quad (2.21)$$

$$\nabla \cdot \mathbf{u}_0 = 0 \quad \text{in } \Omega, \quad (2.22)$$

$$-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot \left[\mu (\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u}) \mathbf{n} - \bar{\mathbf{t}} \right] = p \quad \text{on } \Gamma_N \times [0, T], \quad (2.23)$$

$$\mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u} - \rho (\nabla \mathbf{u}) \mathbf{u} + 2 \nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] = \mathbf{n} \cdot \nabla p \quad \text{on } \Gamma_D \times [0, T]. \quad (2.24)$$

Herein, $\bar{\mathbf{t}} := \tilde{\mathbf{t}} + \chi \mu (\nabla^\top \mathbf{u}) \mathbf{n}$ denotes the given traction prescribed on the Neumann boundary, which is equal to pseudo tractions $\tilde{\mathbf{t}}$ or real tractions \mathbf{t} , depending on the parameter χ . The homogeneous Newtonian case is also included in the above relations, since several terms vanish when $\nabla \mu \equiv 0$, and the formulation by Liu [21] is recovered. Eqn. (2.18) is a Poisson equation in the pressure (PPE) with Neumann and Dirichlet boundary conditions given in Eqns. (2.23) and (2.24) which will be of use in both a coupled solution approach and a split-step scheme. Thm. 2.1.1 states equivalence of the original Navier–Stokes system given in Eqns. (2.1)–(2.2), and (2.9)–(2.11) and the proposed alternative system (2.17)–(2.24).

Theorem 2.1.1. *For sufficiently regular $p, \mathbf{u}, \mathbf{b}, \mathbf{g}, \bar{\mathbf{t}}$, systems (2.1)–(2.2), (2.9)–(2.11) and (2.17)–(2.24) are equivalent.*

Proof. First, show that the standard Navier–Stokes system (2.1)–(2.2), (2.9)–(2.11) implies the alternative system (2.17)–(2.24). We obtain the modified momentum balance equation (2.17) from Eqn. (2.1) rewriting the stress terms as

$$\nabla \cdot \boldsymbol{\sigma} = -\nabla p + \mu [\nabla \cdot (\nabla \mathbf{u}) + \chi \nabla (\nabla \cdot \mathbf{u})] + 2 \nabla^s \mathbf{u} \nabla \mu, \quad (2.25)$$

including the parameter χ being 0 or 1 for the generalised Laplacian and stress-divergence formulations, respectively.

Applying the divergence operator to the momentum equation (2.1) or the just rewritten Eqn. (2.17), the acceleration term vanishes, since we have

$$\nabla \cdot (\partial_t \mathbf{u}) = \partial_t (\nabla \cdot \mathbf{u}) = 0, \quad (2.26)$$

assuming sufficient regularity. The viscous stress contribution is

$$\begin{aligned} & \nabla \cdot (\mu \Delta \mathbf{u}) + \chi \nabla \cdot [\nabla (\nabla \cdot \mathbf{u})] + 2 \nabla \cdot (\nabla^s \mathbf{u} \nabla \mu) \\ &= \nabla \mu \cdot \Delta \mathbf{u} + \mu \nabla \cdot (\Delta \mathbf{u}) + 2 \nabla \cdot (\nabla^s \mathbf{u} \nabla \mu) \\ &= \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u}) - \nabla \times (\nabla \times \mathbf{u})] + \mu \Delta (\nabla \cdot \mathbf{u}) + 2 \nabla \cdot (\nabla^s \mathbf{u} \nabla \mu) \\ &= -\nabla \mu \cdot [\nabla \times (\nabla \times \mathbf{u})] + 2 \nabla \cdot (\nabla^s \mathbf{u} \nabla \mu), \end{aligned} \quad (2.27)$$

using $\nabla \cdot \mathbf{u} = 0$ repeatedly together with the identity

$$\Delta \mathbf{u} \equiv \nabla (\nabla \cdot \mathbf{u}) - \nabla \times (\nabla \times \mathbf{u}), \quad (2.28)$$

such that we directly obtain

$$-\nabla \cdot (\nabla p) = -\Delta p = \nabla \cdot [\rho (\nabla \mathbf{u}) \mathbf{u} - 2 \nabla^s \mathbf{u} \nabla \mu] + [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu - \nabla \cdot \mathbf{b},$$

which is identical to Eqn. (2.18). The Dirichlet condition on the pressure is obtained dotting the Neumann boundary condition of the standard Navier–Stokes system by the unit outward normal \mathbf{n}

$$\begin{aligned} \mathbf{n} \cdot [(-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u}) \mathbf{n}] &= \mathbf{n} \cdot \bar{\mathbf{t}} \\ -p \mathbf{n} \cdot \mathbf{n} = -p &= \mathbf{n} \cdot [\bar{\mathbf{t}} - \mu (\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u}) \mathbf{n}], \end{aligned}$$

and adding $\mu \nabla \cdot \mathbf{u} = 0$ being an arbitrary multiplicity of the continuity equation to get

$$p = -\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot [\mu (\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u}) \mathbf{n} - \bar{\mathbf{t}}], \quad (2.29)$$

which gives Eqn. (2.23) when restricted to the Neumann boundary Γ_N . To complete the first part of the proof, the Neumann condition on the pressure is constructed dotting the momentum balance equation with \mathbf{n} , which results with the identity (2.28) directly to (2.24). Now, the remaining equations are simply Dirichlet conditions on the velocity (2.20), the initial condition for the velocity (2.21) and the continuity equation restricted to $t = 0$ to obtain Eqn. (2.22).

Now, it remains to show that the proposed alternative system (2.17)–(2.24) implies the standard Navier–Stokes system (2.1)–(2.2), (2.9)–(2.11). Start off by taking the divergence of Eqn. (2.17) and add the result to the PPE (2.18), giving

$$\rho \nabla \cdot (\partial_t \mathbf{u}) - \nabla \cdot (\mu \Delta \mathbf{u}) - \chi \nabla \cdot [\mu \nabla (\nabla \cdot \mathbf{u})] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = 0,$$

where we insert

$$\begin{aligned}\nabla \cdot (\mu \Delta \mathbf{u}) &= \nabla \mu \cdot \Delta \mathbf{u} + \mu \nabla \cdot (\Delta \mathbf{u}) = \nabla \mu \cdot \Delta \mathbf{u} + \mu \Delta (\nabla \cdot \mathbf{u}), \\ \nabla \cdot [\mu \nabla (\nabla \cdot \mathbf{u})] &= \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u})] + \mu \nabla \cdot [\nabla (\nabla \cdot \mathbf{u})],\end{aligned}$$

such that we end up with

$$\rho \partial_t \nabla \cdot \mathbf{u} - (1 + \chi) \mu \Delta (\nabla \cdot \mathbf{u}) - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu - \nabla \mu \cdot [\chi \nabla (\nabla \cdot \mathbf{u})] = 0.$$

Then, using Eqn. (2.28), this can be rewritten as

$$\rho \partial_t \nabla \cdot \mathbf{u} - (1 + \chi) \mu \Delta (\nabla \cdot \mathbf{u}) - (1 + \chi) \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u})] = 0,$$

which is a heat equation in the auxiliary variable $\Phi := \nabla \cdot \mathbf{u}$,

$$\rho \partial_t \Phi - (1 + \chi) \nabla \cdot (\mu \nabla \Phi) = 0, \quad (2.30)$$

with variable diffusion coefficient $(1 + \chi) \mu / \rho$. Due to Eqn. (2.22), i.e., $\nabla \cdot \mathbf{u}_0 = 0$, we have $\Phi = 0$ as initial condition at $t = 0$, while boundary conditions for Eqn. (2.30) on Γ_D are derived subtracting the modified momentum balance equation (2.17) dotted with the unit outward normal \mathbf{n} from the pressure Neumann condition (2.24), which gives

$$\mathbf{n} \cdot [\mu \Delta \mathbf{u} + \chi \mu \nabla (\nabla \cdot \mathbf{u}) + \mu \nabla \times (\nabla \times \mathbf{u})] = 0 \quad \text{on } \Gamma_D,$$

which, again using Eqn. (2.28), reduces to

$$\mu(1 + \chi) \mathbf{n} \cdot [\nabla (\nabla \cdot \mathbf{u})] = \mu(1 + \chi) \mathbf{n} \cdot \nabla \Phi = 0 \quad \implies \quad \mathbf{n} \cdot \nabla \Phi = 0 \quad \text{on } \Gamma_D. \quad (2.31)$$

On the Neumann boundary Γ_N , we dot the traction condition (2.20) with \mathbf{n} and add the Dirichlet condition on the pressure (2.23) to directly obtain

$$\mu \nabla \cdot \mathbf{u} = 0 \quad \implies \quad \Phi = 0 \quad \text{on } \Gamma_N. \quad (2.32)$$

To sum up, we have a heat equation in $\Phi := \nabla \cdot \mathbf{u}$ with zero initial, Dirichlet and Neumann conditions, hence, we have $\Phi \equiv 0$. The velocity field \mathbf{u} obtained from the alternative system (2.17)–(2.24) thus implicitly fulfils the continuity condition. At this point, all that remains is showing that the modified momentum equation (2.17) implies the original one (2.1), which trivially follows from using Eqn. (2.25) again. The alternative system fulfils the initial and boundary conditions of the standard Navier–Stokes system, enforces momentum balance and continuity of mass, concluding the proof for both the stress-divergence and Laplacian form of the viscous stress tensor. \square

The alternative system (2.17)–(2.24) can be used to replace the original Navier–Stokes system, but has some fundamentally different properties as shall be seen considering a coupled velocity-pressure formulation in the following section or, alternatively, adopting a split-step approach decoupling the momentum balance equation and the PPE as introduced in Ch. 3.

2.2 Residual-based stabilisation

Adopting a mixed formulation coupling velocity and pressure, care must be taken when selecting velocity and pressure basis functions, as violating the famous Ladyzhenskaya–Babuška–Brezzi (LBB) condition leads to unstable methods [27, 28]. This stability condition necessitates the use of appropriate function spaces, e.g., choosing a polynomial degree one order higher for the velocity interpolation than for the pressure ansatz, so-called Taylor–Hood finite element pairs [29]. In a coupled solution approach, but also when using standard projection-type schemes, the incompressibility constraint does not allow for equal-order interpolation irrespective of the viscous model employed (see, e.g., [27, 28, 30] and references therein).

One possible remedy to circumvent this restriction consists of stabilising the variational formulation and thereby breaking the saddle-point structure of the monolithic velocity–pressure block system. One of the earliest methods of this class by Brezzi and Pitkäranta [31] bypasses the compatibility condition by perturbing the continuity equation, which results in a non-zero pressure–pressure block. In order to improve accuracy, residual-based stabilisations such as the pressure-stabilised Petrov–Galerkin [32] (PSPG) or Galerkin Least-Squares [33] (GLS) methods and similar schemes [32, 34, 35] additionally scale added terms with the momentum balance residual. This leads to small perturbations where the solution is already accurate enough and a relaxation of the incompressibility constraint not only vanishing with decreasing element size h_e , but also with the strong form residual itself.

Within this work, the subscript Ω in $\langle \cdot, \cdot \rangle_\Omega$ is omitted whenever possible, while other domains of integration are clearly defined, keeping the subscripts for these terms only. Then, the standard PSPG method seeks for $(\mathbf{u}, p) \in X_h^u \times X_h^p$, with $X_h^u \subset [H^1(\Omega)]^d$, $X_h^p \subset L^2(\Omega)$, and $\mathbf{u}_h|_{\Gamma_D} = \mathbf{g}_h$, such that

$$\begin{aligned} \langle \mathbf{w}_h, \rho [\partial_t \mathbf{u}_h + (\nabla \mathbf{u}_h) \mathbf{u}_h] \rangle + \langle \nabla^s \mathbf{w}_h, 2\mu (\nabla^s \mathbf{u}_h) \nabla^s \mathbf{u}_h \rangle - \langle \nabla \cdot \mathbf{w}_h, p_h \rangle \\ = \langle \mathbf{w}_h, \mathbf{t} \rangle_{\Gamma_N} + \langle \mathbf{w}_h, \mathbf{b} \rangle, \end{aligned} \quad (2.33)$$

$$\langle q_h, \nabla \cdot \mathbf{u}_h \rangle + \sum_{e=1}^{N_e} \langle \tau_e \nabla q_h, \rho [\partial_t \mathbf{u}_h + (\nabla \mathbf{u}_h) \mathbf{u}_h] + \nabla p_h - \nabla \cdot \mathbf{S}_h - \mathbf{b} \rangle_{\Omega_e} = 0, \quad (2.34)$$

holds for all $(\mathbf{w}_h, q_h) \in X_h^u \times X_h^p$ with $\mathbf{w}_h|_{\Gamma_D} = \mathbf{0}$ and

$$\nabla \cdot \mathbf{S}_h = \mu (\nabla^s \mathbf{u}_h) \Delta \mathbf{u}_h + 2\nabla^s \mathbf{u}_h \nabla \mu (\nabla^s \mathbf{u}_h).$$

Here, real tractions \mathbf{t} are enforced on Γ_N and the stabilisation term is composed of element contributions summed over all N_e elements Ω_e . The stabilisation parameter τ_e is defined in a variety of ways in literature, e.g., as

$$\tau_e := \frac{1}{\rho} \left[\left(\frac{\alpha_1}{\rho h_e^2 \mu_h} \right)^2 + \left(\frac{\alpha_2}{h_e} \|\mathbf{u}_h\| \right)^2 + \left(\frac{\alpha_3}{\Delta t} \right)^2 \right]^{-1/2}, \quad (2.35)$$

with the element size h_e defined as

$$h_e := \text{meas}(\Omega_e)^{1/d} = \left(\int_{\Omega_e} 1 \, d\Omega \right)^{1/d}, \quad (2.36)$$

and where popular choices are $\alpha_3 = 0$ for stationary flow, $\alpha_2 = 0$ in diffusion-dominated scenarios, and consequently $\alpha_2 = \alpha_3 = 0$ for stationary, diffusive flows [27]. Alternatively, one might consider $\alpha_1 = 2\alpha_2 = 2\alpha_3 = 4$ being a popular choice [36–38]. Recommendations also exist for linear elements, $\alpha_1 = 12$ [39], some of which were also extended towards higher-order elements [27, 40].

Viable alternatives include schemes based on artificial compressibility [14, 41], penalty formulations [42], residual-free bubbles [43] or polynomial pressure projections [44]. Nonetheless, residual-based stabilisations remain attractive for practical applications due to their simple implementation without the need for unusual data structures, negligible added computational cost and favourable accuracy. One downside of residual-based stabilisations becomes apparent when using linear elements: the diffusive contribution to the momentum residual cannot be represented appropriately, since the second-order derivatives in

$$\sum_{e=1}^{N_e} \langle \tau_e \nabla q_h, \nabla \cdot \mathbf{S}_h \rangle_{\Omega_e} = \sum_{e=1}^{N_e} \langle \tau_e \nabla q_h, \mu (\nabla^s \mathbf{u}_h) \Delta \mathbf{u}_h + 2 \nabla^s \mathbf{u}_h \nabla \mu (\nabla^s \mathbf{u}_h) \rangle_{\Omega_e},$$

may vanish completely, independent of the rheological law applied. That is, even if $\nabla \mu (\nabla^s \mathbf{u}) \equiv 0$ as in the Newtonian limit, the viscous contribution to the residual is lost. Unfortunately, a loss of accuracy results especially in regions with dominant diffusion [45, 46], even though the PSPG method remains consistent, having $\tau_e \rightarrow 0$ as $h_e \rightarrow 0$. In order to eliminate the spurious pressure boundary layers triggered by the incomplete residual being responsible for the reduced accuracy or even spoiling the solution entirely, one may resort to interior penalty techniques [45, 47], reconstruction of the velocity Laplacian [46, 48], pressure gradient projection [49, 50] or local projection stabilisations [51, 52].

In the context of non-Newtonian fluids, unfortunately, the matter increases in complexity: above mentioned methods relying on or assuming homogeneous viscosity are not directly applicable. Consequently, inf-sup stable finite element pairs, e.g., the classical Taylor–Hood finite element is often resorted to [53–62]. Stabilisation methods based solely on the pressure are directly applicable to the generalised Newtonian case, among them are the penalty method [63–66] or the pressure Poisson stabilisation [31, 67], which were shown to perform rather poorly in the case of inhomogeneous viscosity.

In contrast to that, the pressure gradient projection method by Codina and Blasco [49] does not suffer from such drawbacks, but is considerably more expensive. Moreover, John et al. [17] successfully applied local pressure projection stabilisation [44] to haemodynamic flows with inhomogeneous viscosity. Similar to the Newtonian case, residual-based

formulations such as PSPG and variational multiscale (VMS) methods may also be employed [38, 68–70], but again, their incomplete residual for linear elements can lead to numerical artifacts such as spurious pressure boundary layers and poor conservation in low-order discretisations [1]. One might also treat the viscous stress as an additional tensor-valued unknown [71, 72], which might be necessary in the case of viscoelastic fluids to begin with, but this leads to considerably increased computational costs for three-dimensional flow problems.

Summing up, an improvement of residual-based stabilisations for coupled velocity-pressure formulations of generalised Newtonian fluids in the diffusive-dominant regime is highly needed. As it turns out, we can reformulate terms with second-order derivatives in the stabilised incompressibility-preserving continuity equation as first-order boundary integrals. The resulting structure of the equations is similar to standard residual-based formulations, but the continuity equation is considered in an element-weighted manner, whereas the added stabilisation terms stemming from a consistent PPE are considered globally. The stabilisation parameter is in this sense a scaling parameter blending the continuity equation and the PPE—both enforcing incompressibility. Hence, we are able to substantially increase accuracy by incorporating the full residual while at the same time making the volume integration per element cheaper and widening the range of admissible stabilisation parameters when compared to the traditional PSPG method.

It shall be noted at this point that the proposed stabilisation aims at low Reynolds number flows, where diffusive effects are dominant. This present method is thus not to be confused with other residual-based stabilisations also incorporating stabilising terms to counteract dominant convection or to enhance mass conservation such as streamline upwind Petrov–Galerkin [73, 74] (SUPG), Galerkin Least-Squares [33] (GLS), variational multiscale [38, 68–70] (VMS), artificial diffusion [75] or grad-div stabilisation [76, 77] methods. However, the present method can be combined with some of these concepts to yield methods effective under more general flow conditions, as shall be commented on in Sec. 2.4. The interested reader is referred to, e.g., [27, 78–81] for a more rigorous introduction and in-depth discussion of stabilised finite element methods in the context of incompressible flow problems.

2.3 PPE-based inf-sup stabilisation

The stabilised formulation is built combining the PPE-based system (2.17)–(2.24) with a scaled continuity equation $\beta \nabla \cdot \mathbf{u}$ added to the PPE (2.18), yielding

$$\rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2 \nabla^s \mathbf{u} \nabla \mu = \mathbf{b} - \nabla p \quad \text{in } \Omega \times (0, T], \quad (2.37)$$

$$\beta \nabla \cdot \mathbf{u} + \nabla \cdot [2 \nabla^s \mathbf{u} \nabla \mu - \rho (\nabla \mathbf{u}) \mathbf{u} + \mathbf{b}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = \Delta p \quad \text{in } \Omega \times [0, T], \quad (2.38)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D \times (0, T], \quad (2.39)$$

$$[-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u}] \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N \times (0, T], \quad (2.40)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0, \quad (2.41)$$

$$\mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u} - \rho (\nabla \mathbf{u}) \mathbf{u} + 2 \nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] = \mathbf{n} \cdot \nabla p \quad \text{on } \Gamma \times [0, T], \quad (2.42)$$

with a constant β to be defined later and $\chi \in \{0, 1\}$ to switch between the generalised Laplacian ($\chi = 0$) and stress-divergence forms ($\chi = 1$). Since $\beta \nabla \cdot \mathbf{u}$ is added to the PPE, Thm. 2.1.1 is slightly adapted in the following.

Theorem 2.3.1. *For sufficiently regular $p, \mathbf{u}, \mathbf{b}, \mathbf{g}, \bar{\mathbf{t}}$, systems (2.1)–(2.2), (2.9)–(2.11) and (2.37)–(2.42) are equivalent.*

Proof. Following Thm. 2.1.1, we show that the classical Navier–Stokes system in (2.1)–(2.2) and (2.9)–(2.11) implies the proposed stabilised monolithic system (2.37)–(2.42). This can be shown by applying an identical series of manipulations given in Eqns. (2.25)–(2.28) and adding $\beta \nabla \cdot \mathbf{u}$ to the left-hand side, resulting in the PPE (2.38),

$$\beta \nabla \cdot \mathbf{u} + \nabla \cdot [2\nabla^s \mathbf{u} \nabla \mu - \rho(\nabla \mathbf{u}) \mathbf{u} + \mathbf{b}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = \Delta p.$$

The rest of the first part of the proof remains unaltered up to the Dirichlet condition on the pressure, which is not included in the new system, and the pressure Neumann condition, which is not restricted to Γ_D , but holds on the whole $\Gamma := \partial\Omega$.

Starting from the PPE-based stabilised method (2.37)–(2.42), equivalence is shown applying the divergence to the rewritten momentum balance equation (2.37) and adding it to the PPE (2.38) to get

$$\nabla \cdot (\rho \partial_t \mathbf{u} + \beta \mathbf{u}) - \nabla \cdot (\mu \Delta \mathbf{u}) - \chi \nabla \cdot [\mu \nabla (\nabla \cdot \mathbf{u})] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = 0,$$

which can be rewritten in a similar manner as before, yielding

$$\rho \partial_t \nabla \cdot \mathbf{u} + \beta \nabla \cdot \mathbf{u} - (1 + \chi) \mu \Delta (\nabla \cdot \mathbf{u}) - (1 + \chi) \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u})] = 0.$$

This is a diffusion-reaction equation in $\Phi := \nabla \cdot \mathbf{u}$

$$\rho \partial_t \Phi + \beta \Phi - (1 + \chi) \nabla \cdot (\mu \nabla \Phi) = 0, \quad (2.43)$$

for which initial- and boundary conditions are constructed very much similar to the procedure in Thm. 2.1.1, giving

$$\Phi = \nabla \cdot \mathbf{u} = \nabla \cdot \mathbf{u}_0 = 0 \quad \text{at } t = 0, \quad (2.44)$$

$$\mathbf{n} \cdot \nabla \Phi = 0 \quad \text{on } \Gamma. \quad (2.45)$$

Thus, $\Phi \equiv 0$ is the solution to (2.43)–(2.45), and consequently, the alternative system (2.37)–(2.42) also ensures $\nabla \cdot \mathbf{u} = 0$ implicitly. All that remains is to show that the modified momentum balance equation implies the original form, which directly follows from Eqn. (2.25), thereby completing the proof. \square

Note, however, that with introducing the PPE and replacing the incompressibility constraint by Eqn. (2.38), we also increased regularity requirements on velocity and pressure. In addition to that, a weak form corresponding to the newly devised strong form in Eqns. (2.37)–(2.42) remains to be derived such that discretisation with standard C^0 -continuous Lagrangian finite elements is admissible. Also, second-order derivatives are

still present at a first glance, such that the viscous contribution to the residual might be lost once again. As turns out, an improvement upon the classical PSPG formulation can be derived, addressing these issues simply by rewriting the variational form corresponding to the alternative system (2.37)–(2.42).

2.3.1 Pressure Poisson equation

The starting point for deriving a suitable variational formulation of the PPE with an added term $\beta \nabla \cdot \mathbf{u}$ as given in Eqn. (2.38) is

$$-\Delta p = \nabla \cdot [\rho(\nabla \mathbf{u})\mathbf{u} - 2\nabla^s \mathbf{u} \nabla \mu - \mathbf{b}] + [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu - \beta \nabla \cdot \mathbf{u}.$$

Multiplying with a continuous test function $q \in H^1(\Omega)$ and applying Green's formula results in

$$\begin{aligned} \langle \nabla q, \nabla p \rangle - \langle q, \mathbf{n} \cdot \nabla p \rangle_{\Gamma} &= \langle q, \nabla \cdot [\rho(\nabla \mathbf{u})\mathbf{u} - 2\nabla^s \mathbf{u} \nabla \mu - \mathbf{b}] \rangle \\ &\quad + \langle q, [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu - \beta \nabla \cdot \mathbf{u} \rangle, \end{aligned} \quad (2.46)$$

where integrating the first term in the right-hand side by parts and inserting the pressure Neumann condition (2.42) yields

$$\begin{aligned} \langle \nabla q, \nabla p \rangle &= \langle \nabla q, 2\nabla^s \mathbf{u} \nabla \mu - \rho(\nabla \mathbf{u})\mathbf{u} + \mathbf{b} \rangle - \langle q\mathbf{n}, \mu \nabla \times (\nabla \times \mathbf{u}) + \rho \partial_t \mathbf{u} \rangle_{\Gamma} \\ &\quad + \langle q, [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu \rangle - \langle q, \beta \nabla \cdot \mathbf{u} \rangle. \end{aligned} \quad (2.47)$$

A key point is to rewrite the boundary term in Eqn. (2.47) via the divergence theorem,

$$\begin{aligned} \langle q\mathbf{n}, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle_{\Gamma} &= \langle q, \mu \nabla \cdot [\nabla \times (\nabla \times \mathbf{u})] \rangle + \langle \nabla(q\mu), \nabla \times (\nabla \times \mathbf{u}) \rangle \\ &= \langle \nabla(q\mu), \nabla \times (\nabla \times \mathbf{u}) \rangle \\ &= \langle \nabla q, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle + \langle q, \nabla \mu \cdot [\nabla \times (\nabla \times \mathbf{u})] \rangle, \end{aligned}$$

and additionally using

$$\langle q\mathbf{n}, \rho \partial_t \mathbf{u} \rangle_{\Gamma} = \langle \nabla q, \rho \partial_t \mathbf{u} \rangle + \langle q, \rho \nabla \cdot (\partial_t \mathbf{u}) \rangle = \langle \nabla q, \rho \partial_t \mathbf{u} \rangle.$$

With these two relations, a term in Eqn. (2.47) cancels and the boundary integral is cast as a domain integral such that one ends up with

$$\langle \nabla q, \nabla p \rangle = \langle \nabla q, 2\nabla^s \mathbf{u} \nabla \mu - \rho[\partial_t \mathbf{u} + (\nabla \mathbf{u})\mathbf{u}] + \mathbf{b} - \mu \nabla \times (\nabla \times \mathbf{u}) \rangle - \langle q, \beta \nabla \cdot \mathbf{u} \rangle.$$

Notice, however, that there are still second-order derivatives of \mathbf{u} present. This is similar to the PSPG formulation we took as a baseline and leads to an incomplete representation of the strong form's residual with piecewise linear finite elements. Contrary to the PSPG method, the term in question,

$$\langle \nabla q, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle$$

may, in fact, be integrated by parts once again, giving

$$\langle \nabla q, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle = \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_{\Gamma} + \langle \nabla \times (\mu \nabla q), \nabla \times \mathbf{u} \rangle,$$

which can be reformulated as

$$\begin{aligned} \langle \nabla \times (\mu \nabla q), \nabla \times \mathbf{u} \rangle &= \langle \mu \nabla \times (\nabla q) + \nabla \mu \times \nabla q, \nabla \times \mathbf{u} \rangle = \langle \nabla \mu \times \nabla q, \nabla \times \mathbf{u} \rangle \\ &\equiv \langle \nabla q, (\nabla \times \mathbf{u}) \times \nabla \mu \rangle \equiv \langle \nabla q, (\nabla \mathbf{u} - \nabla^{\top} \mathbf{u}) \nabla \mu \rangle. \end{aligned}$$

Hence, the viscous stress term involving $\nabla \mu$ in the PPE is

$$\begin{aligned} \langle \nabla q, 2\nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u}) \rangle &= \langle \nabla q, (\nabla \mathbf{u} + \nabla^{\top} \mathbf{u}) \nabla \mu - (\nabla \mathbf{u} - \nabla^{\top} \mathbf{u}) \nabla \mu \rangle \\ &= \langle \nabla q, 2(\nabla^{\top} \mathbf{u}) \nabla \mu \rangle, \end{aligned}$$

resulting finally in a suitable weak form of the PPE with first-order derivatives only

$$\begin{aligned} 0 &= \langle \nabla q, \nabla p + \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - 2(\nabla^{\top} \mathbf{u}) \nabla \mu - \mathbf{b} \rangle \\ &\quad + \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_{\Gamma} + \langle q, \beta \nabla \cdot \mathbf{u} \rangle. \end{aligned} \quad (2.48)$$

As can be seen from Eqn. (2.48), the boundary term involving the vorticity $\nabla \times \mathbf{u}$ alone recovers the full viscous stress in the Newtonian case, which is why we refer to this stabilisation method as boundary vorticity stabilisation (BVS). Note here also that for quasi-Newtonian fluids with inhomogeneous viscosity, $\nabla \mu \neq 0$, only a single term is added and hence a single implementation covering both Newtonian and quasi-Newtonian fluid flow problems is easily designed.

Two issues remain with Eqn. (2.48), namely that (i) $\nabla \mu$, with $\mu = \eta(\dot{\gamma}(\nabla^s \mathbf{u}))$ in general depending on the velocities' first-order derivatives increases regularity requirements on the velocity interpolant \mathbf{u} again, and (ii) that the parameter β is still left to be defined.

We overcome issue (i) borrowing ideas from related works in non-Newtonian fluids [19, 82] and introduce the viscosity as an additional continuous unknown. Here, the rheological law expressed in terms of \mathbf{u} as in Eqn. (2.8) can be utilised to project the fluid viscosity onto a continuous space, $\mu \in H^1(\Omega)$, such that

$$\langle r, \mu \rangle = \langle r, \eta(\dot{\gamma}(\nabla^s \mathbf{u})) \rangle \quad \forall r \in L^2(\Omega), \quad (2.49)$$

which is an L^2 -projection onto a C^0 -continuous space and translates to a mass matrix solve at the discrete level.

Regarding (ii), the choice of β , we compare the stabilised continuity equation of the PSPG formulation as presented in (2.34), that is

$$0 = \langle q_h, \nabla \cdot \mathbf{u}_h \rangle + \sum_{e=1}^{N_e} \langle \tau_e \nabla q_h, \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}_h) \mathbf{u}_h] + \nabla p_h - \mu \Delta \mathbf{u}_h - 2 \nabla^s \mathbf{u}_h \nabla \mu - \mathbf{b} \rangle_{\Omega_e}$$

to the PPE in the BVS formulation given in Eqn. (2.48),

$$\begin{aligned} 0 = & \langle q, \beta \nabla \cdot \mathbf{u} \rangle + \langle \nabla q, \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] + \nabla p - 2 (\nabla^\top \mathbf{u}) \nabla \mu - \mathbf{b} \rangle \\ & + \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_\Gamma, \end{aligned}$$

motivating the simple choice

$$\beta|_{\Omega_e} = 1/\tau_e, \quad (2.50)$$

with the standard stabilisation parameter τ_e as defined in Eqn. (2.35) for the PSPG formulation. With this definition of β , we see two important differences of the BVS scheme compared to other residual-based stabilisations, namely, that (i) the continuity equation is considered in an element-weighted manner and (ii) the stabilising terms can be fully reproduced by piecewise linear, standard Lagrangian finite elements due to the viscosity introduced as an additional continuous unknown.

2.3.2 Momentum equation

To derive a suitable variational formulation of the fluid's linear momentum balance equation, we start from the strong form in Eqn. (2.37)

$$\rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2 \nabla^s \mathbf{u} \nabla \mu = \mathbf{b} - \nabla p,$$

again having χ to switch from the generalised Laplacian ($\chi = 0$) to the stress-divergence form ($\chi = 1$). We aim to find $\mathbf{u} \in [H^1(\Omega)]^d$, $\mathbf{u}|_{\Gamma_D} = \mathbf{g}$, such that there holds

$$\begin{aligned} \langle \mathbf{w}, \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mathbf{b} \rangle + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle - \langle \nabla \cdot \mathbf{w}, p \rangle - \langle \mathbf{w}, \bar{\mathbf{t}} \rangle_{\Gamma_N} = 0, \end{aligned} \quad (2.51)$$

for all $\mathbf{w} \in [H^1(\Omega)]^d$ with $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$ and where $\bar{\mathbf{t}} = \tilde{\mathbf{t}} + \chi \mu (\nabla^\top \mathbf{u}) \mathbf{n}$ again denotes a given Neumann datum being either pseudo tractions $\tilde{\mathbf{t}}$ for $\chi = 0$ or real tractions $\mathbf{t} = \boldsymbol{\sigma} \mathbf{n}$ for $\chi = 1$. The stress-divergence form, identical to the weak form of the momentum equation used in the PSPG method given in Eqn. (2.33), trivially follows from integrating by parts the whole stress term and additionally using the identity

$$\nabla \mathbf{w} : \nabla^s \mathbf{u} \equiv \nabla^s \mathbf{w} : \nabla^s \mathbf{u},$$

while recovering the desired boundary terms in the generalised Laplacian form is more involved. Focusing on the viscous stress term, we have for a test function $\mathbf{w} \in [H^1(\Omega)]^d$ with $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$

$$\begin{aligned} -\langle \mathbf{w}, \mu \Delta \mathbf{u} + 2\nabla^s \mathbf{u} \nabla \mu \rangle &= \langle \nabla \mathbf{u}, \nabla(\mu \mathbf{w}) \rangle - \langle \mathbf{w}, (\mu \nabla \mathbf{u}) \mathbf{n} \rangle_{\Gamma_N} - \langle \mathbf{w}, 2\nabla^s \mathbf{u} \nabla \mu \rangle \\ &= \langle \nabla \mathbf{u}, \mu \nabla \mathbf{w} + \mathbf{w} \otimes \nabla \mu \rangle - \langle \mathbf{w}, (\mu \nabla \mathbf{u}) \mathbf{n} \rangle_{\Gamma_N} - \langle \mathbf{w}, 2\nabla^s \mathbf{u} \nabla \mu \rangle \\ &\equiv \langle \nabla \mathbf{w}, \mu \nabla \mathbf{u} \rangle + \langle \mathbf{w}, \nabla \mathbf{u} \nabla \mu \rangle - \langle \mathbf{w}, (\mu \nabla \mathbf{u}) \mathbf{n} \rangle_{\Gamma_N} - \langle \mathbf{w}, 2\nabla^s \mathbf{u} \nabla \mu \rangle \\ &= \langle \nabla \mathbf{w}, \mu \nabla \mathbf{u} \rangle - \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle - \langle \mathbf{w}, (\mu \nabla \mathbf{u}) \mathbf{n} \rangle_{\Gamma_N}, \end{aligned}$$

which gives together with the partially integrated pressure term pseudo tractions $\tilde{\mathbf{t}}$ on Γ_N and is identical to the viscous stress term in Eqn. (2.51) when setting $\chi = 0$. Now, we removed any second-order derivatives on \mathbf{u} , p and μ , allowing us to finally state the full weak formulation.

2.3.3 Coupled variational formulation

Based on the coupled velocity-pressure system with BVS (2.37)–(2.42), appropriate weak forms as presented in Secs. 2.3.1–2.3.2 are available, lowering the continuity requirements by introducing the viscosity as an additional C^0 -continuous field, but restricting us to a continuous pressure. So, for generalised Newtonian fluids, the problem is to find $(\mathbf{u}, p, \mu) \in [X]^d \times X \times X$, with $X \subset H^1(\Omega)$ and $\mathbf{u}|_{\Gamma_D} = \mathbf{g}$, such that there holds

$$\begin{aligned} \langle \mathbf{w}, \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mathbf{b} \rangle + (1 - \nu/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle - \langle \nabla \cdot \mathbf{w}, p \rangle - \langle \mathbf{w}, \tilde{\mathbf{t}} \rangle_{\Gamma_N} = 0, \end{aligned} \quad (2.52)$$

$$\begin{aligned} \langle \nabla q, \nabla p + \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - 2 (\nabla^\top \mathbf{u}) \nabla \mu - \mathbf{b} \rangle \\ + \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_\Gamma + \langle q, \beta \nabla \cdot \mathbf{u} \rangle = 0, \end{aligned} \quad (2.53)$$

$$\langle r, \mu - \eta (\dot{\gamma} (\nabla^s \mathbf{u})) \rangle = 0, \quad (2.54)$$

for all $(\mathbf{w}, q, r) \in [X]^d \times X \times X$, with initial condition $\mathbf{u} = \mathbf{u}_0$ at $t = 0$, $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$ and $\chi = 0$ for the generalised Laplacian and $\chi = 1$ for the stress divergence form. Then, the traction vector $\tilde{\mathbf{t}} = \tilde{\mathbf{t}} + \chi \mu (\nabla^\top \mathbf{u}) \mathbf{n}$ yields pseudo- or real tractions, i.e., $\tilde{\mathbf{t}}$ or \mathbf{t} , respectively. In the case of Newtonian fluids, terms involving $\nabla \mu \equiv 0$ vanish, also rendering the introduced additional field μ superfluous. Consequently, above weak form reduces to finding $(\mathbf{u}, p) \in [X]^d \times X$, in $X \subset H^1(\Omega)$ with $\mathbf{u}|_{\Gamma_D} = \mathbf{g}$, such that

$$\begin{aligned} \langle \mathbf{w}, \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mathbf{b} \rangle + (1 - \nu/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ - \langle \nabla \cdot \mathbf{w}, p \rangle - \langle \mathbf{w}, \tilde{\mathbf{t}} \rangle_{\Gamma_N} = 0, \end{aligned} \quad (2.55)$$

$$\langle \nabla q, \nabla p + \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mathbf{b} \rangle + \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_\Gamma + \langle q, \beta \nabla \cdot \mathbf{u} \rangle = 0, \quad (2.56)$$

for all $(\mathbf{w}, q) \in [X]^d \times X$, with $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$ and $\mathbf{u} = \mathbf{u}_0$ at $t = 0$.

Both the weak forms are suitable for discretisation in space via standard continuous Lagrangian finite elements, leading towards the derivation of a fully discrete scheme to be discussed in the following section.

2.4 Discretisation in space and time

Discretisation in space follows the standard isoparametric finite element concept which we will briefly introduce here. Therein, the spatial domain Ω is approximated by Ω_h , which is the union of N_e shape-regular, non-overlapping finite elements Ω_e

$$\Omega \approx \Omega_h = \bigcup_{e=1}^{N_e} \Omega_e.$$

These simplex or tensor-product finite elements are mapped from corresponding reference elements, denoted by Ω_{ref} . In these reference elements, Lagrangian basis functions denoted by $N_i(\boldsymbol{\xi})$ with reference coordinate $\boldsymbol{\xi}$ are constructed. The basis functions fulfil the Kronecker- δ property, i.e., they take unit value only in a single node, while they are zero in all other nodes. Then, these reference element functions $N_i(\boldsymbol{\xi})$ are utilised to construct a mapping from the reference element Ω_{ref} to a specific element Ω_e in the real domain Ω_h by

$$\boldsymbol{\chi}_e(\boldsymbol{\xi}) := \sum_{i=1}^{N_{\text{ref}}} N_i(\boldsymbol{\xi}) \boldsymbol{x}_i,$$

where N_{ref} is the number of nodes on the respective reference element and \boldsymbol{x}_i are nodal coordinates of the corresponding real element. The core of the isoparametric concept is that the N_{ref} element functions $N_i(\boldsymbol{\xi})$ are at the same time used to construct spaces of Lagrange polynomials of degree k in the reference elements, denoted by $\mathbb{P}^k(\Omega_{\text{ref}})$ and $\mathbb{Q}^k(\Omega_{\text{ref}})$ for simplex and tensor-product elements, giving rise to the standard C^0 -continuous finite element space $X_h(\Omega_h) \subset H^1(\Omega_h)$ composed of mapped element functions,

$$\begin{aligned} X_h(\Omega_h) &:= \left\{ q_h \in C^0(\overline{\Omega_h}) : q_h \circ \boldsymbol{\chi}_e(\boldsymbol{\xi})|_{\Omega_e} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{P}^k(\Omega_{\text{ref}}) \quad \forall \Omega_e \in \Omega_h \right\}, \\ \text{or } X_h(\Omega_h) &:= \left\{ q_h \in C^0(\overline{\Omega_h}) : q_h \circ \boldsymbol{\chi}_e(\boldsymbol{\xi})|_{\Omega_e} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{Q}^k(\Omega_{\text{ref}}) \quad \forall \Omega_e \in \Omega_h \right\}, \end{aligned} \quad (2.57)$$

for simplex or tensor-product elements, respectively. Important consequences of this choice are that (i) restrictions on the function spaces to incorporate essential boundary conditions are easily imposed by directly setting nodal degrees of freedom, (ii) the node functions composed of all mapped element functions of unit value at that node have a limited support leading to sparse system matrices, and (iii) curved domains Ω_h can be approximated with higher-order precision, directly related to the polynomial degree k .

The choice of reference element and associated Lagrange basis, $\mathbb{P}^k(\Omega_{\text{ref}})$ or $\mathbb{Q}^k(\Omega_{\text{ref}})$, as a central element to the spatial discretisation motivates denoting by P_k or Q_k elements the finite elements resulting from a specific choice. When considering incompressible flow problems, we frequently refer to $Q_m Q_n$ or $P_m P_n$ element pairings to interpolate $(\boldsymbol{u}, p) \in [X_h]^d \times X_h$ using the respective reference element with polynomial degrees m and n for the vector-valued velocity and scalar pressure.

Inserting this finite element discretisation into Eqns. (2.52)–(2.53), we obtain the semidiscrete weak form, seeking for $(\mathbf{u}_h, p_h, \mu_h) \in [X_h]^d \times X_h \times X_h$, with $X_h \subset H^1(\Omega_h)$ (2.57) and $\mathbf{u}_h|_{\Gamma_D} = \mathbf{g}_h$, such that

$$\langle \mathbf{w}_h, \rho \partial_t \mathbf{u}_h \rangle + a(\mathbf{w}_h, \mathbf{u}_h, \mathbf{v}_h, p_h, \mu_h, t) = 0, \quad (2.58)$$

$$\langle \nabla q, \rho \partial_t \mathbf{u}_h \rangle + b(q_h, \mathbf{u}_h, \mathbf{v}_h, p_h, \mu_h, t) + c(q_h, \mathbf{u}_h) = 0, \quad (2.59)$$

$$d(r_h, \mathbf{v}_h, \mu_h) = 0, \quad (2.60)$$

for all $(\mathbf{w}_h, q_h, r_h) \in [X_h]^d \times X_h \times X_h$, with $\mathbf{w}_h|_{\Gamma_D} = \mathbf{0}$, introducing shorthand notations for terms including data evaluated at time $t = t^i$, all terms only involving spatial operators and splitting off the element-wise contribution to the incompressibility constraint,

$$\begin{aligned} a(\mathbf{w}, \mathbf{u}, \mathbf{v}, p, \mu, t^i) &:= \langle \mathbf{w}, \rho(\nabla \mathbf{u}) \mathbf{v} - \mathbf{b}|_{t=t^i} \rangle + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ &\quad - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle - \langle \nabla \cdot \mathbf{w}, p \rangle - \langle \mathbf{w}, \bar{\mathbf{t}}|_{t=t^i} \rangle_{\Gamma_N}, \end{aligned} \quad (2.61)$$

$$\begin{aligned} b(q, \mathbf{u}, \mathbf{v}, p, \mu, t^i) &:= \langle \nabla q, \nabla p + \rho(\nabla \mathbf{u}) \mathbf{v} - 2(\nabla^\top \mathbf{u}) \nabla \mu - \mathbf{b}|_{t=t^i} \rangle \\ &\quad + \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_{\Gamma}, \end{aligned} \quad (2.62)$$

$$c(q, \mathbf{u}) := \sum_{e=1}^{N_e} \langle q, \frac{1}{\tau_e} \nabla \cdot \mathbf{u} \rangle_{\Omega_e}, \quad (2.63)$$

$$d(r, \mathbf{v}, \mu) := \langle r, \mu - \eta(\dot{\gamma}(\nabla^s \mathbf{v})) \rangle, \quad (2.64)$$

with τ_e as defined for the PSPG method in Eqn. (2.35). Above, we additionally introduced a velocity vector \mathbf{v} entering the convective term and the rheological law with $\mathbf{v} = \mathbf{u}$ being a straight-forward choice, resulting in a nonlinear problem with coupled viscosity and velocity fields.

Similar to the standard PSPG formulation, we may combine the PPE-based inf-sup stabilisation with other techniques, e.g., to further improve mass conservation via grad-div stabilisation [76, 77] or to counteract instabilities due to dominant convection by SUPG [73, 74] or artificial diffusion [75] methods. Some of the residual-based stabilisation techniques such as the SUPG method might in fact suffer from an incomplete residual as well, but given the flow regime in which they are actually needed, the drawbacks from using an “inviscid” residual are most often negligible. That being said, a transition region might obviously exist, such that some practical applications can lead to configurations where areas with dominant diffusion are present and at the same time convective effects dominate in other parts of the domain. Adding SUPG stabilisation in the latter regions (automatically scaled by an appropriate choice of τ_e) can suppress convective instabilities even in such scenarios, accepting the incomplete residual in the lower-order case, knowing it is of lesser importance. Then, still employing the PPE-based stabilisation can increase accuracy by eliminating spurious pressure boundary layers in zones with dominant diffusion.

For now, we restrict ourselves to low Reynolds numbers, meaning the diffusion-dominant regime and additionally consider grad-div stabilisation [76, 77], even though it is not strictly necessary. The penalty term lending its name to this stabilisation technique is derived from

$$-\langle \mathbf{w}, \nabla(\nabla \cdot \mathbf{u}) \rangle_{\Omega_e} = \langle \nabla \cdot \mathbf{w}, \nabla \cdot \mathbf{u} \rangle_{\Omega_e} + \langle \mathbf{w}, \mathbf{n} \nabla \cdot \mathbf{u} \rangle_{\partial \Omega_e} = \langle \nabla \cdot \mathbf{w}, \nabla \cdot \mathbf{u} \rangle_{\Omega_e},$$

where the boundary term vanishes due to $\nabla \cdot \mathbf{u} = 0$. A corresponding stabilisation is commonly used for LBB-stable and especially for stabilised formulations to enhance mass conservation by adding [27]

$$e(\mathbf{w}_h, \mathbf{u}_h^{n+1}) := \sum_{e=1}^{N_e} \langle \gamma_e \nabla \cdot \mathbf{w}_h, \nabla \cdot \mathbf{u}_h^{n+1} \rangle_{\Omega_e}, \quad (2.65)$$

to the momentum equation (2.58), where the stabilisation parameter γ_e is given as [27]

$$\gamma_e := \begin{cases} h_e/2 & \text{or } \mathcal{O}(h_e) & \text{for equal-order pairs,} \\ 1/10 & \text{or } \mathcal{O}(1) & \text{for inf-sup stable pairs.} \end{cases} \quad (2.66)$$

Discretising in time, we employ the generalised θ -scheme, dividing the time interval of interest $I_t = (0, T]$ from start time $t = 0$ to $t = T$ into N_t steps of possibly non-uniform length $\Delta t^n = t^{n+1} - t^n$, $n = 0, \dots, N_t$ and approximate the time derivative by

$$\partial_t \mathbf{u}|_{t=\theta t^{n+1} + \theta' t^n} \approx \frac{1}{\Delta t^n} (\mathbf{u}^{n+1} - \mathbf{u}^n),$$

with time integration parameters θ and $\theta' := 1 - \theta$, and introducing superscripts indicating time t^i , i.e., $\mathbf{u}^i := \mathbf{u}|_{t=t^i}$ for the velocity, $p^i := p|_{t=t^i}$ for the pressure and likewise for all other fields, data or other quantities to be introduced later.

With this, we finally arrive at the fully discrete problem of finding $(\mathbf{u}_h^{n+1}, p_h^{n+1}, \mu_h^{n+1}) \in [X_h]^d \times X_h \times X_h$ with X_h (2.57), given data \mathbf{u}_h^n and p_h^n from previous time steps and $\mathbf{u}_h^{n+1}|_{\Gamma_D} = \mathbf{g}_h^{n+1}$, such that

$$\begin{aligned} \langle \mathbf{w}_h, \frac{\rho}{\Delta t^n} (\mathbf{u}_h^{n+1} - \mathbf{u}_h^n) \rangle + \theta a(\mathbf{w}_h, \mathbf{u}_h^{n+1}, \mathbf{v}_h^{n+1}, p_h^{n+1}, \mu_h^{n+1}, t^{n+1}) \\ + e(\mathbf{w}_h, \mathbf{u}_h^{n+1}) + \theta' a(\mathbf{w}_h, \mathbf{u}_h^n, \mathbf{v}_h^n, p_h^n, \mu_h^n, t^n) = 0, \end{aligned} \quad (2.67)$$

$$\begin{aligned} \langle \nabla q, \frac{\rho}{\Delta t^n} (\mathbf{u}_h^{n+1} - \mathbf{u}_h^n) \rangle + \theta b(q_h, \mathbf{u}_h^{n+1}, \mathbf{v}_h^{n+1}, p_h^{n+1}, \mu_h^{n+1}, t^{n+1}) \\ + \theta' b(q_h, \mathbf{u}_h^n, \mathbf{v}_h^n, p_h^n, \mu_h^n, t^n) + c(q_h, \mathbf{u}_h^{n+1}) = 0, \end{aligned} \quad (2.68)$$

$$d(r_h, \mathbf{v}_h^{n+1}, \mu_h^{n+1}) = 0, \quad (2.69)$$

for all $(\mathbf{w}_h, q_h, r_h) \in [X_h]^d \times X_h \times X_h$, where $\mathbf{w}_h|_{\Gamma_D} = \mathbf{0}$.

In Eqns. (2.67)–(2.69), the time integration parameters θ and θ' can be chosen to retrieve specific single-step time integration methods. Fixing, e.g., $\theta = 1$ and $\theta' = 0$

yields the implicit Euler timestepping scheme, while $\theta = \theta' = 1/2$ results in the Crank–Nicolson (CN) scheme. A consistent time integration of the pressure term as in so-called pressure-corrected θ -schemes [83], however, requires an initial pressure field $p|_{t=t^0} = p^0$ since $\theta' \neq 0$. Several strategies to overcome this complication have been presented. The initial pressure field is either assumed given, obtained from a preceding ramp-up phase from the quiescent state, recovered solving a suitable PPE with \mathbf{u}_0 , or—as most frequently done—simply computed using a single (small) initial implicit Euler step, which does not require p^0 at all.

Note here also that the element-weighted incompressibility constraint in the PPE (2.68) is enforcing a divergence-free \mathbf{u}_h^{n+1} . This is equivalent to integrating the constraint in time starting from an initially divergence-free velocity field, but does not accumulate errors in mass conservation over time [84, 85]. In analogy, we base the rheological law in Eqn. (2.69) on the auxiliary velocity field evaluated at t^{n+1} , \mathbf{v}_h^{n+1} . This auxiliary velocity field \mathbf{v}_h^{n+1} determines, whether or not the viscosity and velocity are implicitly coupled in a nonlinear fashion and if the convective term introduces the usual nonlinearity as shall be explained in more detail in the following section.

2.5 Nonlinear solver and preconditioning

The discrete spaces X_h are constructed in a standard manner using vector- or scalar-valued, C^0 -continuous nodal basis functions φ^u , φ^p and φ^μ for the respective physical fields. Then, the fully discrete weak form of the BVS scheme given in Eqns. (2.67)–(2.69) leads to (non-)linear systems of equations, depending on the choice of \mathbf{v}_h^{n+1} .

When choosing $\mathbf{v}_h^{n+1} = \mathbf{u}_h^{n+1}$, a nonlinear solution procedure has to be applied, with popular choices in this regard being Newton’s method or a fixed-point iteration, the so-called Picard scheme. Both of these nonlinear solvers require several iterations per time step given a reasonable initial guess and convergence criterion. Herein, we settle for the latter option, because (i) Newton’s method introduces additional coupling terms, which hinder the decoupling of the equation governing viscosity, and (ii) Picard and Newton’s method have been shown to be equally efficient for time-dependent three-dimensional incompressible flow problems in [86].

Hence, to compute the new iterates, i.e., the discrete coefficient vectors $\underline{\mathbf{u}}^{k+1}$, $\underline{\mathbf{p}}^{k+1}$ and $\underline{\mu}^{k+1}$, based on the given initial guess or past iterates, $\underline{\mathbf{u}}^k$, $\underline{\mathbf{p}}^k$ and $\underline{\mu}^k$, a single step of the Picard scheme consists of solving the linearised 3×3 block system

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} & \mathbf{0} \\ \mathbf{C} & \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{M} \end{pmatrix} \begin{pmatrix} \underline{\mathbf{u}}^{k+1} \\ \underline{\mathbf{p}}^{k+1} \\ \underline{\tilde{\mu}} \end{pmatrix} = \begin{pmatrix} \underline{\mathbf{f}} \\ \underline{\mathbf{g}} \\ \underline{\mathbf{h}} \end{pmatrix}, \quad (2.70)$$

where we introduce the intermediate variable $\underline{\tilde{\mu}} := 1/\tilde{\eta} \underline{\mu}^{k+1}$ for the viscosity scaled by $\tilde{\eta} = \eta_\infty$. As can easily be seen, due to the linearisation $\eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^{n+1})) \approx \eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^k))$,

the systems decouple into a projection of the viscosity and a velocity-pressure subsystem with the matrices and vectors defined as

$$[\mathbf{A}]_{ij} := \rho \langle \boldsymbol{\varphi}_i^u, \boldsymbol{\varphi}_j^u + \tilde{\theta} (\nabla \boldsymbol{\varphi}_j^u) \mathbf{u}_h^* \rangle + \tilde{\theta} (1 - \chi/2) \langle \nabla \boldsymbol{\varphi}_i^u + \chi \nabla^\top \boldsymbol{\varphi}_i^u, \mu_h^* \nabla \boldsymbol{\varphi}_j^u + \chi \mu_h^* \nabla^\top \boldsymbol{\varphi}_j^u \rangle \\ - \tilde{\theta} (1 - \chi) \langle \boldsymbol{\varphi}_i^u, (\nabla^\top \boldsymbol{\varphi}_j^u) \nabla \mu_h^* \rangle + \Delta t^n e(\boldsymbol{\varphi}_i^u, \boldsymbol{\varphi}_j^u), \quad (2.71)$$

$$[\mathbf{B}]_{ij} := -\tilde{\theta} \langle \nabla \cdot \boldsymbol{\varphi}_i^p, \boldsymbol{\varphi}_j^p \rangle, \quad [\mathbf{D}]_{ij} := \tilde{\tau}_i \tilde{\theta} \langle \nabla \boldsymbol{\varphi}_i^p, \nabla \boldsymbol{\varphi}_j^p \rangle, \quad [\mathbf{M}]_{ij} := \langle \boldsymbol{\varphi}_i^\mu, \boldsymbol{\varphi}_j^\mu \rangle, \quad (2.72)$$

$$[\mathbf{C}]_{ij} := \tilde{\tau}_i \langle \nabla \boldsymbol{\varphi}_i^p, \rho [\boldsymbol{\varphi}_i^u + \tilde{\theta} (\nabla \boldsymbol{\varphi}_j^u) \mathbf{u}_h^*] \rangle - 2\tilde{\theta} \langle \nabla^\top \boldsymbol{\varphi}_i^u, \nabla \mu_h^* \rangle \\ + \tilde{\tau}_i \tilde{\theta} \langle \nabla \boldsymbol{\varphi}_i^p \times \mathbf{n}, \mu_h^* \nabla \times \boldsymbol{\varphi}_j^u \rangle_\Gamma + \tilde{\tau}_i \Delta t^n c(\boldsymbol{\varphi}_i^p, \boldsymbol{\varphi}_j^u), \quad (2.73)$$

$$[\underline{\mathbf{f}}]_i := \langle \boldsymbol{\varphi}_i^u, \rho \mathbf{u}^n \rangle - \tilde{\theta}' a(\boldsymbol{\varphi}_i^u, \mathbf{u}_h^n, \mathbf{u}_h^n, p_h^n, \mu_h^n, t^n) + \tilde{\theta} \langle \boldsymbol{\varphi}_i^u, \mathbf{b}^{n+1} \rangle + \tilde{\theta} \langle \boldsymbol{\varphi}_i^u, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N}, \quad (2.74)$$

$$[\underline{\mathbf{g}}]_i := \tilde{\tau}_i \langle \boldsymbol{\varphi}_i^p, \rho \mathbf{u}_h^n \rangle + \tilde{\tau}_i \tilde{\theta} \langle \nabla \boldsymbol{\varphi}_i^p, \mathbf{b}^{n+1} \rangle - \tilde{\tau}_i \tilde{\theta}' b(\boldsymbol{\varphi}_i^p, \mathbf{u}_h^n, \mathbf{u}_h^n, p_h^n, \mu_h^n, t^n), \quad (2.75)$$

$$[\underline{\mathbf{h}}]_i := 1/\tilde{\eta} \langle \boldsymbol{\varphi}_i^\mu, \eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^k)) \rangle \quad (2.76)$$

with $\mathbf{u}_h^* := \mathbf{u}_h^k$ and $\mu_h^* := \mu_h^k$, slightly adapted time integration parameters $\tilde{\theta} := \Delta t^n \theta$ and $\tilde{\theta}' := \Delta t^n \theta'$ and a scaling parameter $\tilde{\tau}_i$ for pressure test function $\boldsymbol{\varphi}_i^p$, computed up front as the volumetric average stabilisation parameter of elements adjacent to vertex i with stabilisation parameters τ_e , which effectively rescales the equations to counteract numerical difficulties arising for $\tau_e \rightarrow 0$ as $h_e \rightarrow 0$.

Note here that the zero blocks in the momentum equation, PPE and also the viscosity-velocity block on the lower left of the 3×3 block system (2.70) would be filled if we were to apply Newton's method as mentioned above, which is a strong reason for choosing Picard's method. But even when employing a Picard scheme, linearising the viscosity is only one of a few options. However, it is the single and most natural one introducing zeros in those blocks mentioned. On the other hand, it may not be possible to factor out linear coefficients of the discrete velocity vector from $\eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^{n+1}))$ depending on the considered constitutive equation, which might complicate incorporating other rheological laws and is hence avoided within this work.

System (2.70) thus naturally decouples, allowing for a sequential solve per nonlinear step, starting from either the viscosity projection or velocity-pressure step. We choose to start with the viscosity step, computing μ^{k+1} from \mathbf{u}^k ,

$$\mathbf{M} \underline{\mu} = \underline{\mathbf{h}} \quad \text{and} \quad \underline{\mu}^{k+1} = \tilde{\eta} \underline{\mu}, \quad (2.77)$$

which is the single place where the rheological law enters, rendering adaptations in this regard effortless. Once the viscosity is updated, the just computed $\underline{\mu}^{k+1}$ can be inserted into the velocity-pressure system, i.e., set $\mu_h^* = \mu_h^{k+1}$ and $\mathbf{u}_h^* = \mathbf{u}_h^k$ and solve

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \underline{\mathbf{u}}^{k+1} \\ \underline{\mathbf{p}}^{k+1} \end{pmatrix} = \begin{pmatrix} \underline{\mathbf{f}} \\ \underline{\mathbf{g}} \end{pmatrix} \quad (2.78)$$

for \mathbf{u}_h^{k+1} and p_h^{k+1} until relative or absolute stopping criteria,

$$\left\| \underline{\mathbf{x}}^{k+1} - \underline{\mathbf{x}}^k \right\| \leq \epsilon_{\text{abs}}^{\text{nl}} \quad \text{and} \quad \left\| \underline{\mathbf{x}}^{k+1} - \underline{\mathbf{x}}^k \right\| \leq \epsilon_{\text{rel}}^{\text{nl}} \left\| \underline{\mathbf{x}}^{k+1} \right\|, \quad (2.79)$$

on the combined solution vector of the velocity-pressure system (2.78) are fulfilled. Inserting the most recent viscosity update corresponds to a Gauss-Seidel elimination, whereas alternatively, a Jacobi-iteration might be considered, inserting the old iterate $\mu_h^* = \mu_h^k$ for linearisation. Interestingly, the latter option enables parallel execution, but was not considered, as one expects lower iteration counts using a Gauss-Seidel approach and the viscosity projection step is much cheaper than the velocity-pressure system solve or might even be lumped, rendering its inversion trivial.

To improve the nonlinear solver's convergence behaviour, we apply Aitken's acceleration [87–89]. This relaxation scheme consists of modifying the two velocity-pressure and viscosity subproblem's solutions via

$$\underline{\mathbf{y}}^{k+1} = \omega_k \tilde{\mathbf{y}}^k + (1 - \omega_k) \underline{\mathbf{y}}^k, \quad (2.80)$$

referring to the results after a full solve of (2.70) in nonlinear iteration k by $\tilde{\mathbf{y}}^k$ for both subsystems and with a scalar parameter ω_k , which is recursively defined as [88]

$$\omega_k = -\omega_{k-1} \frac{\underline{\mathbf{r}}^{k-1} \cdot (\underline{\mathbf{r}}^k - \underline{\mathbf{r}}^{k-1})}{\left\| \underline{\mathbf{r}}^k - \underline{\mathbf{r}}^{k-1} \right\|^2} \quad \text{with} \quad \underline{\mathbf{r}}^k = \tilde{\mathbf{y}}^k - \underline{\mathbf{y}}^{k-1}, \quad (2.81)$$

and an initial relaxation parameter ω_0 to be defined.

This fixed-point iteration scheme converges within a few steps, depending on the tolerances set. A standard approach would be to start from the quiescent state and ramping up the solution, which can result in non-convergence, if only the relative stopping criterion were used. Thus, the absolute criterion is effective, if the solution's norm itself is rather small compared to the tolerances in the linear solvers applied. Setting $\epsilon_{\text{abs}}^{\text{nl}}$ accordingly can prevent stalling of the nonlinear solver, and still result in a negligible absolute error $\left\| \underline{\mathbf{x}}^{k+1} - \underline{\mathbf{x}}^k \right\|$, which can become relevant in (pseudo) time-dependent problems using rather small time steps.

However, having a good initial guess, possibly extrapolated from previous (pseudo) time steps, performing just a single iteration of the nonlinear solution scheme might be enough. This is clearly the case for time-independent problems, since we have greater freedom when applying pseudo-timestepping or solving one single nonlinear problem. For in-stationary problems, on the other hand, the construction of an initial guess remains a crucial aspect in order to not spoil temporal accuracy and stability.

In the time-dependent setting, we thus propose a much simpler alternative to solving nonlinear problems (2.70): extrapolate \mathbf{u}_h from known time step data with coefficients

Tab. 2.1: Coefficients β_j^m for extrapolation of order $m = 1, 2$ [90].

j	1	2
$m = 2$	$1 + \frac{\Delta t^n}{\Delta t^{n-1}}$	$\frac{\Delta t^n}{\Delta t^{n-1}}$
$m = 1$	1	—

given in Tab. 2.1, depending on the extrapolation order m in time, via

$$\mathbf{u}_h^{n+1} \approx \mathbf{u}_h^* := \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{u}_h^{n-j} \quad (2.82)$$

This effectively uses \mathbf{u}_h^* as an initial guess for the nonlinear scheme, or—when performing only a single iteration—completely linearises and decouples the overall system (2.70). In the latter case, this leads to a number of simplifications, and has also been shown to be a good compromise between speed, accuracy and stability in the context of Newtonian fluids [91–93] and buoyancy-driven flow problems [94]. Moreover, recovering the viscosity via $\eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^*))$ has two distinct advantages over using an implicit approach: (i) only a simple projection decoupled from the velocity-pressure system and involving only a mass matrix solve is needed, and (ii) exchanging or modifying the rheological law is rendered trivial, since it boils down to adjusting the right-hand side of the projection. Both these arguments directly translate from the discussion regarding the Picard scheme and even lead to an identical sequence of linear subsystems (2.77)–(2.78) to be solved, simply employing only a single iteration of the fixed-point method, while starting from a higher-order accurate initial guess.

Then, the timestepping algorithm including an initialisation phase reads

-
1. *Viscosity initialisation:* Given the initial velocity field \mathbf{u}_h^0 , compute the viscosity at t^0 and t^1 , $\mu_h^0 = \mu_h^1$ by solving Eqn. (2.77) with \mathbf{u}_h^0 , which is identical to a first order extrapolation $\mathbf{u}_h^* = \mathbf{u}_h^0$.
 2. *Implicit Euler step:* Solve the velocity-pressure system (2.78) for \mathbf{u}_h^1 and p_h^1 with $\theta = 1$ and $\theta' = 0$ and a first order extrapolation $\mathbf{u}_h^* = \mathbf{u}_h^0$.
 3. *Timestepping:*
FOR $n = 1, \dots, N_t$
 - a) Extrapolate the velocity based on \mathbf{u}_h^n and \mathbf{u}_h^{n-1} via Eqn. (2.82), giving \mathbf{u}_h^* .
 - b) Compute μ_h^{n+1} via projection (2.77) with \mathbf{u}_h^* .
 - c) Solve the linearised velocity-pressure system (2.78) and time integration parameters θ and θ' as desired to obtain \mathbf{u}_h^{n+1} and p_h^{n+1} .

END FOR

The linear algebraic systems corresponding to the velocity-pressure subsystem (2.78) and viscosity projection step (2.77) are solved in each iteration of the fixed-point scheme

or once per time step for the linearised, fully decoupled approach. Naturally, iterative linear solvers with suitable preconditioners are employed for each solve, where we heavily rely on algebraic multigrid (AMG) methods provided by PETSc Hypre’s BoomerAMG package [95], easily accessible within the finite element toolkit `deal.II` (see, e.g., [96–99]). For the projection to recover the viscosity from velocity, which is a mass matrix problem, a conjugate gradient method (CG) [100] is found effective. Using a single V-cycle of an AMG method with Chebyshev smoother as preconditioner might already be considered an overkill, but is chosen nonetheless for its robustness. As an alternative, we might also lump the mass matrix entirely, and invert the resulting diagonal matrix just performing a vector scaling as will be discussed in Sec. 2.7.

Solving the velocity-pressure subsystem is known to be delicate—a typical standpoint might be (see, e.g. [92, 101]) that coupled solvers treating the velocity-pressure system as a monolithic one are preferred for stationary problems, while projection methods decoupling velocity and pressure unknowns shine in time dependent flow problems [82]. However, the downside of treating decoupled systems being at the core of any projection scheme or similar methods is that the maximum stable time step size is reduced. Depending on the problem at hand, this further restriction might not be critical, e.g., when the desired time integration error dictates the chosen time step size compared to the Courant-Friedrichs-Lewy (CFL) condition, but *can* be.

However, tackling the coupled velocity-pressure system is not an easy task. A standard and more traditional approach might be to use an incomplete LU-factorisation combined with Cuthill-McKee [102] or King [103] reordering algorithms for minimal bandwidth, which is in our case applicable given the non-zero pressure-pressure block. But during the past decades, faster and more robust algorithms have been developed, which are particularly performant in instationary flow problems (see, e.g. [27, 104–108]). Combinations of Krylov subspace solvers (see, e.g., [107, 109, 110]) with geometric or algebraic multigrid methods on individual blocks in physics-based preconditioners are among the best options available up to date. Multigrid methods as scalable solvers and preconditioners are attractive choices, while the algebraic counterpart is entirely based on the discrete matrices and hence easily applicable through open-source scientific software such as BoomerAMG [95] or Trilinos’ ML package [111]. Thus, we focus on the class of Schur-complement-based, so-called physics-based, preconditioners, where the key to an effective preconditioner lies in approximating the inverse of the pressure Schur complement sufficiently. This has been and is still a very active area of research, leading to numerous articles [112–117], reviews [105, 106, 118, 119], and books [107, 109, 110] we recommend to the interested reader for a rigorous introduction and more in-depth discussion.

Already in the 1970’s, Patankar and Spalding [120] introduced the so-called SIMPLE method, which was used as a preconditioner by Vuik et al. [121] and later further developed in [122]. With this method based on the diagonal of the velocity-velocity block, $\text{diag}(\mathbf{A})$, good results can be achieved for instationary problems and/or flows in the diffusive regime, but performance falls off for larger time steps, dominant convection or when elements with high aspect ratio are present. To enrich the approximation of the inverse pressure Schur complement, Silvester et al. [123] and Kay et al. [124] designed the pressure convection-diffusion (PCD) preconditioner building on an operator acting on the pressure space analogously to the velocity operator. While this idea is directly

applicable to both stable and stabilised incompressible flow problems, its purely algebraic counterpart, the least-squares commutator preconditioner [125] was later extended in that regard [126].

Another promising concept roots in consistent manipulation of the block system. The aim is to ease the approximation of the pressure Schur complement considerably, which unfortunately comes at the cost of increasing the numerical effort spent in solving the velocity-velocity block. Impressively, these so-called augmented-Lagrangian-based approaches lead to low iteration counts (almost) entirely independent of the spatial discretisation or present flow regime, including high-Reynolds problems [127]. For stabilised formulations, such schemes were extended in Benzi et al. [128], while adding a standard grad-div term (2.65) with suitable parameter yields similar results for inf-sup stable finite element pairs as shown by Heister and Rapin [129], and yet does not rely on large matrix-matrix products.

Studies comparing these different options in terms of computational cost and robustness revealed that the augmented-Lagrangian preconditioner is in stationary high-Reynolds flow clearly superior to least-squares commutator and PCD alternatives, which themselves showed to be more robust than SIMPLE and related methods [119, 130, 131]. Such a statement does not hold for instationary problems, though, since a strong reaction term scaling with the time step size helps with performance issues even for stretched grids and dominant convection. If we have the CFL condition or a targeted time integration error enforcing a small enough time step size, all of the mentioned methods yield (at least to some extent) satisfactory results. And while the just mentioned ranking is not as strict in the time-dependent case, it does still hold. In the context of practical applications, important aspects besides the achieved iteration counts and robustness such as setup and execution costs, ease of implementation and parallel scalability gain additional weight.

Since the focus herein lies on (highly) variable viscosity in flows of generalised Newtonian fluids, concepts already extended in this regard, e.g., [132–134], are favourable candidates. Borrowing these ideas, slightly modified versions of well-established preconditioners are then directly applicable in the present context. So, within this work we propose to combine some of the aforementioned methods and start from the flexible generalised minimal residual method (FGMRES) [135] as the iterative linear solver. For the velocity-pressure system as given in Eqn. (2.78),

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \underline{\mathbf{u}}^{k+1} \\ \underline{\mathbf{p}}^{k+1} \end{pmatrix} = \begin{pmatrix} \underline{\mathbf{f}} \\ \underline{\mathbf{g}} \end{pmatrix},$$

a physics-based right preconditioner \mathcal{P}^{-1} is then defined as [106]

$$\mathcal{P}^{-1} := \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{S} \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & -\mathbf{B} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{pmatrix}, \quad (2.83)$$

with the pressure Schur complement

$$\mathbf{S} := \mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}, \quad (2.84)$$

where one can easily verify

$$\begin{aligned} & \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & -\mathbf{B} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{pmatrix} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & -\mathbf{B}\mathbf{S}^{-1} \\ \mathbf{0} & \mathbf{S}^{-1} \end{pmatrix} \\ & = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{B}\mathbf{S}^{-1} \\ \mathbf{0} & \mathbf{S}^{-1} \end{pmatrix} = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{C}\mathbf{A}^{-1} & (\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})\mathbf{S}^{-1} \end{pmatrix} = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{C}\mathbf{A}^{-1} & \mathbf{I} \end{pmatrix}, \end{aligned}$$

meaning that exact application of \mathcal{P}^{-1} leads to a triangular block system with ones on the main diagonal, such that the outer Krylov method would converge in at most two iterations [106]. Computing \mathcal{P}^{-1} or rather its application to a vector in an exact fashion, however, is unfeasible since we do not have a matrix representation of \mathbf{S} available. The critical point here is that constructing the exact Schur complement $\mathbf{S} := \mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$ is costly, given that one would need to compute the exact inverse of the full velocity-velocity block matrix \mathbf{A} and additionally perform a triple matrix product, which is in general much denser than \mathbf{D} . Luckily, the action of the inverses \mathbf{A}^{-1} and \mathbf{S}^{-1} within the preconditioner can be approximated when applying \mathcal{P}^{-1} on a given vector within the FGMRES iterations. To this end, we employ a single AMG V-cycle to account for \mathbf{A}^{-1} , whereas designing a good approximation of \mathbf{S}^{-1} is more involved.

Following Turek [104], the velocity-velocity block is split into its components,

$$\mathbf{A} = \rho\mathbf{M}_u + \mathbf{N}_u(\mathbf{u}_h^*, \nabla\mu_h^*) + \mathbf{L}_u(\mu_h^*, \gamma_e),$$

with the velocity mass matrix \mathbf{M}_u , \mathbf{N}_u containing the convective and transposed velocity gradient times viscosity gradient terms and the diffusive contributions plus the grad-div stabilisation gathered in \mathbf{L} . Then, each of these components is accounted for individually, i.e., we consider

$$\mathbf{S}^{-1} \approx \left(\mathbf{D} - \mathbf{C} \operatorname{diag}(\rho\mathbf{M}_u)^{-1} \mathbf{B} \right)^{-1} - \mathbf{M}_p^{-1} \mathbf{F}_p \mathbf{L}_p^{-1} + \mathbf{M}_{\mu,\gamma}^{-1}, \quad (2.85)$$

using the PCD preconditioner by Elman et al. [107], Kay et al. [124] for the convective term and incorporate ideas from [129, 134] to account for variable viscosity and grad-div terms.

The first part involving reaction and stabilisation terms is explicitly computed, which automatically accounts for correct boundary conditions and stabilisation terms and was found to perform better in the small time step limit than an alternative strategy involving approximate inversion of the pressure Laplace matrix \mathbf{L}_p . The pressure matrices are defined as in the Newtonian setting,

$$\begin{aligned} & [\mathbf{M}_p]_{ij} := \langle \varphi_i^p, \varphi_j^p \rangle, \quad [\mathbf{L}_p]_{ij} := \langle \nabla\varphi_i^p, \nabla\varphi_j^p \rangle, \\ \text{and} \quad & [\mathbf{F}_p]_{ij} := \rho \langle \varphi_i^p, \varphi_j^p + \tilde{\theta} \nabla\varphi_j^p \cdot \mathbf{u}_h^* \rangle + \langle \nabla\varphi_i^p, \tilde{\theta} \mu_h^* \nabla\varphi_j^p \rangle, \end{aligned} \quad (2.86)$$

since no notable differences were seen in numerical tests when using a straight forward extension adopting operators corresponding to the specific formulation targeting generalised Newtonian fluids:

$$[\tilde{\mathbf{F}}_p]_{ij} := \langle \varphi_i^p, \rho \varphi_j^p + \tilde{\theta} \nabla \varphi_j^p \cdot (\rho \mathbf{u}_h^* - \nabla \mu_h^*) \rangle + \tilde{\theta} (1 + \chi)^2 \langle \nabla \varphi_i^p, \mu_h^* \nabla \varphi_j^p \rangle, \quad (2.87)$$

such that the standard Newtonian variant is considered from here on. Regarding boundary conditions for these operators, we use

$$p_h = 0 \quad \text{on } \Gamma_N, \quad \text{and} \quad \mathbf{n} \cdot \nabla p_h = 0 \quad \text{on } \Gamma_D,$$

when assembling the pressure Laplace matrix \mathbf{L}_p , while \mathbf{F}_p incorporates Robin boundary conditions of the form [136]

$$-\mu_h^* \mathbf{n} \cdot \nabla p_h + (\mathbf{u}_h^* \cdot \mathbf{n}) p_h = 0 \quad \text{on } \Gamma. \quad (2.88)$$

To account for potentially highly variable viscosity, we take inspiration from [129, 132] and include scaling terms in $\mathbf{M}_{\mu, \gamma}$, which leads to

$$[\mathbf{M}_{\mu, \gamma}]_{ij} := \langle \varphi_i^p, (\mu_h^* + \Delta t^n \gamma_e)^{-1} \varphi_j^p \rangle. \quad (2.89)$$

Comparing matrices as obtained from a standard stabilised discretisation with the present ones, the only terms not vanishing as $\Delta t^n \rightarrow 0$ are found in the pressure-velocity block, \mathbf{C} , where instead of only having $\langle q, \nabla \cdot \mathbf{u} \rangle$, we have instead

$$\rho / \Delta t^n \langle \nabla q, \mathbf{u} \rangle + 1 / \tau_e \langle q, \nabla \cdot \mathbf{u} \rangle,$$

which are in the worst case of equal magnitude. Apparently, one might rewrite the first part as a boundary integral again, but since this peculiarity was also not found to be of much influence, we keep the form as is.

With this, all necessary parts of the preconditioner are defined and all inverses in the Schur complement approximation can be approximated, e.g., by AMG V-cycles. Altogether, a single application of \mathcal{P}^{-1} on a vector thus uses four multigrid cycles on pressure matrices and one on the velocity-velocity block and several matrix-vector products.

2.6 Timestepping schemes

Based on the specific choice of time integration parameters θ and θ' in the generalised θ -scheme, different methods result for advancing the coupled problem in time. Prescribing a fixed uniform time step size for the whole solution process based on accuracy or stability requirements at a single time instant or a subinterval may result in a tremendous waste of resources. But luckily, we are not restricted to a uniform time step size, and in practice, choosing an appropriate time step Δt^n depending on the solution allows capturing rapid changes while keeping the integration error low and improving stability. Thus, adaptive timestepping is a vital ingredient of any effective solver [137–140]. In addition to that,

one might decrease step sizes resulting in a better initial guess for linear and/or nonlinear solvers, or one might even get away with linearising the governing equations in the first place.

To adapt the time step size based on the time integration error, we resort to two different methods building on the generalised θ -scheme: one is a stabilised predictor-corrector approach combining Crank–Nicolson (CN) or alternatively Rannacher (R) timestepping with an explicit Adams–Bashforth method of second order (AB2), while the other is a pressure-corrected fractional-step θ -scheme (FS) with an embedded method of first order.

2.6.1 Predictor-corrector approach

The CN time integration scheme is a particularly popular method, as it is easy to implement, second-order accurate and A-stable [141]. Its non-dissipative nature is an attractive feature when pure convection or convection-dominated problems are considered. Unfortunately though, applying the CN scheme with (too large) fixed time steps can be problematic as thoroughly investigated in literature, including works critically comparing CN to other available timestepping algorithms or operator-splitting methods [101, 142]. Various remedies have been presented to consistently introduce subtle numerical dissipation, counteracting instabilities triggered by rough initial or boundary data. Strategies like introducing intermediate implicit Euler steps [143, 144], periodic averaging [138, 139, 145, 146] or shifting the parameter θ slightly to the implicit side [141] are among the available remedies. Here, we settle for the first strategy, using a variant of Rannacher timestepping and introduce intermediate backward Euler steps periodically between CN steps as an alternative to the “pure” CN scheme.

Employing a predictor-corrector approach (for details, see, e.g., [94, 138, 139, 147]), we control the time step size of the CN or Rannacher (R) scheme by comparing it to an explicit Adams–Bashforth method of second order (AB2), which, however, is based on the last time steps’ implicit solutions, as it would not be stable enough considering it in a fully explicit manner. The AB2 scheme approximates the velocity via

$$\tilde{\mathbf{u}}_h^{n+1} = \mathbf{u}_h^n + \frac{\Delta t^n}{2} \left[\left(2 + \frac{\Delta t^n}{\Delta t^{n-1}} \right) \dot{\mathbf{u}}_h^n - \frac{\Delta t^n}{\Delta t^{n-1}} \dot{\mathbf{u}}_h^{n-1} \right], \quad (2.90)$$

with accelerations at previous time instants, $\partial_t \mathbf{u}_h^n \approx \dot{\mathbf{u}}_h^n$ and $\partial_t \mathbf{u}_h^{n-1} \approx \dot{\mathbf{u}}_h^{n-1}$, incorporating solutions obtained with the implicit CN or R scheme, \mathbf{u}_h^{n+1} and \mathbf{u}_h^n , as

$$\dot{\mathbf{u}}_h^n = \frac{2}{\Delta t^n} (\mathbf{u}_h^{n+1} - \mathbf{u}_h^n) - \dot{\mathbf{u}}_h^{n-1}. \quad (2.91)$$

Then, comparing the solution from the implicit velocity-pressure system, \mathbf{u}_h^{n+1} , with the explicit one, $\tilde{\mathbf{u}}_h^{n+1}$, the new time step size is set to

$$\Delta t^{n+1} = \xi \Delta t^n \left(\frac{\epsilon_u}{e^{n+1}} \right)^{1/3}, \quad \text{with} \quad e^{n+1} = \frac{\|\mathbf{u}_h^{n+1} - \tilde{\mathbf{u}}_h^{n+1}\|}{3(1 + \Delta t^{n-1}/\Delta t^n)}, \quad (2.92)$$

and a specified tolerance ϵ_u on the velocities' truncation error in the L^2 -norm and a safety factor $\xi \in (0, 1]$.

Now, two issues with the AB2 scheme remain to be addressed. First, it is not self-starting, since $\dot{\mathbf{u}}_h^{n-1}$ is not available. Here, simple remedies are to use finite difference formulas based on a few implicit steps and start the explicit method afterwards, or set $\dot{\mathbf{u}}_h^{n-1} \approx \mathbf{0}$, which is of course only correct if the boundary conditions are smoothly ramped in a consistent way starting from the quiescent state [104]. Second, the explicit AB2 integrator is prone to the ringing phenomenon, which describes the approximation fluctuating around the true solution, even for perfectly smooth functions. This effect can be reduced by stabilising, i.e., averaging and re-initialising the solution vectors of the AB2 scheme given the implicit solution vectors (cf. [94, 138, 139])

$$\dot{\mathbf{u}}_h^{n-1} = \frac{1}{2} \left(\dot{\mathbf{u}}_h^n + \dot{\mathbf{u}}_h^{n-1} \right), \text{ and } \dot{\mathbf{u}}_h^n = \frac{1}{\Delta t^n} (\mathbf{u}_h^{n+1} - \mathbf{u}_h^n), \quad (2.93)$$

in every N^* -th step. A similar linear combination might be employed to resolve the issue regarding the initialisation of the AB2 scheme as mentioned above. In the numerical examples presented herein, the averaging steps and intermediate Euler steps of the Rannacher scheme are synchronised although this is not strictly necessary.

Differing from the stabilised time integrators in [94, 138, 139], we do not apply an averaging such as (2.93) to the implicit solutions' velocity components, but resort to Rannacher timestepping, i.e., intermediate implicit Euler steps. The main motivation is to stabilise the pressure in the implicit scheme as well, which is otherwise left untouched. Furthermore, the parameter N^* is considered a fixed, user-defined and (mildly) problem-dependent value. A compromise between accuracy and stability has to be found, since choosing N^* too low leads to excessive numerical dissipation and reduces accuracy too much. Setting N^* too high, on the other side, causes the integrator to stall at a certain step size, even if the target tolerance is not violated. In this case, fluctuations caused by ringing hinder the integrator from increasing the step size further as is demonstrated in numerical examples in Sec. 2.7.

2.6.2 Fractional-step θ -scheme

As an alternative to comparing discretisations from explicit and implicit time integration schemes in a predictor-corrector approach, Runge–Kutta methods with embedded lower-order schemes also allow estimating the error and adapting the time step size accordingly. The beauty of embedded methods lies in the fact that no additional implicit steps are required, but reusing the already computed intermediate steps allows computing another approximation of the solution with lower accuracy. Herein, we compare the stabilised predictor-corrector approach with a second-order accurate and *strongly* A-stable pressure-corrected fractional-step θ -scheme [83]. This is easily achieved by grouping three steps of the generalised θ -scheme to a macro step from t^m to t^{m+1} and selecting time integration parameters and substep lengths according to Tab. 2.2, where

Tab. 2.2: Time integration parameters and step lengths for the FS scheme.

Step	θ	θ'	Δt^n	t^n	t^{n+1}
1	α	β	$\kappa\Delta t^m$	t^m	$t^m + \kappa\Delta t^m$
2	β	α	$\tilde{\kappa}\Delta t^m$	$t^m + \kappa\Delta t^m$	$t^m + (1 - \kappa)\Delta t^m$
3	α	β	$\kappa\Delta t^m$	$t^m + (1 - \kappa)\Delta t^m$	t^{m+1}

the auxiliary coefficients

$$\kappa := 1 - \frac{1}{2}\sqrt{2}, \quad \tilde{\kappa} := 1 - 2\kappa, \quad \alpha := \frac{\tilde{\kappa}}{1 - \kappa}, \quad \beta := 1 - \alpha \quad (2.94)$$

are introduced for ease of notation.

Following Rang [83], the fractional-step θ -scheme is readily interpreted as a Runge–Kutta method presented in the popular format of a Butcher tableau [148]

$$\begin{array}{c|c} \mathbf{c} & \mathbf{A} \\ \hline & \mathbf{b}^\top \end{array} = \begin{array}{c|cccc} & 0 & 0 & 0 & 0 \\ \hline \kappa & \kappa\beta & \kappa\alpha & 0 & 0 \\ \kappa + \tilde{\kappa} & \kappa\beta & (\kappa + \tilde{\kappa})\alpha & \tilde{\kappa}\beta & 0 \\ \hline 1 & \kappa\beta & (\kappa + \tilde{\kappa})\alpha & (\kappa + \tilde{\kappa})\beta & \kappa\alpha \\ \hline & \kappa\beta & (\kappa + \tilde{\kappa})\alpha & (\kappa + \tilde{\kappa})\beta & \kappa\alpha \end{array}, \quad (2.95)$$

where a modified coefficient vector $\tilde{\mathbf{b}}$ [140], given as

$$\tilde{\mathbf{b}} = \begin{pmatrix} 0.11785113033497070959 \\ 0.49509379160690495120 \\ 0.29636243203812433921 \\ 0.09069264621404818692 \end{pmatrix}, \quad (2.96)$$

yields an embedded method of first order, which is entirely based on the intermediate time steps' solutions.

Note here, however, that starting from three steps of the generalised θ -scheme to compute the intermediate solutions, the embedded method can not be employed in a straight forward way. Put as an open problem in [83], the reason for this peculiarity comes from the fact that Runge–Kutta methods are usually formulated in terms of coefficient vectors approximating the time derivative of the solution at intermediate time steps and not in terms of the solution itself, which are computed with the standard generalised θ -scheme. It is, however, possible to recover a single coefficient vector per macro time step using a simple projection, which then allows re-use of existing implementations based on the generalised θ -scheme. To illustrate this work's contribution, we give a very brief introduction to Runge–Kutta methods at this point, but refer the interested reader to [83, 140, 149] and the references therein for a more detailed discussion.

So, for the coupled velocity-pressure formulation of the Navier–Stokes equations, we can write the momentum and continuity equations as

$$\partial_t \mathbf{u} = \mathbf{f}(t, \mathbf{u}, p) := 1/\rho [\mathbf{b} - (\nabla \mathbf{u})\mathbf{u} - \nabla p + \mu \Delta \mathbf{u} + \chi \mu \nabla (\nabla \cdot \mathbf{u}) + 2\nabla^s \mathbf{u} \nabla \mu], \quad (2.97)$$

$$0 = g(t, \mathbf{u}, p) := \nabla \cdot \mathbf{u}, \quad (2.98)$$

where all terms independent of time derivatives are grouped into \mathbf{f} and g , respectively. When employing the PPE-based stabilisation, Eqn. (2.98) is replaced by

$$0 = g(t, \mathbf{u}, p) := \Delta p + \nabla \cdot [\rho(\nabla \mathbf{u})\mathbf{u} - 2\nabla^s \mathbf{u} \nabla \mu - \mathbf{b}] + [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu - \beta \nabla \cdot \mathbf{u}, \quad (2.99)$$

which is the PPE in strong form (2.38). From the current viewpoint, however, Eqn. (2.99) is completely equivalent to Eqn. (2.98) in the sense that the equation governing the pressure variable lacks a pressure time derivative. Consequently, one might use Eqns. (2.98) and (2.99) interchangeably in what follows. Furthermore, we assume fully consistent boundary conditions $\mathbf{u}|_{t=t^0} = \mathbf{u}_0$ and $p|_{t=t^0} = p_0$ given, as this issue has been discussed in the context of generalised θ -timestepping.

Then, an s -stage Runge–Kutta method advances the solution components \mathbf{u} and p in time via

$$\mathbf{u}^{m+1} = \mathbf{u}^m + \Delta t^m \sum_{i=1}^s b_i \mathbf{k}^i, \quad p^{m+1} = p^m + \Delta t^m \sum_{i=1}^s b_i l^i, \quad (2.100)$$

where a total of s auxiliary variables \mathbf{k}^i and l^i are defined via $i = 1, \dots, s$ in

$$\mathbf{k}^i = \mathbf{f} \left(t^m + c_i \Delta t^m, \mathbf{u}^m + \Delta t^m \sum_{j=1}^s \mathbf{A}_{ij} \mathbf{k}^j, p^m + \Delta t^m \sum_{j=1}^s \mathbf{A}_{ij} l^j \right), \quad (2.101)$$

$$0 = g \left(t^m + c_i \Delta t^m, \mathbf{u}^m + \Delta t^m \sum_{j=1}^s \mathbf{A}_{ij} \mathbf{k}^j, p^m + \Delta t^m \sum_{j=1}^s \mathbf{A}_{ij} l^j \right). \quad (2.102)$$

The issue of not having a time derivative on the pressure in neither the original continuity equation nor the PPE for fully incompressible fluids leads to the coefficients l^i not well defined by Eqn. (2.102). To circumvent this complication, we insert the update rule (2.100) into the auxiliary equations (2.101)–(2.102), which results in

$$\mathbf{k}^1 = \mathbf{f}(t^m, \mathbf{u}^m, p^m), \quad 0 = g(t^m, \mathbf{u}^m, p^m), \quad (2.103)$$

$$\mathbf{k}^2 = \mathbf{f}(t^m + \kappa \Delta t^m, \mathbf{u}^I, p^I), \quad 0 = g(t^m + \kappa \Delta t^m, \mathbf{u}^I, p^I), \quad (2.104)$$

$$\mathbf{k}^3 = \mathbf{f}(t^m + (\kappa + \tilde{\kappa}) \Delta t^m, \mathbf{u}^{II}, p^{II}), \quad 0 = g(t^m + (\kappa + \tilde{\kappa}) \Delta t^m, \mathbf{u}^{II}, p^{II}), \quad (2.105)$$

$$\mathbf{k}^4 = \mathbf{f}(t^{m+1}, \mathbf{u}^{III}, p^{III}), \quad 0 = g(t^{m+1}, \mathbf{u}^{III}, p^{III}), \quad (2.106)$$

with the intermediate time steps' solution components of the velocity, \mathbf{u}^I , \mathbf{u}^{II} and \mathbf{u}^{III} , given by

$$\mathbf{u}^I = \mathbf{u}^m + \kappa \Delta t^m (\beta \mathbf{k}^1 + \alpha \mathbf{k}^2), \quad \mathbf{u}^{II} = \mathbf{u}^I + \tilde{\kappa} \Delta t^m (\alpha \mathbf{k}^2 + \beta \mathbf{k}^3), \quad (2.107)$$

$$\mathbf{u}^{III} = \mathbf{u}^{II} + \kappa \Delta t^m (\beta \mathbf{k}^3 + \alpha \mathbf{k}^4), \quad (2.108)$$

and intermediate pressures p^I , p^{II} and p^{III} as

$$p^I = p^m + \kappa \Delta t^m (\beta l^1 + \alpha l^2), \quad p^{II} = p^I + \tilde{\kappa} \Delta t^m (\alpha l^2 + \beta l^3), \quad (2.109)$$

$$p^{III} = p^{II} + \kappa \Delta t^m (\beta l^3 + \alpha l^4). \quad (2.110)$$

As pointed out by Rang [83], one of course obtains $\mathbf{u}^{III} \equiv \mathbf{u}^{m+1}$ and $p^{III} \equiv p^{m+1}$, when inserting the intermediate time steps' solutions into the update rule (2.100), but even better, we see that Eqns. (2.103)–(2.110) are exactly the systems corresponding to the intermediate steps solved in the fractional-step θ -scheme.

The embedded scheme computes the solution at t^{m+1} via

$$\tilde{\mathbf{u}}^{m+1} = \mathbf{u}^m + \Delta t^m \sum_{i=1}^s \tilde{b}_i \mathbf{k}^i, \quad (2.111)$$

but unfortunately, we cannot use any combination of the update formula (2.100), the definitions of the intermediate solutions (2.107)–(2.110) or $\mathbf{A}_{4j} = \mathbf{b}_j$ to reconstruct the \mathbf{k}^i , even though all the relevant solution vectors, i.e., \mathbf{u}^m , \mathbf{u}^I , \mathbf{u}^{II} and $\mathbf{u}^{III} = \mathbf{u}^{m+1}$ are available after solving the full macro time step.

As shown by the authors in [2], the crucial point enabling adaptive timestepping based on the embedded scheme *keeping* the formulation in fluid velocity and pressure is an additional single projection per macro time step to compute the coefficient vector \mathbf{k}^1 as follows. Given the solution $(\mathbf{u}, p, \mu) \in [X_h]^d \times X_h \times X_h$, we can directly use Eqn. (2.103) to project the fluid acceleration \mathbf{k}^1 onto $[X_h]^d$. The variational problem is thus to find $\mathbf{k}^1 \in [X_h]^d$ with $\mathbf{k}^1|_{\Gamma_D} = \partial_t \mathbf{g}$, such that at $t = t^m$ there holds

$$\begin{aligned} \langle \mathbf{w}, \rho \mathbf{k}^1 \rangle &= \langle \mathbf{w}, \mathbf{b} - \rho(\nabla \mathbf{u}) \mathbf{v} \rangle - (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ &\quad + (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle + \langle \nabla \cdot \mathbf{w}, p \rangle + \langle \mathbf{w}, \bar{\mathbf{t}}|_{t=t^m} \rangle_{\Gamma_N}, \end{aligned} \quad (2.112)$$

for all $\mathbf{w} \in [X_h]^d$ with $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$. This can be rewritten with the help of $a(\mathbf{w}, \mathbf{u}, \mathbf{v}, p, \mu, t^i)$ as defined in Eqn. (2.61) as

$$\langle \mathbf{w}_h, \mathbf{k}_h^1 \rangle = -\frac{1}{\rho} a(\mathbf{w}_h, \mathbf{u}_h^m, \mathbf{u}_h^m, p_h^m, \mu_h^m, t^m). \quad (2.113)$$

This corresponds to a mass matrix problem, which is easily solved adopting similar strategies as for the viscosity projection step. Again, a single V-cycle of an AMG method with Chebyshev smoother [95] to precondition the CG solver [100] is found robust and effective, but might be replaced by cheaper alternatives such as Jacobian preconditioning or even lumping the mass matrix, reducing the associated cost further.

The remaining auxiliary variables \mathbf{k}_h^i , $i = 2, 3, 4$ are then easily recovered rewriting Eqns. (2.107)–(2.108) as

$$\begin{aligned} \mathbf{k}_h^2 &= \frac{1}{\alpha} \left(\frac{\mathbf{u}_h^I - \mathbf{u}_h^m}{\kappa \Delta t^m} - \beta \mathbf{k}_h^1 \right), & \mathbf{k}_h^3 &= \frac{1}{\beta} \left(\frac{\mathbf{u}_h^{II} - \mathbf{u}_h^I}{\tilde{\kappa} \Delta t^m} - \alpha \mathbf{k}_h^2 \right), \\ \mathbf{k}_h^4 &= \frac{1}{\alpha} \left(\frac{\mathbf{u}_h^{III} - \mathbf{u}_h^{II}}{\kappa \Delta t^m} - \beta \mathbf{k}_h^3 \right), \end{aligned} \quad (2.114)$$

which finally allows computing the solution at the end of the macro step t^{m+1} with the embedded scheme's coefficients $\tilde{\mathbf{b}}$. The macro time step size of the next step Δt^{m+1} is then selected via [150]

$$\Delta t^{m+1} = \xi \frac{(\Delta t^m)^2}{\Delta t^{m-1}} \sqrt{\frac{\epsilon_u e^m}{(e^{m+1})^2}}, \quad \text{with } e^{m+1} = \|\mathbf{u}_h^{m+1} - \tilde{\mathbf{u}}_h^{m+1}\|, \quad (2.115)$$

where, again, $\xi \in (0, 1]$ is a safety factor and ϵ_u denotes the tolerance in the L^2 -norm of the velocities' truncation error. Note here that the time step is thus adapted at the end of each *macro* time step, which will become relevant in the numerical experiments presented in Sec. 2.7.

To summarise, we enable adaptive timestepping based on a classical implementation of the fractional-step θ -scheme by introducing a single mass matrix solve per macro time step to project the auxiliary variable $\mathbf{k} = \partial_t \mathbf{u}$. This projection could even be decomposed into the d individual components or one might apply mass-lumping, since it is used for time step selection only, but the computational effort is close to negligible in any case—already when comparing the vector-valued mass matrix solve to the three intermediate steps' solves of the linearised, but coupled velocity-pressure system.

In addition to the safety factor ξ , we also repeat time steps that lead to a reduction of the time step size by more than 30% or resulted in divergence of the linear or nonlinear solvers with halved step size. This additional safety measure popular in adaptive timestepping further robustifies the solver, especially if rapid changes in the time step size are encountered, which may lead to reduced quality in the initial guesses computed via extrapolation in time (see, e.g., [139, 147, 151]).

2.7 Computational results

Now, having established the stabilised velocity-pressure formulation, suitable timestepping schemes and preconditioners, this section is devoted to assessing its accuracy and performance in various numerical examples. In particular, we compare the newly devised scheme with available methods from literature, starting from two-dimensional, stationary, Newtonian flow problems working our way up to a realistic setting in haemodynamics. After showcasing convergence in space and time in classical benchmark scenarios and a manufactured solution, we discuss the different timestepping schemes in terms of nonlinear and linear solver iterations per time step in pulsatile flow problems and

within a pseudo-timestepping approach to critically evaluate the scheme's performance. In the last numerical test, we simulate blood flow in an idealised cerebral aneurysm to demonstrate the solvers' capabilities in the cardiovascular context under physiological flow conditions and to highlight relevance of rheological modelling in hemodynamics. In some of the initial examples focusing on the accuracy of the formulation itself, varying stabilisation parameters or computing convergence rates, we will partly use a parallel direct solver [152] since the linear systems are rather small and we can in fact use a direct solver to rule out any ambiguities stemming from linear solver tolerances and to avoid small pseudo timesteps. In the later examples, when investigating solver performance or time integration schemes, the tolerances chosen do play an important role, in which case we will specifically mention them—in any other case, however, tolerances are set appropriately, such that the solution is not affected. The implementation is based on the open-source finite element toolbox `deal.II` [96–99] and heavily relies on its interfaces to the AMG methods provided by [95, 111].

The convergence rates are measured in relative L^2 -norms defined as

$$\|p - p_h\|_0 := \frac{\|p - p_h\|_{L^2(\Omega)}}{\|p\|_{L^2(\Omega)}} \quad \text{and} \quad \|\mathbf{u} - \mathbf{u}_h\|_0 := \frac{\|\mathbf{u} - \mathbf{u}_h\|_{L^2(\Omega)}}{\|\mathbf{u}\|_{L^2(\Omega)}}, \quad (2.116)$$

for stationary problems. If we consider time-dependent cases, we use

$$\begin{aligned} e_Q^{\mathbf{u}} &:= \frac{\|\mathbf{u}_h - \mathbf{u}\|_{L^2(0,T;L^2(\Omega))}}{\|\mathbf{u}\|_{L^2(0,T;L^2(\Omega))}} & \text{and} & & e_{\Omega}^{\mathbf{u}} &:= \frac{\|\mathbf{u}_h - \mathbf{u}\|_{L^2(\Omega)}}{\|\mathbf{u}\|_{L^2(\Omega)}} \Bigg|_{t=T} \\ e_Q^p &:= \frac{\|p_h - p\|_{L^2(0,T;L^2(\Omega))}}{\|p\|_{L^2(0,T;L^2(\Omega))}} & \text{and} & & e_{\Omega}^p &:= \frac{\|p_h - p\|_{L^2(\Omega)}}{\|p\|_{L^2(\Omega)}} \Bigg|_{t=T}. \end{aligned} \quad (2.117)$$

Moreover, we employ the generalised Laplacian form of the momentum equation, i.e., set $\chi = 0$ if not specified otherwise and introduce a parameter α to scale the stabilisation parameter τ_e (2.35) additionally, yielding $\tau_e \leftarrow \alpha\tau_e$.

2.7.1 Kovasznay flow problem

We start off with the classical Kovasznay flow problem [153] as presented in [1], which describes the stationary, laminar flow of an incompressible, Newtonian fluid behind an array of cylinders and features a zero volumetric force, that is, $\mathbf{b} = \mathbf{0}$. The domain $\Omega = (-1/2, 1/2) \times (-1/2, 1/2)$ is considered, wherein the solution is given as

$$\begin{aligned} \mathbf{u} &= \begin{pmatrix} 1 - e^{kx} \cos 2\pi y \\ \frac{k}{2\pi} e^{kx} \sin 2\pi y \end{pmatrix} \quad \text{and} \quad p = 1/2 (e^k - e^{2kx}), \\ \text{with } k &:= \frac{\text{Re}}{2} - \sqrt{\left(\frac{\text{Re}}{2}\right)^2 + (2\pi)^2} \end{aligned}$$

and Re being the Reynolds number. We solve the corresponding pure Dirichlet problem, i.e., $\Gamma_D = \partial\Omega$, with linear triangular elements for $Re = 100$ and stabilisation parameters chosen as $\alpha = \alpha_1 = 1$ and $\alpha_2 = \alpha_3 = \gamma_e = 0$ (stationary, diffusion-dominant case without grad-div stabilisation). The initial grid as shown in Fig. 2.1 is seven times uniformly refined and errors in the relative L^2 -norms for velocity and pressure shown in Fig. 2.2 result. For comparison, errors obtained with the PSPG method and the PPE-based splitting scheme by Johnston and Liu [20], the former one presented in Eqns. (2.34)–(2.34) and the latter obtained when $\alpha \rightarrow \infty$, which is related to the split-step scheme discussed in detail in Ch. 3.

We observe that for both PSPG and the proposed boundary vorticity stabilisation (BVS) velocity and pressure errors converge at the same rates. However, the errors obtained with the latter are considerably smaller. Using linear triangular elements, we directly see the drastic consequences of having an “inviscid” residual in the PSPG formulation: the asymptotic convergence rate is reached two global refinements later than with the BVS method. This behaviour is caused by artificial pressure boundary layers, which are only becoming negligible once the element size h_e is sufficiently small [27]. In this specific example, this leads to an improvement of velocity and pressure errors by an order of magnitude. Comparing to the splitting scheme denoted by PPE in Fig. 2.2, we see convergence rates of order one [20, 154] in the velocity and pressure and a tremendous gain in accuracy achieved by adding the penalty term only.




Fig. 2.1: Kovasznay flow problem: initial grid refined up to seven times uniformly in the course of the convergence study [1].

This effect of spurious pressure boundary layers spoiling the coarse-grid accuracy can be illustrated nicely in the Kovasznay benchmark, since the exact solution features perfectly vertical pressure isolines. Comparing the pressure isolines as obtained with the PSPG and BVS methods at mesh level 5 and $Re = 40$, $\alpha = 100$ in Fig. 2.3, we observe a noticeable impact of the incomplete residual. Note here that the stabilisation parameter was chosen outside of the optimal range to showcase the effect better. This highlights the vast differences between the methods once the stabilisation parameter is chosen somewhat less optimal, which quickly leads to distortions polluting the pressure field in the whole domain for the PSPG method. The pressure as obtained via the BVS method, on the other side, is only mildly affected, and shows distortions in the corners of the domain only.

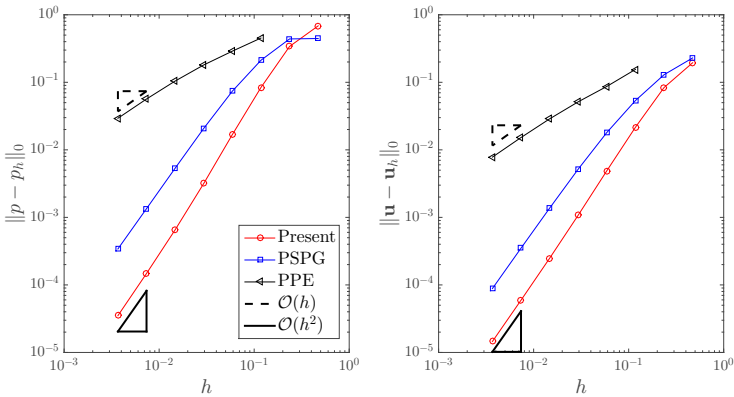


Fig. 2.2: Kovaszny flow problem: errors in the relative pressure (left) and velocity (right) norms comparing the BVS method (“Present”), PSPG and splitting scheme (PPE) [20]. The BVS scheme yields expected orders and outperforms the PSPG method by an order of magnitude [1].




Fig. 2.3: Kovaszny flow problem: pressure isolines obtained with the PSPG (left) or BVS formulation (right), where exact solution’s pressure isolines are perfectly vertical. The BVS scheme shows only mild distortions in the corners, whereas the PSPG solution is polluted by spurious pressure boundary layers [1].

So far we have seen that the incomplete residual in the PSPG method does cause spurious pressure boundary layers and can reduce accuracy by a substantial margin, especially if the stabilisation parameter is not chosen optimally. Now, this naturally leads to the question of how the PSPG and BVS methods compare, when the viscous contribution to the residual is *not* lost. To investigate this matter, we employ equal-order interpolations up to order $k = 4$ using stabilised $Q_k Q_k$ elements and vary the stabilisation parameter α in a uniform 64×64 grid at a Reynolds number of $\text{Re} = 100$. The errors obtained in the relative velocity and pressure norms are depicted in Fig. 2.4 and one clearly observes that (i) the BVS scheme yields errors orders of magnitude lower than the PSPG method, if the stabilisation parameter is out of the optimal range and performs at least as good as the PSPG method, and (ii) the BVS has a much wider range of admissible stabilisation parameters compared to the PSPG method, whereas the latter one even fails to converge for unsuitable α given a zero initial guess for the nonlinear Picard solver.

In both methods, using α too small leads to increased errors due to the lack of stabilisation, after which the suitable range for α follows giving comparable errors which then increase again as $\alpha \rightarrow \infty$. In the PSPG method for polynomial degrees higher than $k = 1$, the diffusive terms in the residual are *not* lost as opposed to the linear case, but choosing a stabilisation parameter α too large leads to an excessive relaxation of the divergence-free constraint, causing the increased errors. On the contrary, letting $\alpha \rightarrow \infty$ in the BVS formulation, the incompressibility constraint is replaced by the fully consistent PPE, which also weakly enforces a divergence-free velocity, but with reduced accuracy [154]. Therefore, the errors obtained with the BVS scheme are considerably lower than those of the PSPG method for large α and vary much smoother.

Additionally, note that the admissible parameter range is considerably wider for the BVS method, which is a crucial property of any stabilisation technique, as selecting the stabilisation parameters is a critical, problem-dependent task often left unattained. The BVS formulation is thus undoubtedly less sensitive to the stabilisation parameter, which for $k = 4$ yields excellent results over the entire parameter range $\alpha \in [10^{-5}, 10^5]$, unlike the PSPG method, for which the error can shoot up as the stabilisation parameter moves away in any direction from the potentially narrow optimal range. If we compare the errors for $\alpha = 1$, i.e., the standard recommended choice for the PSPG method, the BVS using the identical parameter can yield errors of up to one order in magnitude lower in both pressure and velocity, which corresponds well with the previously presented results using $P_1 P_1$ discretisation.

In a final variation of this numerical test, we consider an anisotropic, graded mesh motivated by the fact that capturing the physical boundary layers is vital in practical flow simulations. A coarse discretisation of the domain with a variation in element size of up to 1:9 between neighbouring elements and a maximal aspect ratio of 1:9 as shown in Fig. 2.5(a) is repeatedly refined uniformly as depicted for levels 2 and 3 in Fig. 2.5(b). We set $\alpha = 10$ to compute relative velocity and pressure norms shown in Fig. 2.6, which correspond to solutions on a sequence of 7 grids, where we can once again verify optimal convergence rates and good stability properties, but this time with a more pronounced pre-asymptotic range due to a coarse initial mesh used in the bulk of the domain. Similar effects are to be expected for three-dimensional problems with graded meshes, which are covered in the following.

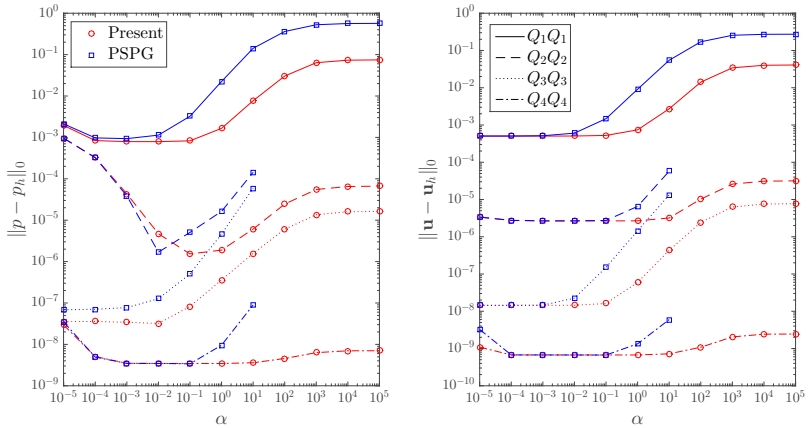


Fig. 2.4: Kovaszny flow problem: pressure (left) and velocity errors (right) in a 64×64 uniform grid of $Q_k Q_k$ elements. The BVS scheme (“Present”) has a considerably wider admissible parameter range than PSPG, which fails to converge for large α [1].




Fig. 2.5: Kovaszny flow problem: anisotropic discretisations used in the refinement study with a maximal aspect ratio of 1:9 and size ratio of up to 1:9 between elements [1].




Fig. 2.6: Kovasznay flow problem: relative velocity and pressure norms obtained with the anisotropic grid and the BVS scheme with $\alpha = 10$. The pre-asymptotic range is larger compared to the uniform grid case, which is due to the large elements used in the bulk of the domain [1].

2.7.2 Poiseuille flow problem in 3D

Within this numerical test taken from [1], we consider a cylindrical pipe in three dimensions with radius R and length L , given as

$$\Omega = \left\{ (x, y, z) \in \mathbb{R}^3 : r := \sqrt{y^2 + z^2} \leq R, 0 \leq x \leq L \right\},$$

and restrict ourselves to stationary flow of a Newtonian fluid, which gives the classical Poiseuille flow with zero body force, $\mathbf{b} = \mathbf{0}$,

$$\mathbf{u} = (u_1(r), 0, 0)^\top, \quad p = \frac{4\mu L \hat{u}_1}{R^2} \left(1 - \frac{x}{L}\right), \quad \text{with} \quad u_1(r) := \hat{u}_1 \left[1 - \left(\frac{r}{R}\right)^2\right],$$

resulting in the volumetric flow rate $Q := \frac{1}{2}\pi \hat{u}_1 R^2$. We prescribe this solution at the inlet at $x = 0$ and set $\mathbf{u} = \mathbf{0}$ at $r = R$, that is, no-slip conditions at the walls. At the outlet, $x = L$, zero pseudo tractions are enforced, $\tilde{\mathbf{t}} = \mathbf{0}$, aiming for minimal upstream disturbance.

First, we investigate the effect of a variation in the stabilisation parameter, for which we employ the mesh depicted in Fig. 2.7. This discretisation using $\approx 92 \times 10^3$ trilinear hexahedral elements yields for both the BVS and PSPG methods $\approx 39 \times 10^4$ degrees of freedom (DoFs) given that the mesh contains $\approx 97 \times 10^3$ nodes. The finite element mesh features a standard boundary layer with elements refined in radial direction towards the boundary at $r = R$, as is typical for pipe flow.

We proceed in setting $L/R = 3$, $\text{Re} = 150$ and choose stabilisation parameters $\alpha_1 = 1$ and $\alpha_2 = \alpha_3 = \gamma_e = 0$, i.e., the standard setting for stationary, diffusion-dominant scenarios without grad-div stabilisation. In this case, we again observe that the proposed BVS scheme outperforms the classical PSPG method in terms of accuracy and sensitivity




Fig. 2.7: Poiseuille flow problem: mesh for the cylindrical pipe with $L/R = 3$, consisting of $\approx 92 \times 10^3$ trilinear hexahedral elements and $\approx 97 \times 10^3$ nodes with refinement towards the boundary at $r = R$ [1].

for various α as Fig. 2.8 highlights. Note here, that non-Cartesian grids in fact do not completely lose the Laplacian contribution to the residual, but nonetheless, PSPG does not perform as well as BVS. Similar to the previous two-dimensional setting using higher-order finite elements, there is a noticeable difference in the solutions obtained. For large enough α , the PSPG relaxes the incompressibility constraint too much, while for the BVS, the continuity equation is replaced by the PPE scheme converging with comparably lower accuracy. That is, the same trends as in the previously presented Kovaszny benchmark are observed.




Fig. 2.8: Poiseuille flow problem: relative pressure (left) and velocity (right) norms with parameter $\alpha \in [10^{-5}, 10^5]$. The optimal parameter range for BVS (“Present”) is considerably wider than for PSPG. In this scenario, BVS outperforms PSPG by two orders of magnitude for α chosen too big [1].

Performing a convergence study, we construct a series of uniformly refined meshes with cross-section as depicted in Fig. 2.9 and an element length of $L/3$ in x -direction. Relative errors in the velocity and pressure in the normalised setting $2R/L = \rho Q/\mu R = 1$ are computed using a direct solver. The resulting errors are listed in Tab. 2.3, clearly showing an estimated convergence rate of two in both solution components, which is optimal as expected for smooth solutions.




Fig. 2.9: Poiseuille flow problem: cross-sections of the meshes at refinement levels 1 to 4 considered in the convergence study [1].

Tab. 2.3: Poiseuille flow problem: relative velocity and pressure errors under uniform refinement and resulting estimated order of convergence (eoc) [1].

Level	$\ \mathbf{u} - \mathbf{u}_h\ _0$	eoc	$\ p - p_h\ _0$	eoc
1	7.2×10^{-2}	–	3.5×10^{-2}	–
2	1.3×10^{-2}	2.4	6.0×10^{-3}	2.5
3	3.3×10^{-3}	2.0	1.5×10^{-3}	2.0
4	8.3×10^{-4}	2.0	3.6×10^{-4}	2.0

2.7.3 Lid-driven cavity flow

A final example of stationary Newtonian fluid flow within this section focuses on a more challenging setup: a non-smooth solution with high Reynolds number, specifically the lid-driven cavity benchmark with $\text{Re} = 5000$ as presented in [1]. The classical setup enforces circulatory flow in the unit square $\Omega = (0, 1) \times (0, 1)$ driven by discontinuous Dirichlet boundary conditions. Setting $\mathbf{u} = (1, 0)^\top$ at the upper boundary combined with no-slip conditions, i.e., $\mathbf{u} = \mathbf{0}$ prescribed at $x = 0$, $y = 0$ and $x = 1$, discontinuities result in the upper corners. This jump in the boundary conditions is resolved employing the regularised variant of the benchmark as proposed by de Frutos et al. [155], however, the pressure still has a strong peak at $x = y = 1$.

The discretisation in space is carried out employing a mesh consisting of 128×128 bilinear Q_1Q_1 elements, while stabilisation parameters are set to $\alpha = \alpha_1 = \alpha_2 = 1$, $\alpha_3 = 0$ and $\gamma_e = 0.3h_e$, i.e., the standard setting for stationary, convection-dominated scenarios *with* grad-div stabilisation. Note here that the grid considered is fine enough to not suffer

from instabilities caused by convective effects, not necessitating the use of other residual-based methods such as Streamline upwind/Petrov–Galerkin [73, 74] (SUPG) or Galerkin Least-Squares [33] (GLS) stabilisations.

Results obtained with the BVS and PSPG methods are shown in Fig. 2.10, where an excellent agreement with reference values of the velocity components at $x = 1/2$ and $y = 1/2$ as published by Erturk et al. [156] and Ghia et al. [157] is found. The pressure fields obtained via the stabilised schemes do not show any spurious oscillations despite of their non-smooth character. Note also the good agreement of the BVS and PSPG schemes, which stems from the combination of (i) a high Reynolds number present, rendering the viscous contribution in the residual of the momentum balance rather small to begin with, (ii) the stabilisation parameter chosen in the expected optimal range ($\alpha = 1$) and (iii) grad-div stabilisation added on top.




Fig. 2.10: Lid-driven cavity flow: comparison of velocity results with reference values given by Erturk et al. [156] and Ghia et al. [157].




Fig. 2.11: Lid-driven cavity flow: pressure distribution in the domain $\Omega = (0, 1) \times (0, 1)$ with a strong peak in the upper right corner [1].

In the numerical tests presented up to now, only stationary flow of Newtonian fluids is considered, showing the improved robustness of BVS compared to PSPG, an enlarged

admissible parameter range and lower errors in both the pressure and velocity solution components. Both the BVS and PSPG formulations can be employed for inf-sup stabilisation of generalised Newtonian fluid flows, which will be demonstrated in the following. In addition to that, the effectiveness of the proposed time integration and linearisation schemes is showcased.

2.7.4 Manufactured solution in 2D

Let us now turn our attention to a test case with non-constant viscosity involving a generalised Newtonian fluid. To investigate temporal and spatial convergence rates and to prescribe interpolated boundary- and initial conditions, we derive fields $\mathbf{u} = (u_1, u_2)^\top$ and p in two space dimensions based on the classical form of the Navier–Stokes equations for incompressible flow of viscous fluids given in Eqns. (2.1)–(2.2). Starting from

$$u_1(\mathbf{x}, t) = \sin x \cos y \cos t \quad \text{and} \quad p(\mathbf{x}, t) = \cos(x y) \cos t,$$

we plug u_1 into the continuity equation (2.2) to get

$$u_2(\mathbf{x}, t) = -\cos t \cos x \sin y \quad \text{and} \quad \dot{\gamma}(\mathbf{x}, t) = \cos t \cos x \cos y,$$

by simply inserting \mathbf{u} into the shear rate's definition (2.7), that is,

$$\dot{\gamma} := \sqrt{1/2 \nabla^s \mathbf{u} : \nabla^s \mathbf{u}}.$$

With the shear rate fixed, we can recover the fluid's viscosity as $\mu(\mathbf{x}, t) = \eta(\dot{\gamma}(\nabla^s \mathbf{u}))$, where we settle for the Carreau law as given in Eqn. (2.8) with $a = 2$ and $\kappa = 1$, i.e.,

$$\eta = \eta_\infty + (\eta_0 - \eta_\infty) \left[1 + (\lambda \dot{\gamma})^2 \right]^{\frac{b-1}{2}}.$$

Then, we finally derive an expression for the body force \mathbf{b} from the balance of linear momentum (with $\chi = 0$) with $\mathbf{b} = (b_1, b_2)^\top$ given the components as

$$\begin{aligned} b_1 = & + 2 \cos t \cos y \sin x \left(\eta_\infty + (\eta_0 - \eta_\infty) \left[1 + (\lambda \dot{\gamma})^2 \right]^{\frac{b-1}{2}} \right) \\ & - \rho \cos y \sin t \sin x - y \sin(x y) \cos t + \rho \cos^2 t \cos x \sin x \\ & + 2\rho(b-1)(\lambda \dot{\gamma})^2 \cos t \cos y \sin x (\eta_0 - \eta_\infty) \left[1 + (\lambda \dot{\gamma})^2 \right]^{\frac{b-3}{2}}, \end{aligned}$$

and

$$\begin{aligned} b_2 = & - 2 \cos t \cos x \sin y \left(\eta_\infty + (\eta_0 - \eta_\infty) \left[1 + (\lambda \dot{\gamma})^2 \right]^{\frac{b-1}{2}} \right) \\ & + \rho \cos x \sin t \sin y - x \sin(x y) \cos t + \rho \cos 2t \cos y \sin y \\ & - 2\rho(b-1)(\lambda \dot{\gamma})^2 \cos t \cos x \sin y (\eta_0 - \eta_\infty) \left[1 + (\lambda \dot{\gamma})^2 \right]^{\frac{b-3}{2}}. \end{aligned}$$

The initial conditions for \mathbf{u} are simply interpolated given the velocity field, but also the pressure is considered given at $t = 0$ to rule out any ambiguities coming from the start-up procedure employed. Therefore, we can consider both \mathbf{u}_h^j and p_h^j , $j = n - 1, n$ given at any step $n = 0, \dots, N_t$, including $n = 0$, that is, we interpolate the exact solution at $t = -\Delta t^0$. Regarding boundary conditions, we enforce inhomogeneous Dirichlet conditions given \mathbf{u} along the lines $x = 0$ and $y = 0$, while on the other boundaries, a nonzero pseudo traction $\tilde{\mathbf{t}}$, derived from the given solution, is inserted into the boundary integral. As for the stabilisation parameters, we set τ_e according to Eqn. (2.35) with $\alpha_1 = 2\alpha_3 = 4$ and $\alpha_2 = 0$, and add grad-div stabilisation with γ_e according to Eqn. (2.66). No sharp gradients are present in the solution by design and we do not suffer from spurious oscillations triggered by dominant convection given the selected parameters. With this in place, we can investigate temporal and spatial convergence rates in a general setting.

Convergence in time Consider the domain $\Omega = (0, 0.1) \times (0, 0.1)$ together with a time interval $I_t = (0, 1]$ and choose the fluid parameters as $\rho = 10^3 \text{ kg/m}^3$, $\eta_0 = 10 \text{ mPa s}$, $\eta_\infty = 1 \text{ mPa s}$, $b = 0.5$ and $\lambda = 1 \text{ s}$. To keep spatial discretisation errors low while decreasing the time step size in the temporal convergence study, we employ a uniform grid of 256×256 Q_1Q_1 elements throughout the convergence study. Then, we double the number of (macro) time steps in each level from an initial single one and compare the results obtained via the pressure-corrected Crank–Nicolson integrator (CN), pressure-corrected fractional-step θ -method (FS) and Rannacher timestepping (R) with the constructed exact solution.

The resulting errors e_Q^u , e_Q^p and e_Q^μ as defined in Eqn. (2.117) are illustrated in Fig. 2.12, where we only included results obtained with the BVS formulation. Note, however, that almost identical errors result for the Q_1Q_1 pair stabilised via BVS or PSPG and Q_2Q_1 elements, due to the fact that the spatial discretisation error is negligible and the minimum error level reached is dominated by the nonlinear solver tolerance of $\epsilon_{\text{rel}}^{\text{nl}} = 10^{-6}$ in Eqn. (2.79). The first thing that comes to notice inspecting Fig. 2.12 carefully is that most of the lines indicate convergence rates close to 2, which is just as expected and optimal considering the second-order accurate time integration schemes used. However, two lines clearly show linear convergence rates in the pressure for the R and FS methods, which we want to further elaborate.

Recall that the R scheme is based on CN time integration, but further introduces intermediate backward Euler steps ($\theta = 1$, $\theta' = 0$) every N^* steps. The second order convergence rates in the velocity and viscosity were only achieved with a fixed number of total Euler steps, meaning that we actually double N^* from level to level (keeping the error introduced by a single intermediate Euler step almost constant). Otherwise, i.e., when keeping $N^* = 2$ fixed, linear convergence rates result as the first-order accurate implicit Euler scheme dominates the error. Thus, when using R timestepping, we start with 4 steps in total, and $N^* = 2$ translates to the third step being an Euler step. Doing so leads to linear convergence in the pressure (+) and a shift in the velocity error (\times) comparing R (–) with the optimally-convergent CN (–) scheme, but second-order rates in the velocity and viscosity. The temporal stability is at the same time noticeably improved compared to the CN scheme.

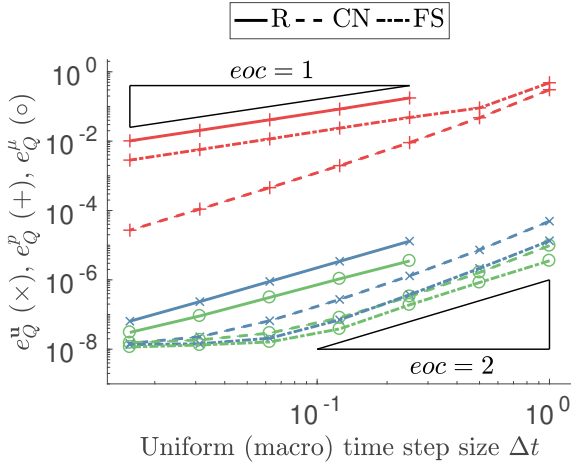


Fig. 2.12: Temporal convergence study: relative errors in the primary variables \mathbf{u} (\times , blue), p ($+$, red) and μ (\circ , green) using a uniform grid of 256×256 Q_1Q_1 BVS elements and pressure-corrected Crank–Nicolson (CN, --), Rannacher (R, –) or fractional-step θ (FS, -·) schemes.

These convergence rates are consistent with the error bounds established in [141], but fully implicit pressure time integration was employed therein, *not* the pressure-corrected variants we employ here. In the present case, when using the pressure-corrected form of R timestepping, we do not observe the expected second-order convergence in the pressure, while the velocity and viscosity errors behave as expected. Comparing to similar temporal stabilisation techniques in [138, 139] and [94], where periodic averaging is introduced every N^* steps, we note that in addition to stabilising the velocities, and thereby counteracting the ringing phenomenon occurring for CN time integration with time steps chosen too large, the intermediate Euler steps act on all involved variables. Thus, possible pressure fluctuations or oscillations in time are also damped in a (globally first-order) and locally second-order accurate way. In practice, the parameter choice for N^* is problem dependent and left to the user, as it controls the specific amount of numerical diffusion, which inevitably increases the overall time integration error compared to basic CN timestepping. For a given number of time steps, it is thus more natural to think of an error corresponding to N^* which *can*, but *does not have to* dominate the time integration error of the basic CN scheme.

The second curve showing linear convergence is the pressure in the FS scheme. This order reduction in the pressure approximation ($+$) for the FS scheme ($-·$) is to be expected, though, owing to the stiffness of the problem (see, e.g., [83] or [158] for a detailed discussion on this matter). Note that the macro time step size is reported in Fig. 2.12 for the FS scheme, meaning we compare p_h^{n+1} at the end of the macro step, not the intermediate steps' solutions, which would not be second-order accurate even in the velocity or viscosity. Summing up, we trade increased temporal stability for accuracy in the pressure variable, which holds for both the FS and the R stepping schemes.

When considering adaptive time step selection, however, even the CN scheme is limited to a first-order accurate pressure, as shown by Jay [159]. This effect might be reduced employing averaging techniques to guarantee a limited increase in time step size [151], since abrupt variations in Δt^n decrease temporal stability and accuracy [160]. The FS scheme suffers from order reduction even in the uniform time step setting as already mentioned, a property which persists when choosing Δt^n nonuniform [83, 159].

In conclusion, all of the presented methods lead to linear convergence in the pressure variable when temporal adaptivity is considered. Consequently, one ends up with linear convergence in time for pressure using any of the above mentioned methods when adaptively choosing the time step size without including any postprocessing. Suitable postprocessing, however, potentially increasing accuracy lies beyond the scope of this work, but might be straight-forwardly included without altering the presented methods. Here, we will focus on performance gains measured in overall time steps and individual field solves for the time interval of interest disconnected from the achieved accuracy.

Convergence in space Now, let us investigate the spatial convergence rates for the case of time-dependent problems. This is to confirm that previous findings for constant viscosity carry over to the generalised Newtonian case. Compared to the foregoing numerical experiment, the setting is slightly adapted, considering $t \in (0, 0.1]$ and integrating in time with 100 CN steps to reduce time integration errors. The spatial approximation error as defined in Eqn. (2.117) is computed at the end of the time interval at $t = T = 0.1$ s, while refining the grid uniformly. Starting from a decomposition of $\Omega = (0, 0.3) \times (0, 0.3)$ into 4×4 elements, we perform a series of uniform refinements to achieve 256×256 elements at the final level. Regarding the ansatz spaces, we employ BVS or PSPG stabilised Q_1Q_1 elements or the LBB-stable Q_2Q_1 pair. The physical parameters $\eta_0 = 10$ mPa s, $\eta_\infty = 1$ mPa s, $\rho = 10$ kg/m³, $b = 0.5$ and $\lambda = 1$ s are chosen. This leads to a Reynolds number in the range of $\text{Re} = 3000$ or $\text{Re} = 300$ depending which one of the viscosity limits, the lower η_∞ or upper η_0 limit, is taken into account.

The spatial errors are reported in Fig. 2.13, showing the expected rates in the velocity, pressure and viscosity fields. Owing to the moderate to high Reynolds number, the BVS (\times) and PSPG (+) schemes yield almost identical errors with rates of 2 for all the involved primary variables, whereas the Taylor–Hood pair (\circ) delivers a cubic convergence rate in the velocity ($-$) and is second order accurate in both the pressure ($--$) and viscosity ($-.$). In this setup, the stabilised methods deliver much worse pressure approximations compared to the inf-sup stable Q_2Q_1 pair, while the difference here depends on the physical parameters. The experimentally observed spatial convergence rates thus naturally extend from the Newtonian case.

Standard velocity-pressure formulations with an inf-sup stable choice of function spaces, i.e., using Taylor–Hood pairs such as Q_iQ_j or P_iP_j with $i > j$, in fact allow using discontinuous pressure spaces, $p \in L^2(\Omega)$. Choosing a higher polynomial degree $i > j$ for the continuous velocity $\mathbf{u} \in [H^1(\Omega)]^d$, convergence of order $j + 1$ in the L^2 -norm of the pressure can be expected. Choosing a stabilised equal-order interpolation such as the BVS scheme or a PSPG stabilised formulation, only convergence of order j can be guaranteed in $\|p_h - p\|_{L^2(\Omega)}$, see [32, 34]. Many of the presented results, however, show

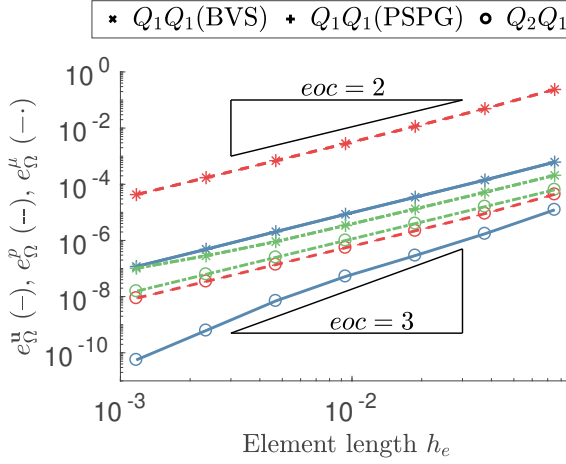


Fig. 2.13: Spatial convergence study: relative errors in the primary variables \mathbf{u} (—, blue), p (—, red) and μ (—, green) using the pressure-corrected Crank–Nicolson scheme with $\Delta t^n = 10^{-3}$ and Q_2Q_1 or Q_1Q_1 elements stabilised via BVS (\times) or PSPG ($+$).

perfectly quadratic convergence using linear interpolation—an effect that is often seen in practice. For a linear pressure ansatz, rates between 1 and 2 are often observed, which depends on the flow regime [27]. The initial higher-order convergence in the pressure variable can not be expected indefinitely, as element sizes where it might be expected depend on the problem parameters as theoretical results suggest [161, 162]. This does not only affect the BVS formulation, but all equal-order methods in general, and hence should be kept in mind when considering stabilised methods such as, e.g., [31, 32, 34, 44, 48, 49].

2.7.5 Carreau fluid channel flow in 2D

Having established the convergence rates for the instationary generalised Newtonian case, let us focus again on the *differences between* the classical PSPG and the newly proposed BVS schemes starting with the stationary case. A constant pressure drop per unit length k along a two-dimensional channel $\Omega = (0, L) \times (-H/2, H/2)$ leads to a developed flow profile of a Carreau fluid, i.e., $a = 2$ and $\kappa = 1$ in Eqn. (2.8), which is given as

$$\mathbf{u}(\mathbf{x}) = \begin{pmatrix} u_1(y) \\ 0 \end{pmatrix}, \quad p(\mathbf{x}) = |k|(L - x), \quad \text{with} \quad u_1(y) = \int_{H/2}^{|y|} f(s) \, ds,$$

where $f(s)$ is governed by

$$\left\{ \eta_{\infty} + (\eta_0 - \eta_{\infty}) \left[1 + (\lambda/2 f(s))^2 \right]^{\frac{b-1}{2}} \right\} f(s) = -|k| s, \quad s \in [0, H/2]. \quad (2.118)$$

The nonlinear equation for $f(s)$ is solved using a fixed-point iteration, and then integrated numerically to obtain \mathbf{u} . Regarding boundary conditions, no-slip conditions are enforced at the upper and lower domain boundaries, $y = -H/2$ and $y = H/2$, while we choose the generalised Laplacian form ($\chi = 0$) and the corresponding pseudo stress in- and outflow conditions at $x = 0$ and $x = L$, respectively, inserting $\hat{\mathbf{t}}$ based on \mathbf{u} and p , which yields $\hat{\mathbf{t}}|_{x=0} = (p_{\text{in}}, 0)^\top$ and $\hat{\mathbf{t}}|_{x=L} = \mathbf{0}$ with a pressure drop of magnitude p_{in} , giving $k = p_{\text{in}}/L$. This (stationary) exact solution is then used to compute approximation errors of the solution components in the relative norms defined in Eqn. (2.116).

We fix the problem dimensions to $L = 3H = 3 \text{ mm}$ and enforce a pressure difference of $p_{\text{in}} = 9 \text{ Pa}$ from $x = 0$ to $x = L$ as typical in the context of blood flow in small vessels. Additionally, rheological parameters representative for the haemodynamic regime are chosen as $\rho = 1050 \text{ kg/m}^3$, $\eta_\infty = 3.45 \text{ mPa s}$, $\eta_0 = 56 \text{ mPa s}$, $b = 0.3568$ and $\lambda = 3.313 \text{ s}$ taken from [163]. The stabilisation parameters are considered as $\alpha_1 = 1$ and $\alpha_2 = \alpha_3 = \gamma_e = 0$, that is, the standard setting for stationary, diffusion-dominant scenarios without grad-div stabilisation. Similar to previous investigations, we first set the additional scaling parameter $\alpha = 1$, meaning, we use a PSPG parameter in the optimal range, and perform a spatial convergence study.

Doing so, we start from a grid of 6×2 quadrilaterals divided along their diagonals to get in total 24 triangular elements as depicted in Fig. 2.14 and perform uniform refinement. The resulting relative L^2 -errors as defined in Eqn. (2.116) are reported in Fig. 2.15. As can be seen, the velocity errors are rather similar for both considered stabilisations, while the pressure approximation is clearly better employing the BVS formulation, even though we employ the recommended $\alpha = 1$. The observed convergence rates are the same in the asymptotic range, but prior to that, the BVS method achieves increased accuracy especially at the domain boundary, where no boundary layers were considered in this numerical test.




Fig. 2.14: Carreau channel flow: coarsest mesh of P_1P_1 elements uniformly refined in the course of the refinement study.

The observed similarity in the velocity error, however, does only hold, if the stabilisation parameter is chosen in the optimal range. This is demonstrated by varying α while keeping the grid fixed at 4 levels of uniform refinement: resulting errors are shown in Fig. 2.16. In analogy to the Newtonian case, we note that a large α used in the PSPG formulation tends to over-relax the incompressibility constraint due to excessive weight




Fig. 2.15: Carreau channel flow: relative errors in the pressure (left) and velocity (right) using P_1P_1 elements stabilised via BVS or PSPG methods.

put on the stabilising residual terms, whereas in the BVS method, using $\alpha \rightarrow \infty$ replaces the continuity equation with the fully consistent, but lower-order accurate PPE. For optimal α , the divergence-free constraint dominates over the stabilisation terms, and, as is to be expected, choosing α too small triggers instabilities as we approach the LLB-stability limit. Note also that the error is increasing when α is chosen too high, as, e.g., going from $\alpha = 0.1$ to $\alpha = 10$ —which actually covers the optimal α —the PSPG method experiences a drastic increase by a factor of 10 in the velocity error, whereas the same norm increases by merely 12% adopting the BVS formulation.




Fig. 2.16: Carreau channel flow: relative pressure (left) and velocity (right) norms with scaling parameter $\alpha \in [10^{-5}, 10^5]$. The optimal parameter range for BVS is wider and the errors up to ≈ 1.5 orders of magnitude lower compared to the PSPG method.

In a nutshell, we see exactly the same trends as in the Newtonian case. In addition to the presented convergence studies, however, this simple setup also allows demonstrating how spurious pressure boundary layers emerge when using standard residual-based sta-

bilisations and lower order elements. This can be beautifully shown starting from the stationary momentum equation ($\partial_t \mathbf{u} = 0$), which simplifies to

$$\nabla p = \mu \Delta \mathbf{u} + \nabla \mathbf{u} \nabla \mu,$$

since we have $\mathbf{b} = \mathbf{0}$, $\nabla \cdot \mathbf{u} = 0$, but also $u_2 = 0$, $\partial_x u_1 = 0$ and $\partial_x \mu = 0$, meaning, the solution does not vary along the channel's longitudinal axis. Then, dividing by $\mu \neq 0$ and taking the divergence, we get

$$\nabla \cdot (1/\mu \nabla p) = \Delta (\nabla \cdot \mathbf{u}) + \partial_x (1/\mu \partial_y u_1 \partial_y \mu) = 0,$$

which is equivalent to requiring

$$\nabla \cdot \mathbf{u} - \tau \nabla \cdot (1/\mu \nabla p) = 0 \tag{2.119}$$

for any arbitrary $\tau \neq 0$. Now, let us compare this with a standard residual-based stabilisation, exemplarily taking PSPG (2.34), where one ends up with

$$\langle q_h, \nabla \cdot \mathbf{u}_h \rangle + \sum_{e=1}^{N_e} \langle \tau_e \nabla q_h, \nabla p_h - \nabla \cdot \mathbf{S}_h \rangle_{\Omega_e}.$$

Considering now a uniform mesh of linear triangular elements (or linear rectangular elements), we additionally have $\nabla \cdot \mathbf{S}_h = \mathbf{0}$ and we can rewrite the stabilised continuity equation using $h = h_e$ and inserting the stabilisation parameter $\tau_e \leftarrow \alpha \tau_e$ based on τ_e (2.35) with $\alpha_1 = 1$ and $\alpha_2 = \alpha_3 = 0$ as

$$\langle q_h, \nabla \cdot \mathbf{u}_h \rangle + \alpha h^2 \langle 1/\mu \nabla q_h, \nabla p_h \rangle. \tag{2.120}$$

This, however, just happens to be the weak form of

$$\nabla \cdot \mathbf{u} - \alpha_1 h^2 \nabla \cdot (1/\mu \nabla p) = 0 \quad \text{in } \Omega, \tag{2.121}$$

$$\mathbf{n} \cdot \nabla p = 0 \quad \text{on } \partial\Omega, \tag{2.122}$$

which unfortunately leads to (2.120) implicitly enforcing a zero Neumann boundary condition on the pressure (2.122) as an unwanted side effect. So, even though Eqn. (2.121) is fully consistent with (2.119), that is, no terms vanish, we suffer from spurious pressure boundary layers increasing the error at the boundary. Naturally, this effect is increased in strength as the stabilisation parameter is chosen larger [47]. This introduces a problem as the stabilisation parameter has to be selected in a large enough range to counteract inf-sup instabilities, but at the same time still small enough to not spoil accuracy in the pressure close to boundaries, which ultimately leads to a limited parameter range as already observed in the numerical tests. Closer inspection of Eqn. (2.121) also reveals that a small enough h , meaning a fine enough mesh, also diminishes the effect. In this

latter case, the polluting zero Neumann boundary condition has a reduced influence on the pressure. In practical applications, boundary layers in the finite element grid can thus help reduce these numerical artifacts to a limited extent.

In our current example, we can showcase this spurious pressure boundary condition by plotting the pressure along the channel's centerline, $y = 0$, for increasing values of α as depicted in Fig. 2.17. At the in- and outlet, the pressure clearly shows a zero gradient, affecting a wider region when using larger α , but decent approximations also for PSPG, if the optimal values are selected. The BVS scheme, on the other side, does not suffer from this issue, since the pressure obtained from the stabilised PPE features fully consistent boundary conditions.

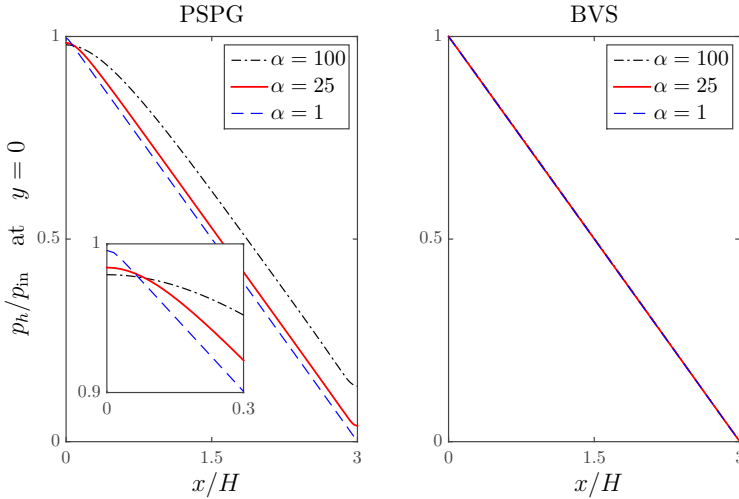
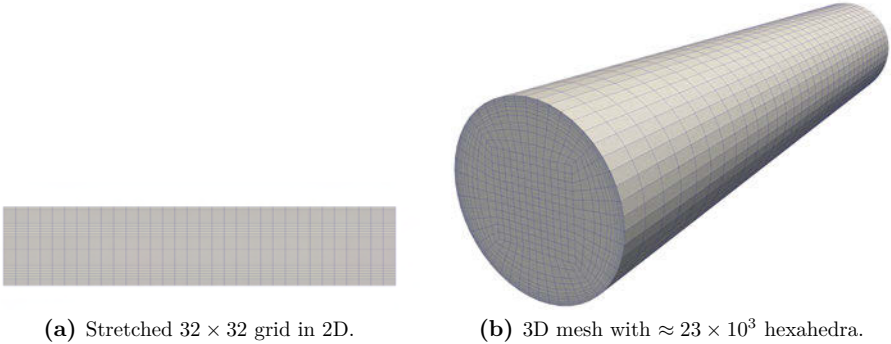


Fig. 2.17: Carreau channel flow: pressure at the channel's center line ($y = 0$) for increasing α . Spurious zero Neumann boundary conditions on the pressure p are induced by PSPG (left), while BVS (right) is not affected.

2.7.6 Time step selection in Carreau channel flow

Now that we have thoroughly investigated the fundamental behaviour of the BVS scheme compared to the PSPG method for Newtonian and generalised Newtonian fluids in stationary and time-dependent flows, also testing convergence rates in space and time, let us turn our attention to adaptive time step selection. We consider straight channels with a length of $L = 50$ mm and a radius of $R = 5$ mm and reasonable finite element grids in two and three space dimensions as depicted in Fig. 2.18.

No-slip conditions, $\mathbf{u} = \mathbf{0}$, are enforced at the cylinder walls, $r = R$, and zero pseudo traction at the outlet, that is, $\tilde{\mathbf{t}} = \mathbf{0}$ at $x = L$. We differentiate two scenarios through the parabolic inflow profile's temporal scaling: in the first case, quasi-stationary flow up to a desired constant volumetric flow rate is approached (pseudo-timestepping), while

(a) Stretched 32×32 grid in 2D.(b) 3D mesh with $\approx 23 \times 10^3$ hexahedra.**Fig. 2.18:** Meshes considered for the Carreau channel flow with adaptive time step selection.

in a second, periodic inflow case, the time step is adapted simply to capture the solution's transient behaviour more efficiently. In both of these settings, the inlet velocity is smoothly ramped from the quiescent state, $u_1|_{t=0} = 0$, via

$$u_1(r, t) = \eta_{u_1} \hat{u}_1 (1 - r^2/R^2) \quad (2.123)$$

$$\text{with } \eta_{u_1} = \eta_r \eta_p, \quad \eta_r = \begin{cases} \sin^2\left(\frac{\pi t}{2T_r}\right) & \text{if } t < T_r, \\ 1 & \text{otherwise,} \end{cases} \quad (2.124)$$

$$\text{and } \eta_p = 1 + \alpha_p \left[\cos^{10}\left(\frac{t - T_r}{T_p} \pi\right) - 1 \right]. \quad (2.125)$$

We start with an initial implicit Euler step and $\Delta t^0 = 10^{-5}$ s, to circumvent the construction of a consistent initial pressure field at $t = 0$. This would not be strictly necessary, though, as (i) we have a smoothly ramped inflow with $\partial_t u|_{t=0} = 0$ (2.124) at the inlet, and (ii) the fluid is accelerated from the quiescent state, i.e., $\mathbf{u}_0 = \mathbf{0}$, $\partial_t \mathbf{u} = \mathbf{0}$ at $t = 0$ in all of Ω , and (iii) the body force \mathbf{b} is zero at $t = 0$, since we have $\mathbf{b} = \mathbf{0}$ at all times. Thus, from the strong form of the momentum balance, one directly obtains $\nabla p = \mathbf{0}$ and together with the outflow condition boiling down to $p = 0$ at $x = L$, we have $p = 0$ at $t = 0$ as the initial, fully consistent pressure field.

Regarding the fluid's rheological parameters, we use upper and lower viscosity limits of $\eta_0 = 250$ mPa s, $\eta_\infty = 3.5$ mPa s, a density of $\rho = 1060$ kg/m³ and further set $b = 0.25$ and $\lambda = 25$ s. The velocity linearisation based on second-order accurate extrapolation (2.82) is employed, decoupling the viscosity completely from the velocity-pressure system as discussed thoroughly in Sec. 2.5. The velocity-pressure system is solved via the preconditioned FGMRES solver using a (rather crude) relative convergence criterion of 10^{-4} , which is admissible here, since we use \mathbf{u}^n and p^n as initial guess and rather small time step sizes. For all shown examples, we additionally consider a safety factor in Eqns. (2.92) and (2.115) for adaptive time step selection as $\xi = 0.98$ and specify the targeted L^2 -error on the velocity ϵ_u as needed.

Quasi-stationary solution Consider first the two-dimensional setup, where we generate a quasi-stationary solution by smoothly ramping up the inflow from zero to a stationary value of $\hat{u}_1 = 0.01$ m/s within the first $T_r = 0.5$ s and set $\alpha_p = 0$, skipping the periodic part. Then, the time step selection merely controls, how fast the steady state is reached in a pseudo-timestepping scheme. We control the L^2 -norm of the velocity truncation error by a tolerance of $\epsilon_u = 10^{-3}$, keeping the safety factor at $\xi = 0.98$, but do not recompute any steps, even if the solver does not converge or a reduction in the pseudo time step size of more than 30% occurred.

Let us first demonstrate the stabilising effect of Rannacher (R) stepping compared to Crank–Nicolson (CN), first of which adds intermediate implicit Euler steps every N^* steps, as discussed in Sec. 2.6. To rule out any possible sources of error or temporal instability coming from the inf-sup stabilisation, we will employ Q_2Q_1 elements without any stabilisation. Then, combine the aforementioned time integration methods within a predictor-corrector approach with an explicit Adams–Bashforth method of second order, to obtain what is herein referred to as the R-AB2 and CN-AB2 schemes, respectively. So, we can introduce slight numerical damping via $N^* = 10$ in R-AB2 and already see a drastic effect on the time step sizes: while the CN-AB2 scheme does not surpass a time step size of $\Delta t^n \approx 0.018$ s due to the ringing phenomenon (oscillating around this time step size), the R-AB2 scheme reaches $\Delta t^n \approx 0.82$ s within the interval $t \in (0, 0.1$ s]. Notice here in particular that the time step size shown in Fig. 2.19 increases after each stabilisation step (intermediate Euler step) of the R-AB2 scheme, whereas the CN-AB2 stalls at a certain level.

Regarding the iteration counts of the linear solvers, we observe that both the velocity-pressure system’s iteration count $N_{u,p}$ and viscosity projection’s preconditioned CG iterations N_μ are nicely bounded in the range of 5 to 15 steps after the initialisation and ramp up phase. In Fig. 2.19, the corresponding graphs are depicted as well, from which we can also see that the iteration counts vary due to variable quality in the initial guesses, but even a 10 times higher time step size reached with R-AB2 stepping does not noticeably deteriorate the preconditioner’s quality due to the small problem at hand and dominant reaction terms for all time step sizes.

Thus, CN-AB2 is not a very attractive option, since one does not have a chance reaching the theoretical limit $\Delta t^n = \infty$, which is why we will not consider this option in the following numerical tests. On the other side, however, the fractional-step θ -scheme (FS) with its embedded method can be used in such a pseudo-timestepping approach as well. We change the setup slightly by considering a much longer time period compared to the first initial test, using $t \in (0, 200]$ now, and employ Q_1Q_1 elements in the BVS formulation. The linear solvers are limited to 200 iterations, after which the last iterate is simply taken as the solution—an acceptable technique in pseudo-timestepping, since we are only interested in the quasi-stationary state reached at the end of some pseudo time interval.

In Fig. 2.20, time step sizes and iteration counts are summarised. Until the target time $T = 200$ s is reached, the time step size is increased fast and cut back regularly owing to aggressive tolerance settings for pseudo-timestepping. In any of the time integration schemes, non-convergence within 200 iterations of the linear solver applied to the

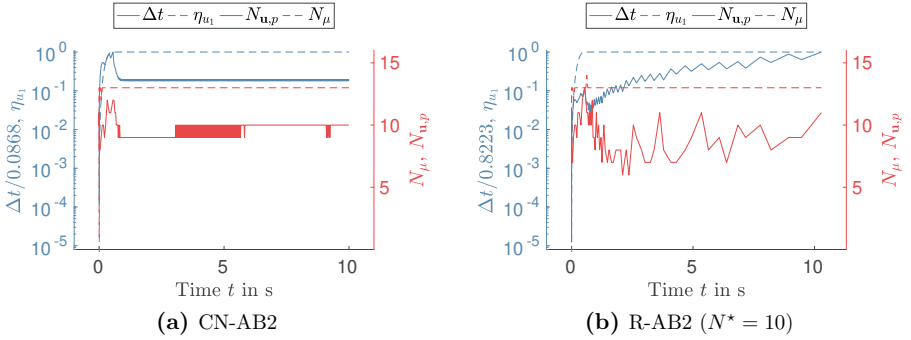


Fig. 2.19: Quasi-stationary Carreau channel flow: temporal inlet velocity scale η_{u_1} , scaled time step size Δt (left axis, blue) and linear solver iterations of velocity-pressure ($N_{u,p}$) and viscosity (N_μ) systems (right axis, red), all using Q_2Q_1 elements. The ringing phenomenon hinders the CN-AB2 integrator from reaching bigger Δt , while the stabilisation allows R-AB2 to surpass such a limit.

velocity-pressure system leads to an increased error estimate, causing a reduction of the time step size. In the following steps, the linear systems are solved adequately again, until the time step grows, potentially leading to a worsened initial guess (if the stationary solution is not yet reached) and more importantly: smaller reaction terms. This cycle of growth, divergence, cutting back the time step size and growth repeats itself, until the targeted time $T = 200$ s is passed, i.e., $t^n > T$ for R-AB2 or $t^m > T$ for the FS scheme, respectively.

Depending on how fast the time step can grow back to this limiting value, we reach T within ≈ 150 or ≈ 270 steps with the R-AB2 and FS scheme, respectively. These step counts refer to the total number of steps of a generalised θ -scheme, meaning that the FS scheme actually needs $273/3 = 91$ macro time steps. However, since the macro step size within the FS scheme is only adapted every 3 substeps, we see a higher number of total steps being used, even if the maximum macro step size is ≈ 23.04 s, so, on average $23/3 \approx 7.67$ s per substep, compared to ≈ 4.77 s reached with the R-AB2 integrator—a 60% increase from the R-AB2 scheme. This additionally leads to the fact that we reach the stability limit regarding linear solver convergence more often, causing the solver to reduce the time step size each time. For the R-AB2 integrator, on the other side, an increase in the step size can clearly be seen every N^* -th R-AB2 step with oscillatory time step sizes between.

Summing up, the FS scheme adapts slower to the target tolerance, but results in smoother transitions overall and larger maximum fractional step sizes, but the R-AB2 scheme bypasses the time step limit seen in CN-AB2 and allows adapting the step size each time step. If the problem were “more” nonlinear, we might choose a smaller tolerance ϵ_u or introduce a maximum step size and not reach 200 linear iterations at all. Judging solely from the total pseudo time steps needed, R-AB2 is superior to FS in this test.

Another interesting detail worth commenting is how considering the generalised $\tilde{\mathbf{F}}_p$, given in Eqn. (2.87), compares to the standard \mathbf{F}_p (2.86) as commonly used for Newtonian

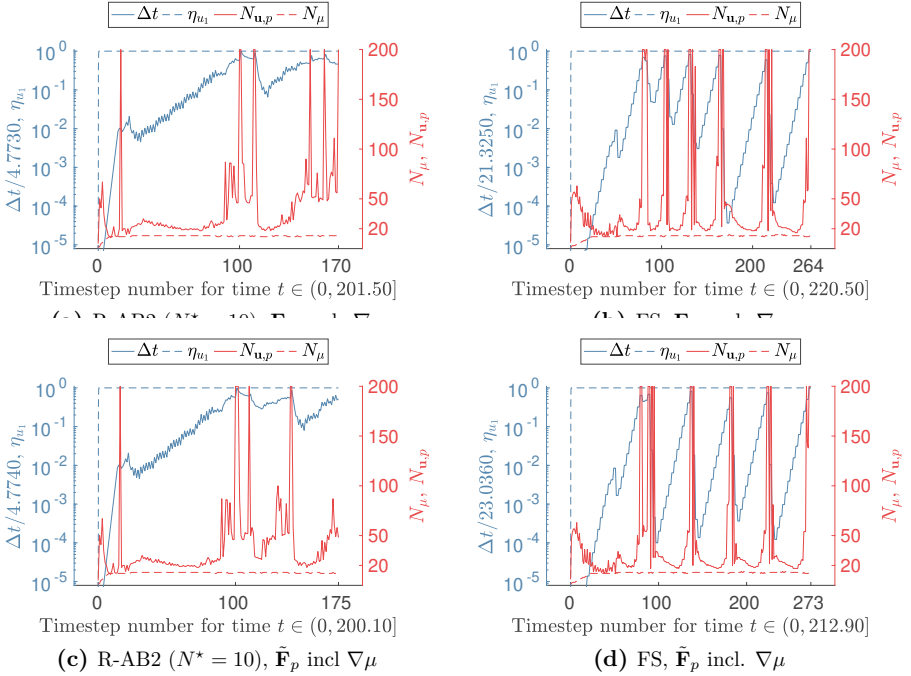


Fig. 2.20: Quasi-stationary Carreau channel flow: inlet velocity scale, η_{u_1} , scaled (macro) time step size (left axis, blue) and linear solver iterations of velocity-pressure ($N_{u,p}$) and viscosity (N_{μ}) solves (right axis, red), all using Q_1Q_1 elements and the BVS formulation. Step size is regularly cut back at linear solver divergence and recovers over time, which happens faster for the reckless R-AB2. Top and bottom rows use different operators \mathbf{F}_p (2.86) or $\tilde{\mathbf{F}}_p$ (2.87) in the PCD preconditioner.

Tab. 2.4: Quasi-stationary Carreau channel flow: number of time steps and iteration counts, format: total (**mean**), for R-AB2 and FS integrators with standard PCD preconditioner neglecting $\nabla\mu$ and a generalised variant.

	time steps N_t	$N_{u,p}$	relative $N_{u,p}$
\mathbf{F}_p excl. $\nabla\mu$, R-AB2	171	7157 (41.85)	1.02 (1.05)
\mathbf{F}_p excl. $\nabla\mu$, FS	265	12314 (46.47)	1.75 (1.16)
$\tilde{\mathbf{F}}_p$ incl. $\nabla\mu$, R-AB2	176	7032 (39.95)	1.00 (1.00)
$\tilde{\mathbf{F}}_p$ incl. $\nabla\mu$, FS	274	12212 (44.57)	1.74 (1.12)

fluids (see the discussion in Sec. 2.5 for further details). These two options are included in Fig. 2.20, which for $\chi = 0$ only differ in the viscosity gradient term, $\nabla\mu$, acting similar to a convective velocity. Comparing the top and bottom rows in Fig. 2.20, only mild differences in the iteration counts are observed for both the FS and R-AB2 schemes, showing that the adapted $\tilde{\mathbf{F}}_p$ does not improve convergence significantly.

Comparing the individual time integration schemes, we measure that the FS leads to an increase of $\approx 12\%$ in the mean iteration count needed for the velocity-pressure system solve compared to R-AB2. This can be traced back to the higher maximum step size for FS, corresponding to a less dominant reaction term, and consequently, worsened efficiency of the PCD preconditioner. We end the discussion on this pseudo-timestepping example with Tab. 2.4, which summarises the total and relative (i.e., per solve or time step) iteration counts. Tab. 2.4 also highlights the difference in executed time steps, but on the other side only mild differences per solve for both time integration schemes (FS and R-AB2) and the preconditioner variants employing \mathbf{F}_p and $\tilde{\mathbf{F}}_p$.

Periodic Carreau channel flow Now, we switch to three space dimensions and a discretisation with 23×10^3 hexahedra as depicted in Fig. 2.18, tackling a time-dependent problem by setting $\hat{u}_1 = 0.05$ m/s, $T_r = 0.5$ s, $\alpha_p = 0.25$ and $T_p = 1.0$ s. For the resulting periodic pipe flow of Carreau fluid in the interval from $t = 0$ to $T = 3$ s, we have a Reynolds number computed as in the Newtonian case of $\text{Re} = \rho u_{\text{mean}} 2R / \eta_\infty \approx 75$. Again, we employ $Q_1 Q_1$ interpolation stabilised via the BVS method and set the tolerance on the L^2 -norm of the velocity to $\epsilon_u = 10^{-4}$. Conversely to the previously presented quasi-static problem, time steps that lead to a time step size reduction of more than 30% are recomputed with halved (macro) step size.

Fig. 2.21 summarises the results obtained for the FS and R-AB2 schemes, where the latter introduces intermediate Euler steps every $N^* = 15$ time steps. Again, we show iteration counts for all substeps of the FS scheme, while the time step size indicated refers to the (macro) time step size. Clearly, the inflow scale η_{u_1} and resulting drastic change in the solution causes a reduction in the time step size as soon as the velocity is increased, while in the phases with constant inflow, the step size is increased by factors of up to 20. Regarding the step size variations, one can observe that the R-AB2 scheme leads to stronger fluctuations, while the FS scheme delivers a much smoother transition.

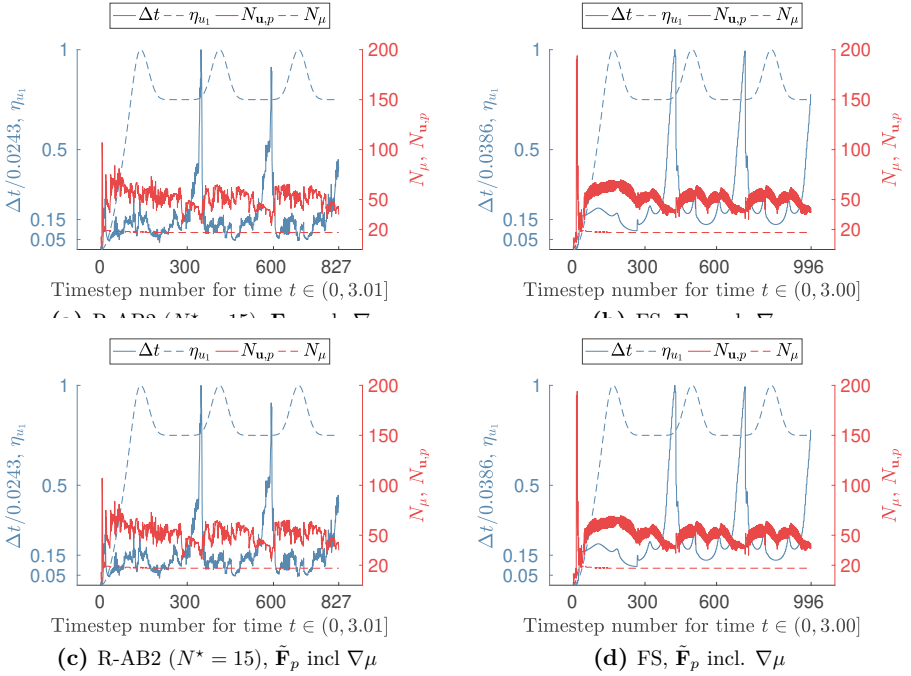


Fig. 2.21: Periodic Carreau channel flow: inlet velocity scale, η_{u_1} , scaled (macro) time step size (left axis, blue) and linear solver iterations of velocity-pressure ($N_{u,p}$) and viscosity (N_μ) solves (right axis, red), all using Q_1Q_1 elements and the BVS formulation. Both R-AB2 (left) and FS (right) show a decrease in the time step size by a factor of 10 or 20 during peak flow, while $\tilde{\mathbf{F}}_p$ and \mathbf{F}_p PCD variants perform the same.

The tolerance of $\epsilon_u = 10^{-4}$ results in rather large time steps when the solution is almost constant, where the R-AB2 scheme reaches a maximum step size of $\Delta t = 2.43 \times 10^{-2}$ s, and the FS method—due to its slower adaptation every third time step only—achieves a roughly 47% lower maximum time step size of $\Delta t \approx 3.86/3 \times 10^{-2} \approx 1.29 \times 10^{-2}$ s.

Comparing the two timestepping schemes, the FS scheme needs 996 time steps to reach $T = 3$ s, whereas the R-AB2 scheme needs 828 steps. Note that depending on the specific inflow scaling, either one of the two schemes can result in less time steps being necessary. The more the inflow varies in time, or, in a general sense, the more the solution changes from time step to time step, the more likely it is that the R-AB2 scheme requires fewer time steps overall, given that it adapts the time step size more often. If long enough periods of time feature almost constant solutions, the FS scheme can deliver bigger time steps as seen in the previous example. In addition to that, setting an appropriate tolerance ϵ_u or imposing a maximum time step size or CFL number has a large influence on which scheme turns out to be superior as well.

In any case, we have constant N_μ and nicely bounded $N_{u,p}$, latter of which depend mildly on the time step size and resulting reaction term. During the ramp up phase, a vastly

changing solution and worse initial guess together with the relative convergence criterion being harder to fulfil due to the increased influence of round-off errors leads to increased iteration counts. However, hardly any steps are repeated because of non-convergence within 200 iterations in the PCD-preconditioned FGMRES solver. The adapted $\tilde{\mathbf{F}}_p$ taking the viscosity gradient into account within the operator on the pressure space does not influence the iteration count noticeably and is thus seen as rather irrelevant for time-dependent problems. In general, smaller tolerances would lead to smoother variations in the time step size and lower iteration counts, whereas these experiments presented here aim to test the solvers performance in a challenging scenario with large variations in the time step and high viscosity gradients.

To sum up, we report the same trends as Silvester et al. [123] and Elman et al. [107]: for (pseudo) time-dependent problems, the PCD preconditioner performs well, i.e., yields acceptable iteration counts for a large range of Reynolds numbers and stretched grids. The most important aspect, though, is unfortunately the choice of a small enough time step, such that the reaction term corresponding to a scaled mass matrix is large enough.

2.7.7 Flow through an idealised aneurysm

In this final numerical test, we investigate the solver’s performance in a more realistic setting, that is, blood flow through an idealised cerebral aneurysm in the internal carotid artery under physiological flow conditions [164] as considered, e.g., in [17]. The computational domain is constructed starting from a reference cylinder $\hat{\Omega}$ with radius 2.5 mm and unit length 1 m,

$$\hat{\Omega} := \left\{ (\hat{x}, \hat{y}, \hat{z}) \in \mathbb{R}^3 : r := \sqrt{\hat{y}^2 + \hat{z}^2} \leq 2.5 \times 10^{-3}, 0 \leq \hat{x} \leq 1 \right\},$$

“inflating” the cross-section via

$$\hat{y} \leftarrow \hat{y} \left(1 + 6 \sin^{30}(\pi \hat{x}) \right), \quad \hat{z} \leftarrow \hat{z} \left(1 + 4 \sin^{30}(\pi \hat{x}) \right),$$

and applying an additional map

$$\mathbf{x}(\hat{\mathbf{x}}) : \mathbb{R}^3 \rightarrow \mathbb{R}^3, \quad \mathbf{x}(\hat{\mathbf{x}}) := \begin{pmatrix} 0.16 \hat{x} \\ 0.0015 \cos 5\hat{x} \\ 0.009 \sin 7\hat{x} \end{pmatrix}.$$

This leads to the idealised geometry depicted in Fig. 2.22, where the cross-sections at the in- and outlet remain circular, but the radii expand smoothly from 2.5 mm to $a = 17.5$ mm, meaning, up to a seven-fold increase in diameter at maximum. The finite element grid consists of $\approx 174 \times 10^3$ hexahedra, $\approx 180 \times 10^3$ nodes and consequently $\approx 901 \times 10^3$ unknowns using Q_1Q_1 elements. Using the LBB-stable Q_2Q_1 pair based on the same number of hexahedra, one would end up with $\approx 46 \times 10^5$ DoFs—an increase by a factor of five.

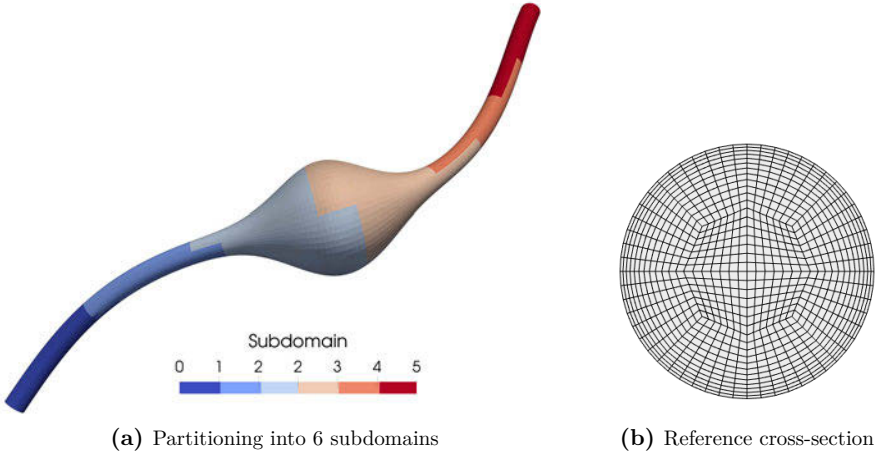


Fig. 2.22: Blood flow in idealised aneurysm: discretisation using $\approx 174 \times 10^3$ hexahedra.

Regarding boundary conditions, we enforce zero velocities $\mathbf{u} = \mathbf{0}$ on the vessel wall, while zero pseudo tractions, i.e., $\tilde{\mathbf{t}} = \mathbf{0}$ are prescribed at the outlet. A more realistic scenario would include models of the downstream vasculature, accounting for its resistance and capacitance in order to enforce a variable and physiologically meaningful pressure level, see Sec. 4.6. Again, the quiescent state $\mathbf{u}|_{t=0} = \mathbf{0}$ is chosen as initial condition, smoothly ramping up the velocity profile in inlet normal direction very much similar to Eqns. (2.123)–(2.124) via

$$\mathbf{u}(t, r) = -\mathbf{n} \eta_{u_1} \hat{u}_1 (1 - r^2/R^2), \quad (2.126)$$

with the unit outward normal \mathbf{n} and radius r on the inlet face, but replace the time-periodic scale η_p by a truncated Fourier series,

$$\eta_p = 0.6658 + 1/50 \sum_{k=1}^7 \left[a_k \cos\left(2\pi k \frac{t}{T_p}\right) + b_k \sin\left(2\pi k \frac{t}{T_p}\right) \right].$$

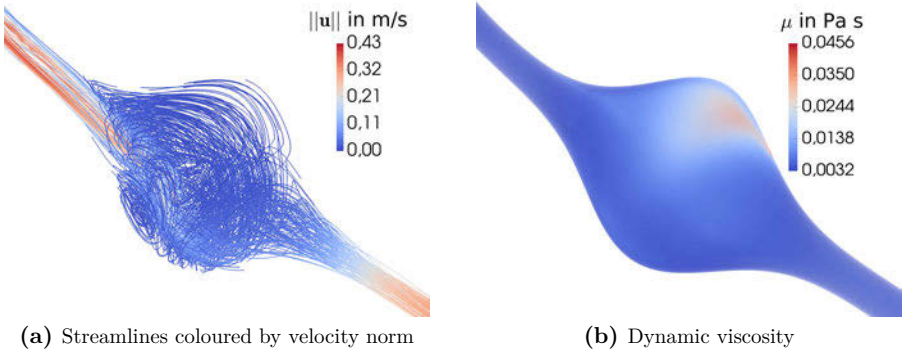
Choosing the Fourier coefficients a_k and b_k according to Tab. 2.5, additionally setting $T_p = 0.917$ s, $T_r = T_p/2$ and targeting a peak inflow velocity of $\hat{u}_1 \approx 0.5$ m/s completes the spatial and temporal scaling of the inlet velocity profile based on measurements by Thiriet [164]. Overall, three cardiac cycles are considered in the time interval starting from $t = 0$ to $T = 3$ s.

The rheological parameters to model the shear-thinning behaviour of blood are selected as $\eta_0 = 45$ mPa s, $\eta_\infty = 3.2$ mPa s, $\lambda = 10.03$ s and $b = 0.344$, following Gambaruto et al. [165] together with $\rho = 1060$ kg/m³. Taking the in- and outlet radius and the viscosity in the Newtonian limit as reference, this yields a Reynolds number of $\text{Re} = \rho \hat{u}_1 R / \eta_\infty \approx 414$. Thus we expect a laminar, but complex flow pattern with recirculations and regions with stagnant flow, triggering large gradients in the shear rate and apparent fluid viscosity. These effects are caused by the combination of the present geometry and the periodic

Tab. 2.5: Fourier coefficients for maximum inflow velocity, taken from [164].

k	1	2	3	4	5	6	7
a_k	-5.27	-5.10	-2.81	-2.15	-0.12	0.79	0.27
b_k	9.52	0.72	0.29	-3.25	-0.95	-1.17	-0.50

inflow scaling. A snapshot of a representative state at the third and final peak systole at $t \approx 2$ s is shown in Fig. 2.23, which highlights pronounced circulatory flow patterns. The streamlines are colored by the velocity magnitude and indicate vast differences in speed and direction of flow, while the corresponding viscosity spans the whole admissible range given the large shear rates.

**Fig. 2.23:** Blood flow in idealised aneurysm: recirculating flow in aneurysm at peak systole, i.e., $t \approx 2$ s, triggering large viscosity gradients.

However, we want to emphasise at this point that we are not interested in phenomenological effects *per se*, but rather in the numerical treatment and algorithms that can cope with the occurring phenomena. Thus, the solution—as fascinating as it is—is not of greater importance within this work, but a critical investigation of the solver and its components is.

So, focusing on the solver, shortly summarise the settings, before going into further details. In space, Q_1Q_1 elements in the BVS formulation are employed, while in time, we compare the FS and R-AB2 ($N^* = 15$) schemes. The tolerance on the L^2 -norm of the truncation error used in adaptive timestepping is set to $\epsilon_u = 5 \times 10^{-3}$ for both the FS and the R-AB2 integrators, which use an initial implicit Euler step with a length of $\Delta t^0 = 10^{-5}$ s. Safety measures are switched on, meaning the safety factor $\xi = 0.98$ is employed and (macro) steps are repeated with halved step size, when divergence of the (non-)linear solver occurs or a step size reduction of more than 30% is suggested by the time step selection scheme.

The velocity-pressure and viscosity subsystems are decoupled based on a second-order accurate extrapolation of the velocity. We focus on a comparison of (i) the basic variant

resolving the nonlinearities via a Aitken-accelerated Picard scheme, solving the velocity-pressure and viscosity-projection systems repeatedly with (ii) the linearised and fully decoupled algorithm, solving each of these systems once per time step. The initial guess for the linear solvers is taken as the last time step's solution or the last iterate in the nonlinear solver, whenever applicable. Note here that it is not to be confused with the velocity considered for linearisation and decoupling the velocity-pressure and viscosity subsystems. We aim for an error reduction by 10^4 in the linear solver within 200 iterations in the linearised scheme, while when employing the nonlinear solver, we get away with a higher relative tolerance of 10^{-2} in the linear solvers, while still reducing the nonlinear solver's residual by a factor of 10^3 .

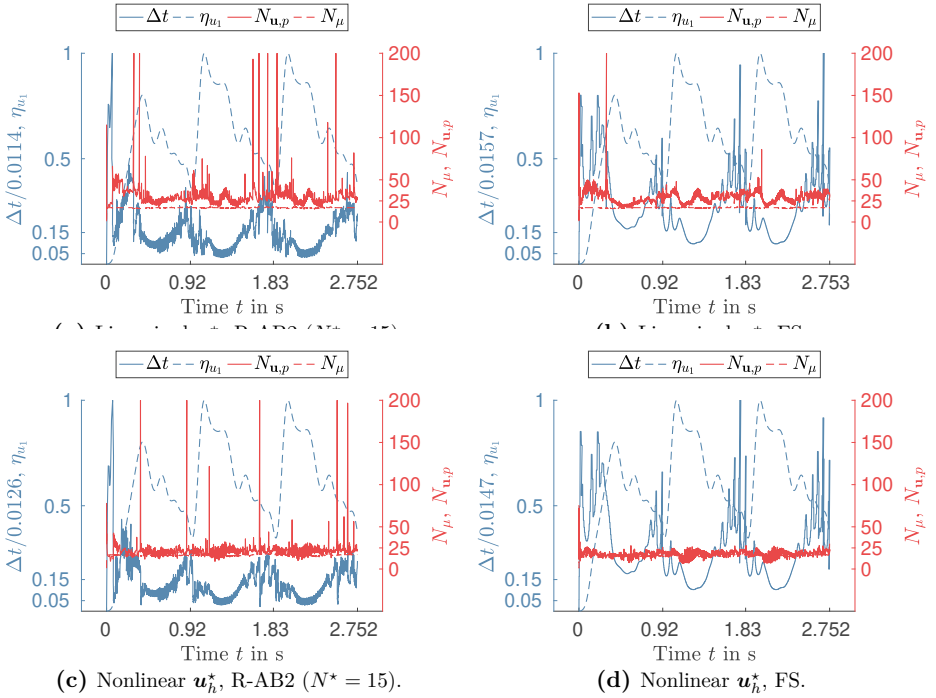


Fig. 2.24: Blood flow in idealised aneurysm: adaptive time stepping and inflow scale η_{u_1} (left axis, blue) and linear solver iterations (right axis, red) using linearised variant (top row) or Aitken-accelerated Picard scheme (bottom row).

In Fig. 2.24, the resulting time step sizes with the inflow velocity scale η_{u_1} and (mean) iteration counts (over all nonlinear steps of the time step) needed to solve the velocity-pressure ($N_{u,p}$) and viscosity (N_μ) subsystems up to the desired accuracy are presented. The total number of executed time steps until three cycles are completed are almost identical: (2353) 2335 for the (nonlinear) R-AB2 option compared to (2553) 2567 for the (nonlinear) FS integrator. In the present scenario, both schemes manage to adapt to the rapidly varying inflow, independent of the linearised (top row) or nonlinear (bottom row) variants of the solver being used. During peak systole, the time step decreases,

reaching similar levels for both integrators, and increases as the flow rate is reduced. Since the diastolic phase is not as pronounced as in the previously discussed numerical tests considering periodic flow, we do not observe a plateau being reached, but rather a temporary increase in step size within a short period of time. After the initialisation phase, the minimum step size for the R-AB2 scheme is ≈ 0.5 ms, whereas the FS scheme ends up having a minimum macro step size of ≈ 1.5 ms, meaning the individual steps are of comparable size, while also the overall patterns match well. The predictor-corrector scheme, however, shows oscillations directly related to N^* , which are also much more pronounced compared to the fractional-step θ -scheme. The iterations in the linear solvers are slightly higher for the linearised scheme, owing to the tighter tolerance criterion. Also, divergence is observed less often for the nonlinear variants, which one might expect. Interestingly, R-AB2 leads to more failed solution attempts compared to FS scheme possibly caused by FS leading to smoother transitions of the time step size and related improved quality of the extrapolations.

Moreover, drastic changes in the solution due to rapidly changing inflow rates induce higher linear and also nonlinear iteration counts especially when $\partial_t \eta_{u_1} \approx 0$, caused by a decreasing quality of the initial guess when the *trend* in the solution changes. As a matter of fact, using a second-order accurate (i.e., linear) extrapolation as initial guess lead to an increase in divergent steps, which is why we settled here for the more secure option. That is, we take $\mathbf{u}_h^* = \mathbf{u}_h^n$ as the initial guess for the nonlinear solver, which is not possible in the linearised and fully decoupled variant, as it would limit the scheme to first order temporal accuracy.

For a better comparison of the linear and nonlinear solver variants, iteration counts, (non-)linear steps and time steps needed to complete the three cardiac cycles are given in Tab. 2.6 together with the mean time per velocity-pressure subsystem solve and the resulting relative computing time. The linearisation pays off in the example considered—performing the viscosity projection step and velocity-pressure subsystem solve only once per time step with tighter tolerances yields slightly higher linear iteration counts, but spares us performing multiple iterations in the fixed-point iteration. In the nonlinear variants, we thus end up performing more iterations in the linear solvers overall. Despite the fact that we solve a single system up to twice as fast, we need more than two nonlinear iterations on average, which results in a speed-up of up to 3 when using the linearised variants.

Comparing the R-AB2 and FS timestepping schemes, we see that they perform equally well in terms of steps needed and mean (non-)linear iteration counts. The smoother variation in the time step size and resulting better initial guess in the FS scheme seems to make up for the less aggressive adaptation compared to the R-AB approach, such that the number of steps does not differ significantly—contrary to the previously presented examples. So, summing up, the use of linearised schemes can pay off considerably and does neither influence accuracy nor stability too much. This, however, turns out to depend on the physical parameters in the problem at hand, the chosen tolerance ϵ_u and also the spatial discretisation, which possibly limits the admissible time step size via a CFL condition.

Tab. 2.6: Performance comparison of (non-)linear R-AB2 ($N^* = 15$) and FS schemes: time steps N_t , system solves, linear solver iteration count in velocity-pressure and viscosity projection steps (total in thousands and mean in brackets), mean time spent in the velocity-pressure system solve and relative computing time.

	R-AB2 (lin.)	FS (lin.)	R-AB2 (nonlin.)	FS (nonlin.)
time steps N_t	2336	2568	2354	2554
system solves	2.34 (1.00)	2.57 (1.00)	8.53 (3.62)	8.68 (3.40)
$N_{u,p}$	68.43 (29.29)	73.00 (28.43)	190.97 (22.40)	159.97 (18.44)
N_μ	39.32 (16.83)	43.14 (16.90)	138.71 (16.27)	138.36 (15.95)
time/ u - p solve	39.95 s	30.72 s	21.94 s	18.27 s
comp. time	1.03	1.00	2.90	2.83

2.8 Summary and conclusion

Starting from an introduction of the governing equations for, e.g., haemodynamics or polymer flow, that is, the Navier–Stokes equations for incompressible flows incorporating generalised Newtonian rheology, the preceding sections summarise the development of a stabilised velocity-pressure formulation. It builds upon a pressure Poisson equation derived from the fluid’s momentum balance, for which fully consistent boundary conditions are formulated. We also show that the PPE adapted to the case of generalised Newtonian fluids implicitly enforces the divergence-free constraint onto the velocity field and hence mass conservation is ensured.

Following this coupled solution approach, the continuity equation is replaced by an equation combining $\beta \nabla \cdot \mathbf{u}$ and the PPE, such that the saddle-point structure of the velocity-pressure system is broken. Therefore, lower equal-order interpolation is enabled and proper reformulation of the PPE allows employing standard C^0 -continuous finite elements. The required stabilisation parameter is proposed in analogy to well-established relations originally designed for the pressure-stabilised Petrov–Galerkin method and thus scales with the element size, strong form residual and physical parameters. An important aspect in this regard, however, is that the newly proposed method does not suffer from spurious pressure boundary layers as we can represent the shear stress contribution to the strong form residual even when using *linear* interpolation of the involved fields. Additionally, the stabilisation parameter’s admissible range is enlarged substantially, such that velocity and pressure errors can be reduced significantly. This advantage is rooted in the construction of the equation governing the pressure, as stabilisation parameters chosen too large merely lead to a transition from the classical divergence-free constraint to the consistent PPE, which implicitly enforces continuity by design. In the PSPG method on the other side, a stabilisation parameter chosen too large puts excessive weight on the strong form residual, triggering pressure boundary layers.

Applying the coupled velocity-pressure formulation to time-dependent and more practical problems, we further introduce several modifications of the basic scheme to improve

performance. The viscosity is decoupled from the overall system by extrapolating the velocity in (pseudo-)time, such that the fluid momentum balance equation can be linearised as well. This leads to great flexibility in the rheological law, which is easily replaced solely by adapting the right-hand side of the viscosity projection step. The velocity-pressure and viscosity steps are then solved once per time step, where the former can be suitably preconditioned adopting (and slightly modifying) ideas from classical Schur-complement-based approaches. A basic single-step formulation using a generalised θ -method then allows adapting the time step size by means of (i) a predictor-corrector approach combining Rannacher timestepping and a stabilised explicit Adams–Bashforth predictor or (ii) a fractional step θ -scheme with an embedded lower-order scheme grouping three steps of the generalised θ -method.

Several numerical experiments of academic nature demonstrate accuracy and robustness of the presented approach, showcasing the expected convergence rates and improvements compared to established schemes. On the other side, tests applying the proposed method to an idealised cerebral aneurysm highlight the scheme’s potential for practical application in haemodynamics, where the proposed linearisations, decoupling strategies, adaptive timestepping and physics-based preconditioning via standard off-the-shelf algebraic multigrid methods available as open-source scientific software yield substantial performance gains in terms of overall time steps executed and/or number of individual subproblem solves required to complete a given time interval of interest.

As already mentioned previously, the BVS shifts towards a PPE-based split-step scheme for high stabilisation parameters, which inherently enforces mass conservation and breaks the saddle-point structure. Thus, a closely related alternative to the proposed velocity-pressure formulation is a split-step scheme, which completely decouples velocity and pressure spaces. Such an approach will naturally involve further linearisations to allow solving for the fluid velocity and pressure individually. The following chapter introduces such a split-step scheme suitable for generalised Newtonian fluid flow, such that a comparison to the monolithic BVS method is rendered possible.

3 Incompressible viscous flows: split-step approach

Contrary to the coupled solution approach presented in the previous section, where velocity and pressure unknowns are retrieved by solving a single monolithic (block-)system, split-step or sometimes also called time-splitting or fractional-step schemes aim to decouple the involved physical fields. Note, however, that *fractional-step scheme* is now to be understood distinctively different than in the previous section. To avoid confusion with the fractional-step θ -scheme, we thus refer to the methods avoiding the monolithic velocity-pressure system as split-step schemes or time-splitting methods.

The main motivation for such a split lies in circumventing the monolithic system and thereby easing preconditioner design tremendously and/or enable the use of different linear solvers. The resulting subproblems usually consist of one or several advection-diffusion-reaction equations, Poisson problems and L^2 -projections, all of which are standard problems in science and engineering, such that highly efficient solution algorithms are readily available [30, 166]. Split-step methods can be further subdivided into algebraic splitting schemes, projection methods and consistent splitting schemes, where we refer the reader to the excellent works of Guermond et al. [30, 167] and the monograph by Karniadakis and Sherwin [166] for an overview and a more in-depth discussion.

Algebraic splitting schemes perform the decomposition of variables on the fully discrete level, while the other two families of methods perform the splitting in the continuous setting. The most popular schemes fall in the latter category, these so-called projection methods can be further classified into velocity-correction (see, e.g., [168–170]) and pressure-correction schemes (see, e.g., [171–175]). The fundamental ingredient in both of these variants is a Helmholtz–Leray decomposition of the velocity vector, writing the momentum balance equation as

$$\partial_t \mathbf{u} + \nabla p = \partial_t \hat{\mathbf{u}} \quad \text{with} \quad \nabla \cdot \mathbf{u} = 0,$$

where we have a solenoidal part $\partial_t \mathbf{u}$, an irrotational part ∇p and the remaining terms of the momentum equation are simply gathered in $\partial_t \hat{\mathbf{u}}$. With the help of $\nabla \cdot (\partial_t \mathbf{u}) = 0$ given sufficient regularity, we can derive a Poisson equation in the pressure and recover the velocity \mathbf{u} in a following projection step.

Unfortunately, deriving consistent, let alone higher-order accurate boundary conditions for the subproblems involved in projection schemes is far from trivial. The original method by Chorin [171], but also the incremental pressure-correction scheme [172] suffered from a nonphysical pressure close to boundaries due to inherently inconsistent

pressure boundary conditions. A remedy was presented by Timmermans et al. [175], introducing a correction term to the pressure step in the so-called rotational pressure-correction scheme. While the incremental and rotational pressure-correction schemes both achieve second-order accuracy for the velocity, Guermond et al. [167] showed that for three-dimensional problems including Neumann boundaries, even the improved rotational form runs into a contradiction: the factor scaling the divergence correction term in the pressure step should equal 1 to recover consistent pressure boundary conditions, but stability requirements dictate a factor smaller than $2/3$. Although the order of convergence is not spoiled *per se*, suppressing the loss of accuracy close to the boundaries is still a matter of ongoing research requiring further measures to be taken [176, 177].

An alternative to projection schemes was presented by Liu [21], who extended the PPE-based method of Johnston and Liu [20], allowing also for traction boundary conditions being imposed on Neumann boundaries. In Liu's method [21], the continuity equation is completely replaced by an alternative, yet, equivalent system very much similar to Eqns. (2.17)–(2.24). To obtain fully decoupled and linearised equations for both velocity and pressure, the key aspect is to additionally extrapolate the pressure and the convective velocity in the momentum balance equation in time. Thus, after solving for the velocity or its individual components separately, the pressure is computed based on the PPE, which is subject to fully consistent boundary conditions. Several variations of this scheme including higher-order methods have been devised and were shown to perform well in challenging scenarios of incompressible flow [154, 178–183].

Extending those well-established concepts towards non-homogeneous viscosity as is the case for generalised Newtonian fluids turns out to be rather complex. Recent developments include the works by Deteix and Yakoubi [184], which generalised a rotational projection scheme to account for variable viscosity, or by Plasman et al. [185], additionally allowing for natural boundary conditions. However, when considering for a shear-rate dependent viscosity, the discretised stress tensor containing velocity gradients necessary in the pressure-correction step needs to be smooth enough—an inconvenience that can be overcome by a series of projections onto a continuous space as introduced by Deteix and Yakoubi [82]. This is straight-forward and accurate, but unfortunately comes at an increased cost: besides a vector-valued advection-diffusion-reaction equation and two Poisson problems, one ends up with more than ten scalar projections to perform additions of various discrete quantities. Even if we decide to lump all mass matrices in these projection steps, there is also another problem with projection methods in general: although not as easily seen as in coupled formulations and sometimes even ignored, an inf-sup condition still applies to projection schemes. This might not have consequences if a stabilising term is present in the pressure step, but in general, one has to expect increased errors for large time steps and decreased stability for small enough time steps [30, 166, 177].

Against this background, we derive a PPE-based split-step scheme considering for generalised Newtonian fluids with fully consistent Dirichlet and Neumann boundary conditions based on the alternative system already presented in Eqns. (2.17)–(2.24). This can be interpreted as an extension of the Newtonian splitting scheme by Liu [21] combined with the generalised Newtonian PPE introduced by Pacheco and Steinbach [19]. Both a generalised Laplacian and a stress-divergence form imposing real tractions or

pseudo tractions via Neumann terms are available, while the number of steps stays at a minimum and accuracy and stability are not compromised. This flexible scheme allows employing equal-order finite elements and inf-sup stable pairs. In addition to that, the rheological law, i.e., the viscosity projection step, is effortlessly exchanged. Here, we summarise the joint work by the authors [3] introducing the basic scheme and a suitable variational formulation in Sec. 3.1, present (higher-order) accurate linearisation, extrapolation and time integration variants in Sec. 3.2 before testing the scheme in various academic, benchmark and applied scenarios in Sec. 3.5.

As a starting point for our derivations, we choose the PPE formulation as given in (2.17)–(2.24), repeating it here for the convenience of the reader:

$$\begin{aligned}
\rho[\partial_t \mathbf{u} + (\nabla \mathbf{u})\mathbf{u}] - \mu[\Delta \mathbf{u} + \chi \nabla(\nabla \cdot \mathbf{u})] - 2\nabla^s \mathbf{u} \nabla \mu + \nabla p &= \mathbf{b} && \text{in } \Omega \times (0, T], \\
-\Delta p + \nabla \cdot [2\nabla^s \mathbf{u} \nabla \mu - \rho(\nabla \mathbf{u})\mathbf{u}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu &= -\nabla \cdot \mathbf{b} && \text{in } \Omega \times [0, T], \\
\mathbf{u} &= \mathbf{g} && \text{on } \Gamma_D \times (0, T], \\
(-p\mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u}) \mathbf{n} &= \bar{\mathbf{t}} && \text{on } \Gamma_N \times (0, T], \\
\mathbf{u} &= \mathbf{u}_0 && \text{at } t = 0, \\
\nabla \cdot \mathbf{u}_0 &= 0 && \text{in } \Omega, \\
-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot [\mu(\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u}) \mathbf{n} - \bar{\mathbf{t}}] &= p && \text{on } \Gamma_N \times [0, T], \\
\mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u} - \rho(\nabla \mathbf{u})\mathbf{u} + 2\nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] &= \mathbf{n} \cdot \nabla p && \text{on } \Gamma_D \times [0, T].
\end{aligned}$$

which is equivalent to the classical velocity-pressure form of the Navier–Stokes equations for incompressible flows according to Thm. 2.1.1. This system elevates regularity requirements on both velocity and pressure, as we have already discussed when using it as a starting point for the BVS stabilisation.

Compared to the weak form of the PPE with $\beta \nabla \cdot \mathbf{u}$ added in the velocity-pressure formulation (see Sec. 2.3.3), the main difference is now that we aim to compute the pressure *given* velocity \mathbf{u} and viscosity μ . To determine a unique pressure, the essential boundary condition on Γ_N is incorporated as opposed to the BVS scheme, which only uses a natural boundary condition. In a first step, we thus focus on finding a variational formulation allowing for C^0 -continuous function spaces, easing spatial discretisation in complex three-dimensional geometries.

3.1 A weak form for C^0 -interpolation

The weak form as a basis for the split-step scheme is constructed by a sequence of steps, namely, (i) rewriting all second derivatives in the PPE, (ii) projecting the Dirichlet condition on the pressure on Γ_N , (iii) applying the divergence theorem in a standard way to the momentum balance equation and, finally, (iv) introducing the viscosity as an additional continuous field. The price to pay for replacing the explicit, seemingly trivial divergence-free constraint with a consistent and invertible PPE is thus an involved right-hand side of the discretised PPE and additional L^2 -projections of the viscosity and the pressure boundary condition in Ω and on Γ_N , respectively. In the following, we shed some light on each of these mentioned steps.

3.1.1 Pressure Poisson equation

Similar to Sec. 2.3.1, we start off by multiplying the PPE (2.18) by $q \in H^1(\Omega)$, which has a zero trace on Γ_N , $q|_{\Gamma_N} = 0$, and perform integration by parts to get

$$\langle \nabla q, \nabla p \rangle - \langle q, \mathbf{n} \cdot \nabla p \rangle_{\Gamma_D} = \langle q, \nabla \cdot [(\rho \nabla \mathbf{u})\mathbf{u} - 2\nabla^s \mathbf{u} \nabla \mu - \mathbf{b}] + [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu \rangle,$$

where we can insert

$$\mathbf{n} \cdot \nabla p = \mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u} - \rho(\nabla \mathbf{u})\mathbf{u} + 2\nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] \quad \text{on } \Gamma_D \times [0, T],$$

that is, the Neumann boundary condition on the pressure (2.24), to yield

$$\begin{aligned} & \langle \nabla q, \nabla p + \rho(\nabla \mathbf{u})\mathbf{u} - \mathbf{b} \rangle + \langle q\mathbf{n}, \rho \partial_t \mathbf{u} \rangle_{\Gamma_D} \\ &= \langle \nabla q, 2\nabla^s \mathbf{u} \nabla \mu \rangle + \langle q, [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu \rangle - \langle q\mathbf{n}, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle_{\Gamma_D}. \end{aligned} \quad (3.1)$$

The boundary term on Γ_D involving the double curl of the velocity in Eqn. (3.1) is rewritten with the help of Gauss' divergence theorem

$$\begin{aligned} \langle q\mathbf{n}, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle_{\Gamma_D} &= \langle q\mathbf{n}, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle_{\Gamma} \\ &= \langle q, \mu \nabla \cdot [\nabla \times (\nabla \times \mathbf{u})] \rangle + \langle \nabla (q\mu), \nabla \times (\nabla \times \mathbf{u}) \rangle \\ &= \langle \nabla (q\mu), \nabla \times (\nabla \times \mathbf{u}) \rangle \\ &= \langle \nabla q, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle + \langle q, \nabla \mu \cdot [\nabla \times (\nabla \times \mathbf{u})] \rangle, \end{aligned}$$

such that we are left with

$$\langle \nabla q, 2\nabla^s \mathbf{u} \nabla \mu \rangle - \langle \nabla q, \mu \nabla \times (\nabla \times \mathbf{u}) \rangle$$

on the right-hand side of Eqn. (3.1). Similar to the BVS setting, we can get rid of the second derivatives on \mathbf{u} by first integrating by parts

$$\langle \mu \nabla q, \nabla \times (\nabla \times \mathbf{u}) \rangle = \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_{\Gamma} + \langle \nabla \times (\mu \nabla q), \nabla \times \mathbf{u} \rangle,$$

and using the relation

$$\begin{aligned} \langle \nabla \times (\mu \nabla q), \nabla \times \mathbf{u} \rangle &= \langle \nabla \mu \times \nabla q + \mu \nabla \times (\nabla q), \nabla \times \mathbf{u} \rangle \\ &= \langle \nabla \mu \times \nabla q, \nabla \times \mathbf{u} \rangle \\ &\equiv \langle \nabla q, (\nabla \times \mathbf{u}) \times \nabla \mu \rangle \\ &\equiv \langle \nabla q, (\nabla \mathbf{u} - \nabla^\top \mathbf{u}) \nabla \mu \rangle, \end{aligned}$$

which is an identical sequence of steps presented during the derivation of the BVS form. Thus, we end up with a simplified weak form given by

$$\langle \nabla q, \nabla p \rangle = \langle \nabla q, \mathbf{b} - (\rho \nabla \mathbf{u})\mathbf{u} + 2(\nabla^\top \mathbf{u}) \nabla \mu \rangle + \langle \mathbf{n} \times \nabla q, \mu \nabla \times \mathbf{u} \rangle_{\Gamma} - \langle q\mathbf{n}, \rho \partial_t \mathbf{u} \rangle_{\Gamma_D}.$$

Taking a closer look at the boundary term on Γ , we can formulate Lemma 3.1.1:

Lemma 3.1.1. *For $\mathbf{v} \in [L^2(\Gamma)]^d$ and q such that $\mathbf{n} \times \nabla q \in [L^2(\Gamma)]^d$, with $q|_{\Gamma_N} = 0$, there holds*

$$\langle \mathbf{n} \times \nabla q, \mathbf{v} \rangle_{\Gamma} = \langle \mathbf{n} \times \nabla q, \mathbf{v} \rangle_{\Gamma_D}. \quad (3.2)$$

Proof. The gradient ∇q can be decomposed into a normal and a tangential part as

$$\nabla q = (\mathbf{n} \cdot \nabla q) \mathbf{n} + (\mathbf{n} \times \nabla q) \times \mathbf{n}.$$

Since we have $q = 0$ on $\Gamma_N = \Gamma \setminus \Gamma_D$, the tangential component $(\mathbf{n} \times \nabla q) \times \mathbf{n}$ is identical to zero on Γ_N . The contribution stemming from the normal component restricted to Γ_N also vanishes, since we have

$$\mathbf{n} \times [(\mathbf{n} \cdot \nabla q) \mathbf{n}] = (\mathbf{n} \cdot \nabla q) \mathbf{n} \times \mathbf{n} \equiv 0.$$

Thus, altogether we have

$$\langle \mathbf{n} \times \nabla q, \mathbf{v} \rangle_{\Gamma} = \langle \mathbf{n} \times \nabla q, \mathbf{v} \rangle_{\Gamma_D} + \langle \mathbf{n} \times \nabla q, \mathbf{v} \rangle_{\Gamma_N} = \langle \mathbf{n} \times \nabla q, \mathbf{v} \rangle_{\Gamma_D}, \quad (3.3)$$

as stated initially. \square

With this, the weak form of the PPE further simplifies to

$$\begin{aligned} \langle \nabla q, \nabla p \rangle &= \langle \nabla q, \mathbf{b} - (\rho \nabla \mathbf{u}) \mathbf{u} + 2 \left(\nabla^{\top} \mathbf{u} \right) \nabla \mu \rangle \\ &\quad + \langle \mathbf{n} \times \nabla q, \mu \nabla \times \mathbf{u} \rangle_{\Gamma_D} - \langle q \mathbf{n}, \rho \partial_t \mathbf{u} \rangle_{\Gamma_D}, \end{aligned} \quad (3.4)$$

finally free of second derivatives on the velocity \mathbf{u} . Notice, however, that Eqn. (3.4) contains first derivatives of the viscosity μ . Consequently, higher-order derivatives on \mathbf{u} enter through the back door as the rheological law typically depends on the shear rate $\dot{\gamma}(\nabla^s \mathbf{u})$ for the generalised Newtonian fluids we are interested in. Very much similar to the BVS stabilisation, we yet again overcome this issue by introducing the viscosity as an additional continuous field and recover the viscosity weakly. That is, we project the viscosity given the rheological law and velocity \mathbf{u} , searching for $\mu \in H^1(\Omega)$, such that

$$\langle r, \mu \rangle = \langle r, \eta(\dot{\gamma}(\nabla^s \mathbf{u})) \rangle \quad \forall r \in L^2(\Omega). \quad (3.5)$$

This is a simple L^2 -projection, which corresponds to a standard mass matrix solve on the discrete level and is completely identical to the viscosity projection step in the coupled scheme. Projecting the scalar viscosity or the components of the viscous tensor to circumvent higher-order regularity requirements on the velocity is a handy trick that has been employed in similar works targeting non-Newtonian fluids [19, 82].

3.1.2 Pressure Dirichlet condition

The essential boundary condition on the pressure (2.23) reads

$$p = -\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot \left[\mu \left(\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u} \right) \mathbf{n} - \bar{\mathbf{t}} \right] \quad \text{on } \Gamma_N \times [0, T].$$

Within a standard finite element setup, this boundary condition naturally enters the pressure ansatz space and thus allows us to uniquely determine the pressure based on \mathbf{u} and p to be begin with. However, the boundary data we wish to enforce is discontinuous, given the fact that it involves first derivatives of the velocity \mathbf{u} and the unit outward normal \mathbf{n} . A simple solution to this dilemma as proposed by Liu [21], and also used in the shear-rate projection method of Plasman et al. [185], is to introduce an intermediate variable ζ defined as

$$\zeta := -\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot \left[\mu \left(\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u} \right) \mathbf{n} - \bar{\mathbf{t}} \right], \quad (3.6)$$

and projecting it onto a continuous space on the boundary. That is, find $\zeta \in H^{1/2}(\Gamma_N)$, such that $\forall s \in L^2(\Gamma_N)$ there holds

$$\langle s, \zeta \rangle_{\Gamma_N} = \langle s, -\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot \left[\mu \left(\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u} \right) \mathbf{n} - \bar{\mathbf{t}} \right] \rangle_{\Gamma_N}. \quad (3.7)$$

In a discrete setting, the nodal values of the corresponding vector $\underline{\zeta}$ are then directly prescribed on the Neumann boundary. This projection is negligible in terms of computational cost, as it is trivially preconditioned and solved, but even more importantly, the linear system size is tiny compared to the other steps, as the dimension of the problem is one order lower *and* ζ is defined on Γ_N only.

In our experiments, the time spent solving the linear system corresponding to projecting the pressure Dirichlet data is $< 1\%$, but one might consider one of the following (even cheaper) alternatives: (i) lumping the corresponding mass matrix, (ii) restricting the problems to the support of each node i to project the nodal DoF $[\underline{\zeta}]_i$ to set $[\underline{\mathbf{p}}]_i$ or (iii) even using an averaged value on each node. In case of multiple outlets, the projections are also independent, which could be exploited as well. Moreover, prescribing mean pressures, e.g., computed from models of the downstream vasculature, one might even consider completely neglecting the projection and thereby the viscous contribution on Γ_N . Given the small problem sizes, efficient solution methods available and relatively simple implementation, within this work we stick to one (global) mass matrix system and solve for all coefficients in $\underline{\zeta}$ simultaneously, even if multiple disconnected sections composing Γ_N are present (e.g., multiple outlets in the cardiovascular setting), as this strategy was found uncritical in all of the considered examples.

3.1.3 Momentum equation

The weak form of the linear momentum balance equation is completely identical to the coupled velocity-pressure formulation (see Sec. 2.3.2) and thus not repeated in its

Tab. 3.1: Coefficients α_j^m in backward differentiation formulae with variable time step size and orders $m = 1, 2$ [90].

j	0	1	2
$m = 2$	$\frac{2\Delta t^n + \Delta t^{n-1}}{\Delta t^n(\Delta t^n + \Delta t^{n-1})}$	$-\frac{\Delta t^n + \Delta t^{n-1}}{\Delta t^n \Delta t^{n-1}}$	$\frac{\Delta t^n}{\Delta t^{n-1}(\Delta t^n + \Delta t^{n-1})}$
$m = 1$	$\frac{1}{\Delta t^n}$	$-\frac{1}{\Delta t^n}$	—

entirety. Starting from Eqn. (2.17),

$$\rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2\nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b},$$

with $\chi = 0$ for the generalised Laplacian form and $\chi = 1$ for the the stress-divergence form, multiply with $\mathbf{w} \in [H^1(\Omega)]^d$, $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$ and integrate by parts. Then, the variational problem is to find $\mathbf{u} \in [H^1(\Omega)]^d$ with $\mathbf{u}|_{\Gamma_D} = \mathbf{g}$, such that there holds

$$\begin{aligned} \langle \mathbf{w}, \rho [\partial_t \mathbf{u} + (\nabla \mathbf{u}) \mathbf{u}] - \mathbf{b} \rangle + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle - \langle \nabla \cdot \mathbf{w}, p \rangle - \langle \mathbf{w}, \bar{\mathbf{t}} \rangle_{\Gamma_N} = 0, \end{aligned} \quad (3.8)$$

for all $\mathbf{w} \in [H^1(\Omega)]^d$ with $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$ and where $\bar{\mathbf{t}} = \tilde{\mathbf{t}} + \chi \mu (\nabla^\top \mathbf{u}) \mathbf{n}$ denotes a given Neumann datum being either pseudo tractions $\tilde{\mathbf{t}}$ for $\chi = 0$ and the generalised Laplacian form or real tractions $\mathbf{t} = \boldsymbol{\sigma} \mathbf{n}$ for $\chi = 1$ and the stress-divergence form. Now, all second-order derivatives have been completely eliminated from the weak form, and we can proceed with timestepping, linearisation and decoupling velocity and pressure in the split-step scheme.

3.2 Discretisation in time

Keeping the notation unchanged, we decompose the time interval of interest $I_t = (0, T]$ from $t = 0$ to end time T into N_t steps of possibly variable length $\Delta t^n = t^{n+1} - t^n$ with $n = 0, \dots, N_t$. The time derivative $\partial_t \mathbf{u}$ is approximated by a possibly higher-order backward differentiation formula (BDF)

$$\partial_t \mathbf{u}|_{t=t^{n+1}} \approx \sum_{j=0}^m \alpha_j^m \mathbf{u}^{n+1-j} = \alpha_0^m \mathbf{u}^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m \mathbf{u}^{n-j}, \quad (3.9)$$

with coefficients listed in Tab. 3.1. Then, we can write the time-discrete weak form governing the velocity as: find $\mathbf{u}^{n+1} \in [X]^d$ with $X := H^1(\Omega)$ and $\mathbf{u}^{n+1}|_{\Gamma_D} = \mathbf{g}^{n+1}$ given data from previous time steps \mathbf{u}^{n+1-j} , $j = 1, \dots, m$, such that

$$\begin{aligned} 0 = \langle \mathbf{w}, \rho \sum_{j=0}^m \alpha_j^m \mathbf{u}^{n+1-j} + \rho (\nabla \mathbf{u}^{n+1}) \mathbf{u}^{n+1} - \mathbf{b}^{n+1} \rangle - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}^{n+1}) \nabla \mu^{n+1} \rangle \\ + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu^{n+1} \nabla \mathbf{u}^{n+1} + \chi \mu^{n+1} \nabla^\top \mathbf{u}^{n+1} \rangle - \langle \nabla \cdot \mathbf{w}, p^{n+1} \rangle - \langle \mathbf{w}, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N}, \end{aligned} \quad (3.10)$$

holds for all $\mathbf{w} \in [X]^d$, with $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$. In the weak form of the PPE, the BDF is applied to the boundary term involving $\partial_t \mathbf{u}$, which on Γ_D is identical to $\partial_t \mathbf{g}$. Thus, we aim to find $p \in X$ incorporating continuous Dirichlet data $p^{n+1}|_{\Gamma_N} = \zeta^{n+1}$, such that

$$\begin{aligned} \langle \nabla q, \nabla p^{n+1} \rangle &= \langle \nabla q, \mathbf{b}^{n+1} - (\rho \nabla \mathbf{u}^{n+1}) \mathbf{u}^{n+1} + 2 (\nabla^\top \mathbf{u}^{n+1}) \nabla \mu^{n+1} \rangle \\ &\quad + \langle \mathbf{n} \times \nabla q, \mu^{n+1} \nabla \times \mathbf{u}^{n+1} \rangle_{\Gamma_D} - \langle \mathbf{q} \mathbf{n}, \rho \sum_{j=0}^m \alpha_j^m \mathbf{g}^{n+1-j} \rangle_{\Gamma_D}, \end{aligned} \quad (3.11)$$

for all $q \in X$, $q|_{\Gamma_N} = 0$. The projection of boundary data seeks for $\zeta^{n+1} \in H^{1/2}(\Gamma_N)$, such that for all $s \in L^2(\Gamma_N)$, there holds

$$\begin{aligned} \langle s, \zeta^{n+1} \rangle_{\Gamma_N} &= \langle s \mathbf{n}, \mu^{n+1} (\nabla \mathbf{u}^{n+1} + \chi \nabla^\top \mathbf{u}^{n+1}) \mathbf{n} - \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N} \\ &\quad - \langle s, \mu^{n+1} \nabla \cdot \mathbf{u}^{n+1} \rangle_{\Gamma_N}, \end{aligned} \quad (3.12)$$

and the viscosity projection is completely identical to the one used in the BVS setting, namely, find $\mu \in H^1(\Omega)$, such that

$$\langle r, \mu^{n+1} \rangle = \langle r, \eta(\dot{\gamma}(\nabla^s \mathbf{u}^{n+1})) \rangle \quad \forall r \in L^2(\Omega). \quad (3.13)$$

Regarding the function spaces, let us remark here that the mathematical analysis for a setting in H^1 , and subsequent discretisation via standard continuous Lagrangian finite elements, using either equal-order interpolation or a Taylor–Hood pair for velocity and pressure still remains an open problem. However, strong numerical evidence in both the Newtonian and non-Newtonian case indicates good stability properties of such PPE-based methods [154, 178–183].

Comparing to the BVS scheme, the involved weak forms seem oddly familiar—in fact, we merely skip the added $\beta \nabla \cdot \mathbf{u}$ in the PPE and apply Dirichlet boundary conditions for the pressure. Formulating a monolithic system based on Eqns. (3.10)–(3.13) is rather complex, since we would need to incorporate $p = \zeta$ on Γ_N one way or another, e.g., via Lagrange multipliers or a penalty method. The key idea in split-step schemes, however, is to entirely *avoid* solving such a coupled system. The desired decoupling of the primary variables can be reached by extrapolating in time, which is in fact very much similar to the approach employed in the velocity-pressure formulations to decouple the viscosity from the remaining Navier–Stokes system. Using extrapolations of the form

$$\begin{aligned} \mathbf{u}^{n+1} \approx \mathbf{u}^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{u}^{n-j}, \quad p^{n+1} \approx p^* = \sum_{j=0}^{m-1} \beta_{j+1}^m p^{n-j}, \\ \mu^{n+1} \approx \mu^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mu^{n-j}, \end{aligned} \quad (3.14)$$

with coefficients given in Tab. 2.1, a variety of options in constructing a split-step scheme arise. Here, we settle for a variant first solving a linearised momentum equation with extrapolated convective velocity, viscosity and pressure (gradient), project the viscosity

based on the just computed \mathbf{u}^{n+1} , construct continuous boundary data for the pressure on Γ_N given $(\mathbf{u}^{n+1}, \mu^{n+1})$ and then finally solve for the pressure.

Such a scheme is iteration-free *per se* and numerical evidence suggests that an optional iterative loop enclosing these steps is not strictly necessary, but might robustify the method in some scenarios. The specific variant of the higher-order split-step scheme we consider within this work can thus be summarised in the following steps:

-
1. *Viscosity initialisation*: Compute the viscosity μ^0 by solving Eqn. (3.13) with \mathbf{u}^0 .
 2. *Pressure initialisation*: Compute the pressure p^0 by first projecting the boundary data ζ^0 via (3.12) and then solving the PPE (3.11).
 3. *Initial lower-order steps*: Until enough time step data is gathered, solve the split-step scheme with reduced order $\hat{m} = 1, \dots, m - 1$.
 4. *Timestepping*:

FOR $n = m - 1, \dots, N_t$

- a) Extrapolate in time via Eqn. (3.14), yielding \mathbf{u}^* , p^* and μ^* .
- b) Find $\mathbf{u}^{n+1} \in [X]^d$, $\mathbf{u}^{n+1}|_{\Gamma_D} = \mathbf{g}^{n+1}$ such that

$$\begin{aligned}
0 = & \langle \mathbf{w}, \rho \sum_{j=0}^m \alpha_j^m \mathbf{u}^{n+1-j} + \rho (\nabla \mathbf{u}^{n+1}) \mathbf{u}^* - \mathbf{b}^{n+1} \rangle - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}^{n+1}) \nabla \mu^* \rangle \\
& + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu^* \nabla \mathbf{u}^{n+1} + \chi \mu^* \nabla^\top \mathbf{u}^{n+1} \rangle \\
& - \langle \nabla \cdot \mathbf{w}, p^* \rangle - \langle \mathbf{w}, \tilde{\mathbf{t}}^{n+1} \rangle_{\Gamma_N},
\end{aligned} \tag{3.15}$$

holds for all $\mathbf{w} \in [X]^d$, $\mathbf{w}|_{\Gamma_D} = \mathbf{0}$.

- c) Recover the viscosity μ^{n+1} through (3.13) given \mathbf{u}^{n+1} .
- d) Project ζ^{n+1} via (3.12) based on \mathbf{u}^{n+1} and μ^{n+1} .
- e) Solve the PPE (3.11) for p^{n+1} using \mathbf{u}^{n+1} , μ^{n+1} and ζ^{n+1} .
- f) Update the time step size Δt^n and corresponding coefficients α_j^m and β_j^m .

END FOR

In the above algorithm, Eqn. (3.15) governing the velocity couples all velocity components, regardless of $\chi = 0$ for the generalised Laplacian formulation or $\chi = 1$ for the stress-divergence form. One might want to decouple the individual velocity components, which is easily achieved in the generalised Laplacian form, by additionally considering the component-coupling part of the viscous stress term, i.e., $(\nabla^\top \mathbf{u}) \nabla \mu$, fully explicit. Doing so allows seeking for individual velocity components $u_i^{n+1} \in X$, with $u_i^{n+1}|_{\Gamma_D} = g_i^{n+1}$, such that

$$\begin{aligned}
& \langle w_i, \rho \alpha_0^m u_i^{n+1} + \rho \mathbf{u}^* \cdot \nabla u_i^{n+1} \rangle + \langle \nabla w_i, \mu^* \nabla u_i^{n+1} \rangle \\
= & \langle w_i, b_i^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m u_i^{n-j} + \nabla \mu^* \cdot \partial_{x_i} \mathbf{u}^* \rangle + \langle \partial_{x_i} w_i, p^* \rangle + \langle w_i, \tilde{t}_i^{n+1} \rangle_{\Gamma_N},
\end{aligned} \tag{3.16}$$

holds for all $w_i \in X$, $w_i|_{\Gamma_D} = 0$, which is just a scalar advection-diffusion-reaction equation. Independent of the order of time integration, we end up with a standard problem, available *off-the-shelf* in many open-source finite element software packages. Being such a standard problem, AMG methods are often easier tuned in this setup, possibly rendering the solution process more efficient. This matter is further investigated in the corresponding numerical results Sec. 3.5, where we examine the effect of the linearisation on the critical time step size. This is motivated by the fact that semi-implicit treatment of the viscous term might induce an additional, potentially limiting time step constraint.

Another popular simplification in this context is to consider the convective term fully explicit, that is, $\mathbf{u}^* \cdot \nabla u_i^*$ instead of $\mathbf{u}^* \cdot \nabla u_i^{n+1}$ in Eqn. (3.16). This introduces a standard hyperbolic time step restriction, but leads to a symmetric and positive definite system matrix, such that tailored, hence faster, numerical tools can be employed. Again, we will investigate the effect of this option on the overall scheme in Sec. 3.5, since a potential speed-up awaits, if the triggered problem dependent instabilities are not spoiling the scheme's stability too much.

Regarding a comparison of the stress-divergence and generalised Laplacian forms, apart from allowing a decoupling of the velocity components, it is at this point not clear how the schemes compare, even if a coupled vector-valued momentum equation is considered. In the proof of Thm. 2.1.1, which states equivalence of the standard Navier–Stokes system and our PPE-based reformulation, we ended up with a heat equation in the auxiliary variable $\Phi := \nabla \cdot \mathbf{u}$ reading

$$\rho \partial_t \Phi - (1 + \chi) \nabla \cdot (\mu \nabla \Phi) = 0,$$

with zero initial, Dirichlet and Neumann boundary conditions. This might indicate improved mass conservation properties at the discrete level employing the stress-divergence form ($\chi = 1$) due to the increased diffusion coefficient, although the numerical experiments in Sec. 3.5.2 do *not* confirm that. For a more in-depth discussion, theoretical and numerical comparison between these two forms, we refer the interested reader to the works by Limache et al. [186] and Pacheco et al. [26].

As already pointed out, the generalised Laplacian form is particularly well suited for pure flow simulations, as it features natural boundary conditions involving pseudo tractions for truncated domains. However, one might also correct for the difference in pseudo and real tractions adding a boundary term to the left-hand side in scenarios like fluid–structure interaction (with Robin interface conditions) or multi-phase flow, which motivates keeping both forms in what follows.

3.3 Improving mass conservation

When the divergence-free constraint is replaced by the consistent PPE, an inf-sup stabilisation is employed, or a continuous pressure approximation with Taylor–Hood elements is chosen, the resulting velocity field is not *exactly*, but merely *weakly* divergence-free.

Although the PPE enforces incompressibility as well, Liu et al. [178] observed that including a Leray projection considerably increases stability for highly non-smooth solutions. Similar to classical projection methods (see, e.g., [30, 166, 167, 169] and references therein), the idea is to simply split the velocity as computed from the momentum balance equation (3.15) via a Helmholtz–Leray decomposition into

$$\mathbf{u} = \hat{\mathbf{u}} + \nabla\psi, \quad (3.17)$$

where $\hat{\mathbf{u}}$ is solenoidal. Taking the divergence of (3.17), we have

$$\begin{aligned} \nabla \cdot \mathbf{u} &= \nabla \cdot \hat{\mathbf{u}} + \nabla \cdot \nabla\psi, \\ -\nabla \cdot \mathbf{u} &= -\Delta\psi, \end{aligned}$$

and construct a complete boundary value problem adding suitable conditions,

$$-\Delta\psi = -\nabla \cdot \mathbf{u} \quad \text{in } \Omega, \quad (3.18)$$

$$\mathbf{n} \cdot \nabla\psi = 0 \quad \text{on } \Gamma_D, \quad (3.19)$$

$$\psi = 0 \quad \text{on } \Gamma_N. \quad (3.20)$$

Unfortunately now, the divergence-free velocity field $\hat{\mathbf{u}} = \mathbf{u} - \nabla\psi$ does not fulfil the boundary conditions of the original problem—the momentum balance equation—which can easily be shown by verifying that

$$\begin{aligned} \nabla \cdot \hat{\mathbf{u}} &= \nabla \cdot \mathbf{u} - \nabla \cdot \nabla\psi = \nabla \cdot \mathbf{u} - \nabla\psi = 0 && \text{in } \Omega, \\ \mathbf{n} \cdot \hat{\mathbf{u}} &= \mathbf{n} \cdot \mathbf{u} - \mathbf{n} \cdot \nabla\psi = \mathbf{n} \cdot \mathbf{u} && \text{on } \Gamma_D, \\ \boldsymbol{\tau} \cdot \hat{\mathbf{u}} &= \boldsymbol{\tau} \cdot \mathbf{u} - \boldsymbol{\tau} \cdot \nabla\psi = \boldsymbol{\tau} \cdot \mathbf{u} && \text{on } \Gamma_N, \end{aligned}$$

for any tangential vector $\boldsymbol{\tau}$ on Γ_N . That is, we have to choose between the originally computed velocity field \mathbf{u} , on which the PPE inherently imposed $\nabla \cdot \mathbf{u} = 0$, or the newly constructed $\hat{\mathbf{u}}$, on which $\nabla \cdot \hat{\mathbf{u}} = 0$ is explicitly enforced, but does only partially fulfil the desired boundary conditions. As discussed by Guermond et al. [30], both fields converge at the same rates, such that from an accuracy standpoint, both choices are equally well suited. Herein, the motivation for using a projection is to suppress a possible build-up of errors in mass conservation as a consequence of introducing a split-step scheme rather than explicitly enforcing the incompressibility constraint, where we want to refer the reader to the original work by Liu et al. [178] for a detailed discussion.

Another downside of the classical Leray projection is that the divergence-free velocity field $\hat{\mathbf{u}} = \mathbf{u} - \nabla\psi$ is discontinuous across element edges if standard C^0 -continuous Lagrangian finite elements are used to solve the auxiliary Poisson equation in ψ . As a consequence, yet another projection onto a continuous space has to be employed, increasing the numerical effort. On top of that, the projected field $\hat{\mathbf{u}}_h \in [X_h]^d \subset [H^1(\Omega)]^d$ is then also only *weakly* divergence-free.

In order to circumvent these shortcomings, we follow [21] and employ what is referred to as *divergence damping* [179, 183]. Divergence damping is a compromise between the velocity \mathbf{u} as computed from the split-step scheme and the Leray-projected divergence-

free velocity $\hat{\mathbf{u}}$ not fulfilling the desired boundary conditions. We can suppress error accumulation by applying the Leray projection onto the old time steps' velocities only, such that the additional vector-valued projection of the velocity is not needed and hence even reducing the associated numerical effort. The time derivative term discretised via a BDF including divergence damping is then approximated as

$$\partial_t \mathbf{u}|_{t=t^{n+1}} \approx \alpha_0^m \mathbf{u}^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m \hat{\mathbf{u}}^{n-j} = \alpha_0^m \mathbf{u}^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m (\mathbf{u}^{n-j} - \nabla \psi^{n-j}),$$

which we can directly plug into the balance of linear momentum (3.16), yielding

$$\begin{aligned} & \langle \rho \mathbf{w}, \alpha_0^m \mathbf{u}^{n+1} + (\nabla \mathbf{u}^{n+1}) \mathbf{u}^* \rangle \\ & + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu^* \nabla \mathbf{u}^{n+1} + \chi \mu^* \nabla^\top \mathbf{u}^{n+1} \rangle - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}^{n+1}) \nabla \mu^* \rangle \\ & = \langle \mathbf{w}, \mathbf{b}^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m (\mathbf{u}^{n-j} - \nabla \psi^{n-j}) \rangle + \langle \nabla \cdot \mathbf{w}, p^* \rangle + \langle \mathbf{w}, \tilde{\mathbf{t}}^{n+1} \rangle_{\Gamma_N}, \end{aligned} \quad (3.21)$$

or, for the generalised Laplacian form with semi-implicit viscous terms,

$$\begin{aligned} & \langle w_i, \rho \alpha_0^m u_i^{n+1} + \rho \mathbf{u}^* \cdot \nabla u_i^{n+1} \rangle + \langle \nabla w_i, \mu^* \nabla u_i^{n+1} \rangle \\ & = \langle w_i, b_i^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m (u_i^{n-j} - \partial_{x_i} \psi^{n-j}) \rangle + \nabla \mu^* \cdot \partial_{x_i} \mathbf{u}^* + \langle \partial_{x_i} w_i, p^* \rangle + \langle w_i, \tilde{t}_i^{n+1} \rangle_{\Gamma_N}, \end{aligned} \quad (3.22)$$

sparing the projection of $\hat{\mathbf{u}}$ onto a continuous space. This technique can be interpreted as penalising violations of mass conservation. Thereby, this simple approach considerably improves mass balance and additionally increases temporal stability of the split-step scheme. Per time step, we thus only add computational overhead in terms of the Poisson problem seeking for $\psi \in X$ with $\psi|_{\Gamma_D} = 0$, such that

$$\langle \nabla q, \nabla \psi \rangle = - \langle q, \nabla \cdot \mathbf{u} \rangle \quad \forall q \in X, q|_{\Gamma_N} = 0. \quad (3.23)$$

Note here, however, that the boundary conditions for ψ in Eqns. (3.19) and (3.20)—which might have seemed intuitive, but somewhat arbitrary at first—now upon closer inspection turn out to be well suited for a combination of the two Poisson problems being the PPE and Eqn. (3.23). In the decoupled and linearised momentum balance equation we have

$$\nabla p^* + \sum_{j=0}^{m-1} \alpha_{j+1}^m \nabla \psi^{n-j},$$

which we summarise in a new, pressure-like variable \hat{p}^* . It is particularly important here that neither $\psi|_{\Gamma_N} = 0$ spoils the pressure Dirichlet condition, nor does $\mathbf{n} \cdot \nabla \psi|_{\Gamma_D} = 0$ introduce inconsistencies in the Neumann condition on the pressure.

Therefore, one can (i) drop the terms involving ψ in the momentum balance equation, and (ii) combine the Poisson problems, seeking for $\hat{p}^* \in X$ with $\hat{p}^*|_{\Gamma_N} = \zeta^*$,

$$\langle \nabla q, \nabla \hat{p}^* \rangle = -\langle q, \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \nabla \cdot \mathbf{u}^{n-j} \rangle + \sum_{j=0}^{m-1} \beta_{j+1}^m l(q, \mathbf{u}^{n-j}, \mu^{n-j}, t^{n-j}) \quad (3.24)$$

for all $q \in X$, where $q|_{\Gamma_N} = 0$. Here, l evaluated at $t = t^i$ is defined as

$$\begin{aligned} l(q, \mathbf{u}, \mu, t^i) := & \langle \nabla q, \mathbf{b}|_{t=t^i} - \rho(\nabla \mathbf{u})\mathbf{u} + 2(\nabla^\top \mathbf{u}) \nabla \mu \rangle \\ & + \langle \mathbf{n} \times \nabla q, \mu \nabla \times \mathbf{u} \rangle_{\Gamma_D} - \langle q \mathbf{n}, \rho \partial_t \mathbf{g}|_{t=t^i} \rangle_{\Gamma_D}, \end{aligned} \quad (3.25)$$

where we might again use a BDF to approximate $\partial_t \mathbf{g}|_{t=t^i} \approx \sum_{j=0}^m \alpha_j^m \mathbf{g}^{n+1-j}$ or replace the extrapolation of $\partial_t \mathbf{g}$ with the BDF approximation at t^{n+1} in (3.24). The boundary projection variable is also consistently extrapolated in time, such that the problem is to find $\zeta^* \in H^{1/2}(\Gamma_N)$, such that for all $s \in L^2(\Gamma_N)$, there holds

$$\begin{aligned} \langle s, \zeta^* \rangle_{\Gamma_N} = & - \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s, \mu^{n+1} \nabla \cdot \mathbf{u}^{n-j} \rangle_{\Gamma_N} \\ & + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s \mathbf{n}, \mu^{n-j} (\nabla \mathbf{u}^{n-j} + \chi \nabla^\top \mathbf{u}^{n-j}) \mathbf{n} - \bar{\mathbf{t}}^{n-j} \rangle_{\Gamma_N}. \end{aligned} \quad (3.26)$$

Now, inspecting the indices and thus the required time step data, one sees that only old time step data is being used to compute the combined pressure variable \hat{p}^* . Thus, we can switch the sequence of steps in the split-step scheme formally—recovering the pressure extrapolation and Leray projection variables still implicitly (at old time steps)—to obtain a scheme with divergence damping and combined pressure-like variable as:

-
1. *Viscosity initialisation:* Compute the viscosity μ^0 by solving Eqn. (3.13) with \mathbf{u}^0 .
 2. *Initial lower-order steps:* Until enough time step data is gathered, solve the split-step scheme with order $\hat{m} = 1, \dots, m-1$.
 3. *Timestepping:*
FOR $n = m-1, \dots, N_t$
 - a) Compute continuous pressure boundary data ζ^* via (3.26).
 - b) Solve the modified Poisson equation (3.24) for \hat{p}^* .
 - c) Extrapolate in time via Eqn. (3.14), yielding \mathbf{u}^* and μ^* .
 - d) Solve the momentum balance equation for \mathbf{u}^{n+1} directly via (3.21), or for individual components through (3.22), neglecting any terms involving ψ .
 - e) Recover the viscosity μ^{n+1} via (3.13) given \mathbf{u}^{n+1} .
 - f) Update the time step size Δt^n and corresponding coefficients α_j^m and β_j^m .

END FOR

The number of steps within the split-step scheme does thus not increase from the base method without divergence damping, consisting of (i) a projection of the pressure Dirichlet data ζ , (ii) a PPE for the combined Leray and pressure variable \hat{p} , (iii) a single vector-valued or, alternatively, d decoupled advection-diffusion-reaction equation(s) in the velocity (components) and lastly, (iv) a projection of the viscosity onto a continuous space.

Despite the fact that divergence damping is employed to additionally penalise build-up of errors in mass conservation, we want to emphasize that the presented scheme—similar to its Newtonian counterpart with constant viscosity proposed by Liu [21]—is distinctly different from projection schemes, as it is based on a consistent PPE, inherently enforcing a solenoidal velocity field. Owing to the fact that we do not project the velocity onto a weakly divergence-free space and use the resulting scalar quantity to update the pressure, but rather compute it from an appropriate Poisson problem, the mathematical properties change. As a consequence, the proposed scheme is not subject to an inf-sup condition, meaning, equal-order interpolation is admissible [154], whereas classical projection schemes require an LBB-compatible choice of function spaces for velocity and pressure [185]. The most challenging aspect in the presented split-step scheme is thus handling the more involved right-hand side of the PPE, as terms involving third-order derivatives of the velocity naturally arise. Luckily, we are able to reformulate those terms, such that the final scheme can be discretised with standard C^0 -continuous Lagrangian finite elements, as shall be seen next.

3.4 Solving the linear systems

Discretisation in space is carried out in the usual manner, constructing continuous nodal basis functions φ^u , φ^p , φ^μ and φ^ζ , spanning the respective function spaces. For any test and ansatz function (or its components), continuous functions from $X_h \subset X := H^1(\Omega)$ in the domain and the respective trace on Γ_N for the pressure Dirichlet data projection are employed. Then, the fully discrete weak forms in the higher-order split-step scheme with divergence damping directly lead to systems of linear equations in the nodal degrees of freedom (DoFs). As previously presented, the first step consists of projecting boundary data on Γ_N to recover the nodal coefficients of $\hat{\zeta}_h^*$, gathered in the discrete solution vector $\underline{\hat{\zeta}}^*$, by solving

$$\mathbf{M}_\zeta \underline{\hat{\zeta}}^* = \underline{\mathbf{e}} \quad (3.27)$$

followed by the the pressure Poisson problem with $\hat{p}_h^* = \hat{\zeta}_h^*$ on Γ_N ,

$$\mathbf{L} \underline{\hat{p}}^* = \underline{\mathbf{f}}, \quad (3.28)$$

to compute the nodal vector $\underline{\hat{p}}^*$ corresponding to \hat{p}_h^* . The pressure-like \hat{p}_h^* is then used together with extrapolated quantities \mathbf{u}_h^* and μ_h^* to assemble the linear system arising from linear momentum balance with $\mathbf{u}_h^{n+1} = \mathbf{g}_h^{n+1}$ on Γ_D ,

$$\mathbf{A} \mathbf{u}^{n+1} = \mathbf{g}, \quad (3.29)$$

which we solve for \mathbf{u}^{n+1} constituting \mathbf{u}_h^{n+1} and finally recover the viscosity, solving a scaled mass matrix problem

$$\mathbf{M}_\mu \tilde{\underline{\mu}} = \underline{\mathbf{h}} \quad \text{and} \quad \underline{\mu}^{k+1} = \tilde{\eta} \tilde{\underline{\mu}}, \quad (3.30)$$

where we take $\tilde{\eta} = \eta_\infty$, very much similar to the viscosity projection step in the coupled solver used to tackle the velocity-pressure formulation. The above used matrices and vectors are defined as

$$[\mathbf{M}_\zeta]_{ij} := \langle \varphi_i^\zeta, \varphi_j^\zeta \rangle_{\Gamma_N}, \quad [\mathbf{L}]_{ij} := \langle \nabla \varphi_i^p, \nabla \varphi_j^p \rangle, \quad (3.31)$$

$$\begin{aligned} [\mathbf{e}]_i &:= - \sum_{j=0}^{m-1} \beta_{j+1}^m \langle \varphi_i^\zeta, \mu_h^{n+1} \nabla \cdot \mathbf{u}_h^{n-j} \rangle_{\Gamma_N} \\ &\quad + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle \varphi_i^\zeta \mathbf{n}, \mu_h^{n-j} (\nabla \mathbf{u}_h^{n-j} + \chi \nabla^\top \mathbf{u}_h^{n-j}) \mathbf{n} - \bar{\mathbf{t}}^{n-j} \rangle_{\Gamma_N} \end{aligned} \quad (3.32)$$

$$\begin{aligned} [\mathbf{f}]_i &:= - \sum_{j=0}^{m-1} \alpha_{j+1}^m \langle \varphi_i^p, \rho \nabla \cdot \mathbf{u}_h^{n-j} \rangle \\ &\quad + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle \nabla \varphi_i^p, \mathbf{b}^{n-j} - (\rho \nabla \mathbf{u}_h^{n-j}) \mathbf{u}_h^{n-j} + 2 (\nabla^\top \mathbf{u}_h^{n-j}) \nabla \mu_h^{n-j} \rangle \\ &\quad + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle \mathbf{n} \times \nabla \varphi_i^p, \mu \nabla \times \mathbf{u}_h^{n-j} \rangle_{\Gamma_D} - \sum_{j=0}^m \alpha_j^m \langle q \mathbf{n}, \rho \mathbf{g}^{n+1-j} \rangle_{\Gamma_D} \end{aligned} \quad (3.33)$$

$$\begin{aligned} [\mathbf{A}]_{ij} &:= \langle \rho \varphi_i^u, \alpha_0^m \varphi_j^u + (\nabla \varphi_j^u) \mathbf{u}_h^* \rangle - (1 - \chi) \langle \varphi_i^u, (\nabla^\top \varphi_j^u) \nabla \mu_h^* \rangle \\ &\quad + (1 - \chi/2) \langle \nabla \varphi_i^u + \chi \nabla^\top \varphi_i^u, \mu^* \nabla \varphi_j^u + \chi \mu_h^* \nabla^\top \varphi_j^u \rangle \end{aligned} \quad (3.34)$$

$$[\underline{\mathbf{g}}]_i := \langle \varphi_i^u, \mathbf{b}^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \mathbf{u}_h^{n-j} \rangle + \langle \nabla \cdot \varphi_i^u, p_h^* \rangle + \langle \varphi_i^u, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N} \quad (3.35)$$

$$[\mathbf{M}_\mu]_{ij} := \langle \varphi_i^\mu, \varphi_j^\mu \rangle, \quad [\underline{\mathbf{h}}]_i := 1/\tilde{\eta} \langle \varphi_i^\mu, \eta (\dot{\gamma} (\nabla^s \mathbf{u}_h^{n+1})) \rangle \quad (3.36)$$

Matrices \mathbf{M}_ζ , \mathbf{L} , \mathbf{M}_μ and \mathbf{A} are standard matrices in the finite element context and thus easily preconditioned and solved. Regarding the linear solver, a CG method [100] is employed for the symmetric and positive definite mass matrix and Poisson problems, whereas if \mathbf{A} is non-symmetric, as is the case when convective and/or the non-symmetric diffusive terms are considered implicit, an FGMRES solver [135] is used. Efficient preconditioners for these standard problems are constructed through open-source software packages such as BoomerAMG [95] or Trilinos' ML package [111], on the latter of which we heavily rely on. More often than not, default settings for the individual problem

types are predefined and can be further tuned to achieve satisfying results as shown in a series of numerical examples in the following section.

3.5 Computational results

Let us now investigate on the accuracy, stability and efficiency in numerical experiments. We consider first a manufactured solution to establish estimated orders of convergence and then move on to the two-dimensional lid-driven cavity problem, where several variants of the scheme are compared in terms of temporal stability. Then, a three-dimensional abdominal aortic aneurysm (AAA) is considered, where we evaluate performance of the split-step scheme and compare it to the coupled solution approach using BVS stabilisation as presented in Sec. 2.2. Both the coupled solver and the split-step scheme are implemented with help of the finite element toolbox `deal.II` [97, 98], using Lagrangian tensor-product finite elements denoted as Q_1Q_1 or Q_2Q_1 for the linear, equal-order case and a Taylor–Hood pair, respectively. The space for viscosity projection is in any case the same as the one used for pressure approximation. Unless stated otherwise, full projections (i.e., no lumped mass matrix solves) and the generalised Laplacian form with semi-implicit diffusion and implicit, linearised convective term are considered, decoupling individual velocity components, but keeping the convective term within \mathbf{A} .

As already mentioned in the context of linearising the convective term in the momentum equation in a coupled velocity-pressure formulation, treating the convective terms semi-implicitly introduces a standard CFL condition. Numerical evidence suggests, that the split-step scheme, extrapolating several solution components, does *not further* limit the maximum admissible time step size [3, 21]. Thus, when adaptive timestepping is considered, we settle for a puristic time step control written as

$$\Delta t^{n+1} = \max \left[\Delta t_{\max}, \Delta t^n 0.98 \min \left(1.02, \frac{\text{CFL}_{\max}}{\max_{e=1, \dots, N_e} \text{CFL}_e} \right) \right], \quad (3.37)$$

merely targeting a predefined maximum CFL number, CFL_{\max} , based on the element CFL number denoted by CFL_e , which is defined on each element as

$$\text{CFL}_e = \max_{i=1, \dots, d} \frac{|u_i| \Delta t^n}{h_i}, \quad (3.38)$$

with element length h_i being the maximum element node distance in direction x_i divided by the polynomial degree.

3.5.1 Manufactured solution in 2D

Consider the two-dimensional unit domain $\Omega = (0, H) \times (0, H)$ with a side length of $H = 1$, wherein the manufactured solution and rheological law are given as

$$p(\mathbf{x}, t) = 2 \sin(2 - 2x)f(t), \quad \mathbf{u}(\mathbf{x}, t) = \begin{pmatrix} f(t) \sin 2y \sin^2 x \\ -f(t) \sin 2x \sin^2 y \end{pmatrix},$$

$$\eta(\dot{\gamma}) = \left[1 + (10\dot{\gamma})^2\right]^{-1/4},$$

modelling a shear-thinning fluid. We consider the pure Dirichlet problem, i.e., $\Gamma_D = \Gamma$ and, consequently, $\Gamma_N = \emptyset$. Depending on the chosen function $f(t)$ to ramp up the solution, we can achieve diminishing or dominant spatial or temporal errors as desired. In this whole section, unlike in the other numerical examples, the divergence damping is not used to investigate convergence rates of the basic scheme more clearly. First, let us test spatial convergence rates, setting

$$f(t) = 1 - e^{-2t},$$

such that the solution reaches a steady state as $t \rightarrow \infty$. Starting from a coarse mesh with four uniform quadrilateral elements, uniform refinements are performed and the relative errors defined as

$$\|p - p_h\|_T := \left(\frac{\|p - p_h\|_{L^2(\Omega)}}{\|p\|_{L^2(\Omega)}} \right) \Big|_{t=T}, \quad \|\nabla \mathbf{u} - \nabla \mathbf{u}_h\|_T := \left(\frac{\|\nabla \mathbf{u} - \nabla \mathbf{u}_h\|_{L^2(\Omega)}}{\|\nabla \mathbf{u}\|_{L^2(\Omega)}} \right) \Big|_{t=T},$$

are computed at $t = T = 10$, at which point the steady state is approximately reached due to the function $f(t)$ leading to $1 - f(t) \approx 2 \times 10^{-9}$.

For discretisation in time, the BDF1 scheme with a uniform time step size of $\Delta t = 0.01$ is sufficient, since spatial errors at the end of the considered interval dominate. Fig. 3.1 summarises the results of the spatial convergence study, where the expected rates are observed. The Q_1Q_1 pair achieves convergence rates of 1 in both the pressure's relative L^2 -norm and the relative H^1 -semi-norm of the velocity, while the LBB-stable Taylor-Hood pair features convergence rates of 2 in those norms.




Fig. 3.1: Manufactured solution: spatial errors obtained with Q_1Q_1 and Q_2Q_1 elements when reaching the stationary state at $t = T = 10$.

As already mentioned several times during the derivation of the spit-step solver, the option of lumping the mass matrix in the viscosity projection step seems attractive, since

it allows swapping a mass matrix solve with a simple vector scaling. Thus, the spatial rates when taking a lumped or the full mass matrix \mathbf{M}_μ (3.36) are compared in Fig. 3.2. Results for Q_1Q_1 elements with and without lumped mass matrices in the viscosity projection step show similar results, whereas the Taylor–Hood pair having a higher rate is more affected: while the pressure rate does not suffer from any noticeable drop-off, the velocity rate is slightly decreased starting from the finer grid levels considered within this test. This indicates that depending on the problem at hand, lumping the mass matrix in the viscosity projection step might be a worthwhile compromise between accuracy and speed, as inversion of a lumped mass matrix is trivial.

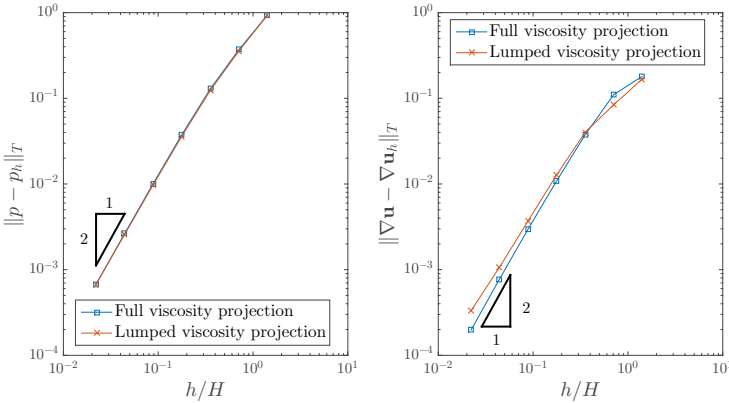


Fig. 3.2: Manufactured solution: spatial errors obtained with Q_2Q_1 elements when reaching the stationary state comparing full or lumped mass matrix solves in the viscosity projection step.

Finally, the temporal convergence rates are investigated with $f(t) = \sin^2 t$ in the time interval from $t = 0$ to $t = 10$, computing the $L^\infty(Q)$ -norms being the maximum error in the space-time domain, $Q := \Omega \times (0, T]$, obtained either via a BDF1 scheme with Q_1Q_1 elements or BDF2 integration combined with the Taylor–Hood pair Q_2Q_1 . The finest mesh constructed for the spatial convergence studies is selected to reduce spatial errors and an initial time step of $\Delta t = 0.16$ is halved in each refinement step until $\Delta t = 1.25 \times 10^{-3}$ is reached. The resulting errors are depicted in Fig 3.3, where the expected orders of convergence in the considered norms can clearly be identified. Note here, that the maximum element CFL number as defined in (3.38) ranges from 0.1 to 15 in this setup, indicating already good stability properties of the split-step scheme.

So, summing up these results concerning experimental convergence rates in space and time, the expected rates were observed in the numerical tests. Additionally, lumping the mass matrix seems to have only a mild effect on the velocity errors, and can thus be considered a worthwhile option for practical applications. During the temporal convergence study, rather high CFL numbers achieved suggest good stability properties, which are further investigated in the following numerical test.




Fig. 3.3: Manufactured solution: investigating temporal convergence rates in maximum space-time errors using a fine grid and BDF1 and BDF2 schemes with Q_1Q_1 and Q_2Q_1 elements, respectively.

3.5.2 Lid-driven cavity flow

To compare the stability properties of the different linearisation options, we resort to the classical lid-driven cavity benchmark in its regularised form as proposed by de Frutos et al. [155]. In two space dimensions, the unit domain $\Omega = (0, H) \times (0, H)$ with $H = 1$ m is subject to a ramped horizontal velocity at the lid, i.e., at $y = H$, which is given by $u_1 = \hat{u}_1 \eta_r(t) \eta_x(x)$, where

$$\eta_r(t) = \begin{cases} \sin^2 \frac{\pi t}{2T_r} & \text{if } t < T_r, \\ 1 & \text{otherwise,} \end{cases} \quad \text{and} \quad \eta_x(x) = \begin{cases} 1 - \cos^4 \frac{\pi x}{2L} & \text{if } x < L, \\ 1 - \cos^4 \frac{\pi(x-H)}{2L} & \text{if } x > H - L, \\ 1 & \text{otherwise,} \end{cases} \quad (3.39)$$

with $T_r = 1$ s and the regularisation width $L = H/10$ m. A pure Dirichlet problem is considered, prescribing $\mathbf{u} = \mathbf{0}$, that is, no-slip conditions at the remaining parts of the boundary. Regarding the rheological law, we consider a Carreau fluid setting $\kappa = 1$ and $a = 2$ in the general form given in Eqn. (2.8), obtaining

$$\eta(\dot{\gamma}) = \eta_\infty + (\eta_0 - \eta_\infty) \left[1 + (\lambda \dot{\gamma})^2 \right]^{\frac{b-1}{2}}. \quad (3.40)$$

Lower and upper viscosity limits are chosen as $\eta_\infty = 1.0$ mPa s, $\eta_0 = 100$ mPa s and the remaining parameters are set to $\lambda = 10$ s and $b = 0.25$. Together with a fluid density of $\rho = 1.0$ kg/m³ and a maximum velocity $\hat{u}_1 = 10$ m/s, this leads to a Reynolds number of $\text{Re} = \rho \hat{u}_1 H / \eta_\infty \approx 10^4$. The computational domain is discretised via 128×128 uniform bilinear Q_1Q_1 elements or 64×64 Q_2Q_1 elements leading to an equal number of velocity DoFs. Using these uniform meshes, latter of which is depicted in Fig. 3.4, boundary layers are deliberately not resolved to test the sensitivity of the split-step scheme with respect to large gradients on comparably coarse meshes.




Fig. 3.4: Lid-driven cavity flow: mesh of 64×64 Q_2Q_1 elements in 4 subdomains.

This setting results in large gradients in the velocity, challenging our solver by the viscosity varying by two orders of magnitude. Similar to the element CFL number, we also define the element Reynolds number as

$$\text{Re}_e = \max_{i=1,\dots,d} \frac{\rho |u_i| h_i}{\mu}, \quad (3.41)$$

for later reference. A stationary state as shown in Fig. 3.5 using Q_2Q_1 interpolation is reached after an initial ramp up phase, where the time step size is selected according to (3.37), with a target maximum element CFL number of $\text{CFL}_{\max} = 5$ and $\Delta t_{\max} = \infty$.

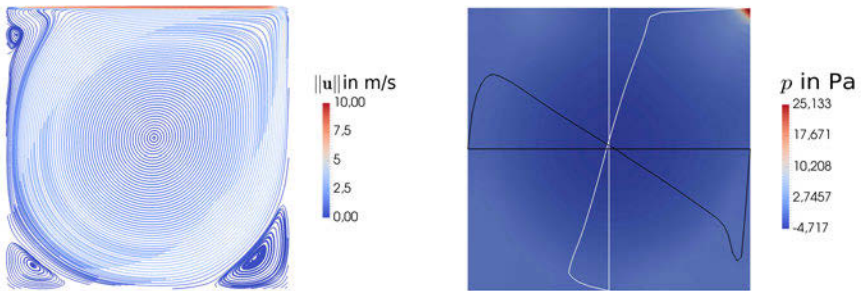


Fig. 3.5: Lid-driven cavity flow: selected streamlines and pressure p overlaid by scaled velocity components $0.05u_1$ at $x = H/2$ (white) and $0.1u_2$ at $y = H/2$ (black).

Now, the different approaches regarding linearisation of the momentum equation, given in Eqns. (3.21) and (3.22), are compared using a uniform time step size which results in a maximum CFL_e of ≈ 10 when reaching the stationary state at $t \approx 1$ s. A CG solver is employed whenever applicable, while in any other case an FGMRES method is used to

reduce the residual by a factor of 10^6 . We report iteration counts needed for the viscosity projection step (N_μ), the pressure Poisson step (N_p) and the vector-valued momentum balance or the alternative mean iteration count when solving the individual components (N_u or \bar{N}_{u_i} , respectively). Fig. 3.6 shows that all considered variants are stable when using Q_1Q_1 elements and BDF1 time integration under a standard CFL condition. Given the high target CFL number of ≈ 10 , some schemes diverge, given different constants in the respective CFL conditions. Iteration counts increase with the element CFL and Reynolds numbers reaching their maxima of ≈ 10 and ≈ 50 , respectively, as the boundary condition is fully ramped up.

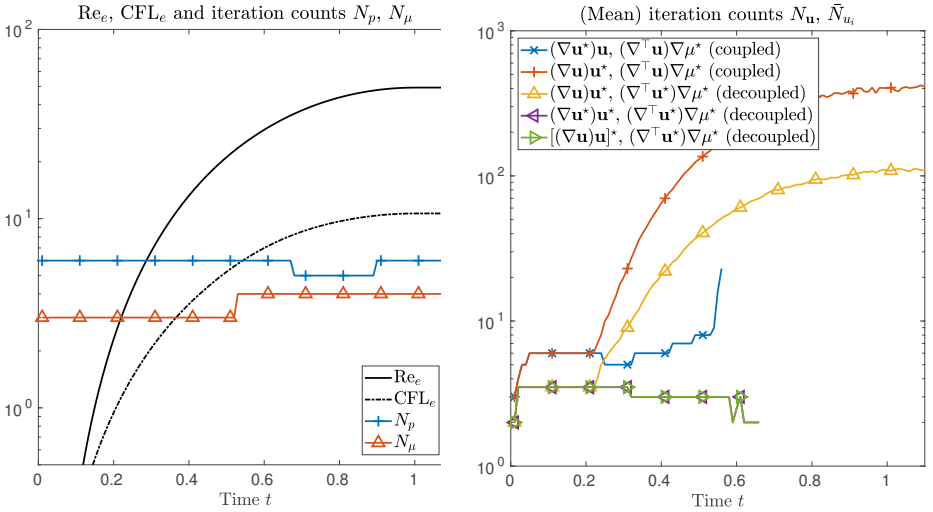


Fig. 3.6: Lid-driven cavity flow: element Reynolds and CFL numbers, iteration counts of the PPE (N_p) and viscosity projection (N_μ) steps (left) and comparison of convective and viscous term variants (right) using BDF1 timestepping and Q_1Q_1 elements.

Inspecting the initial range closer, we see that as long as $CFL_e < 1$ holds, the AMG preconditioner performs well. For the PPE and viscosity projection steps, iteration counts stay constant regardless of interpolation order or time integration scheme used, but the momentum balance step shows vastly different behaviour depending on the chosen linearisation variant. The standard linearisation of the convective term, $(\nabla \mathbf{u}^{n+1}) \mathbf{u}^*$, is considerably more stable than the other options considered here. The choices being fully explicit in the convective term, $(\nabla \mathbf{u}^*) \mathbf{u}^*$ or $[(\nabla \mathbf{u}) \mathbf{u}]^*$, lead to temporal instabilities as soon as $CFL_e = 1$ is surpassed, finally leading to divergence of the linear solver at $t \approx 0.6$ s. The fully explicit variants extrapolating the convective term as a whole or the velocity vector lead to almost identical results and might be worth considering, if the time step size is small enough. The main reason for doing so is that it allows using a CG solver, if the viscous term is treated semi-implicitly as well. Using the most stable variant of the convective term, treating $\nabla^\top \mathbf{u} \nabla \mu$ fully explicit, i.e. $\nabla^\top \mathbf{u}^* \nabla \mu^*$, does not introduce further restrictions on the time step size in our experiments.

Now, introducing again a lumped mass matrix solve in the viscosity projection step, we observe no influence on the iteration counts of the linear solvers and no instabilities are introduced. As can be seen from Fig. 3.7 (left), the iteration counts are almost identical, such that introducing lumping in the viscosity projection step and effectively replacing a scalar-valued mass matrix solve with a simple vector scaling is indeed a worthwhile alternative. However, the effects of lumping the viscosity projection step's mass matrix are largely problem dependent and thus might have consequences if used in other scenarios.




Fig. 3.7: Lid-driven cavity flow: impact of various linearisation variants of the convective term and lumped viscosity projection step (“lumped”) in the generalised Laplacian form (left) and differences to the stress-divergence form (right).

Comparing the generalised Laplacian and stress-divergence forms of the momentum equation, we see that a decoupling of the velocity components leads to distinct linearisations needed. For the generalised Laplacian form, considering $(\nabla^\top \mathbf{u}) \nabla \mu$ fully explicit, shifting it to the right-hand side is enough, while in the stress-divergence form, we end up with $2\mu^* \nabla^s \mathbf{w} : \nabla^s \mathbf{u}$ leading to a symmetric matrix, but also to a more involved coupling of the individual components. In order to achieve a decoupled variant, we first rewrite the term as $\mu^* \nabla \mathbf{w} : \nabla^s \mathbf{u}$ (which is completely equivalent) and then shift the non-symmetric, component-coupling part, $\mu^* \nabla \mathbf{w} : \nabla^\top \mathbf{u}$, to the right-hand side. What remains to be assembled into the stiffness matrix is then an integral $\langle \nabla \mathbf{w}, \mu^* \nabla \mathbf{u} \rangle$ for both forms, but the contribution to the right-hand side is either $\langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu^* \rangle$ for the generalised Laplacian form or $-\langle \nabla \mathbf{w}, \mu^* \nabla^\top \mathbf{u} \rangle$ for the stress-divergence form.

As the results summarised in Fig. 3.7 indicate, all variants are stable, yielding low iteration counts given small enough time steps up to $\text{CFL}_e \approx 8.5$. Iteration counts for the two variants do not differ much prior to that point, while considering individual components has a much bigger influence. At $\text{CFL}_e \approx 8.5$, the decoupled variant of the

stress-divergence form diverges, which does not occur using the generalised Laplacian form. Depending on the application at hand, the linearised terms allowing for a split into d scalar-valued subproblems might be more (or even less) important and thus lead to different maximum allowed element CFL numbers before the linear solvers give in.

When a second-order scheme is employed, that is, BDF2 timestepping combined with linear extrapolation, the maximum reachable CFL number is reduced as can be seen from Fig. 3.8. By design of this numerical test, a maximum $\text{CFL}_e \approx 9$ can be achieved, but the FGMRES solver diverges shortly after. However, same trends as observed using the first-order scheme are seen: first, the iteration counts in the viscosity projection and PPE solves are nicely bounded, while the fully explicit convective term reduces the admissible time step size drastically (reaching $\text{CFL}_e \approx 2$ with temporal instabilities clearly visible in the solution). Second, decoupling the velocity components leads to a reduced mean iteration count and the decoupled version of the stress-divergence form diverges earlier than the decoupled version of the generalised Laplacian variant. Third, lumping the mass matrix in the viscosity projection step has yet again no effect on the remaining algorithm in terms of iteration counts. That is, altogether, the same trends as in the previous series of tests are observed.

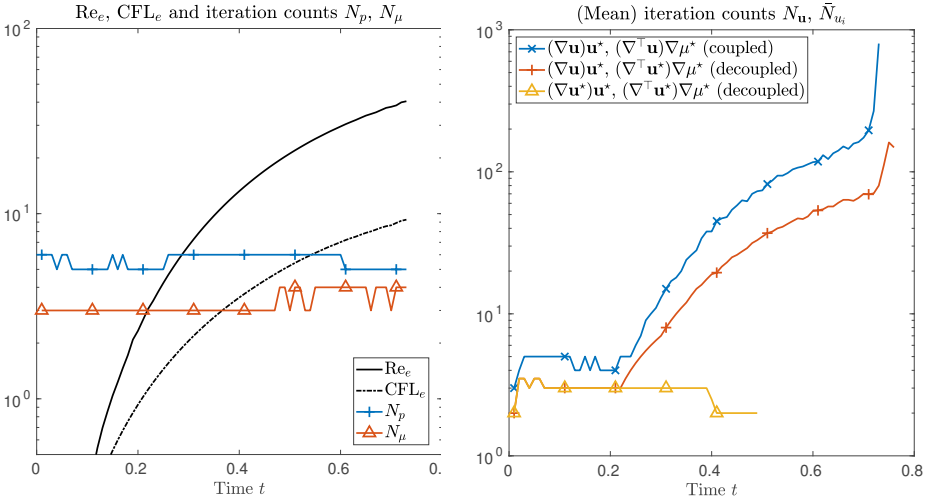


Fig. 3.8: Lid-driven cavity flow: element Reynolds and CFL numbers, iteration counts of the PPE (N_p) and viscosity projection (N_μ) steps (left) and comparison of convective and viscous term variants (right) using BDF2 timestepping and Q_1Q_1 elements.

Lastly, the second-order accurate scheme, i.e., BDF2 and linear extrapolation is combined with Q_2Q_1 elements. The spatial discretisation uses a quarter of the elements only, but has the same number of velocity DoFs as the tests using Q_1Q_1 elements—the element CFL number as defined in (3.38) incorporates the polynomial degree by halving the element length in each direction. Similar trends as in the already presented computational experiments with BDF2 and linear extrapolation are observed. Fig. 3.9 shows iteration counts for the viscosity projection and PPE steps staying nicely bounded, while

at $\text{CFL}_e \approx 9$ also the most stable linearised variants of the scheme diverge. Lumping the mass matrix in the viscosity projection step has again no influence on the observed iteration counts or maximum reachable CFL_e , and the decoupled variants lead to lower iteration counts while not introducing too strict time step constraints. The decoupled variant of the stress-divergence again diverges slightly earlier than the generalised Laplacian counterpart.

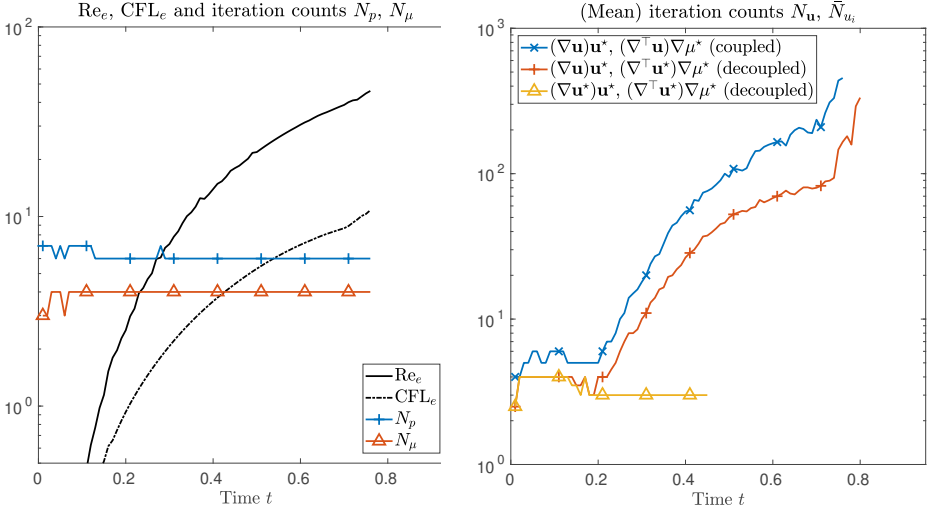


Fig. 3.9: Lid-driven cavity flow: element Reynolds and CFL numbers, iteration counts of the PPE (N_p) and viscosity projection (N_μ) steps (left) and comparison of convective and viscous term variants (right) using BDF2 timestepping and Q_2Q_1 elements.

Summarising the results obtained in these tests, we see that a standard CFL condition has to be fulfilled when using any variant of the split-step scheme. Comparing the first and second-order accurate schemes, we observe greater stability when using BDF1 together with $\mathbf{u}^* = \mathbf{u}^n$ and similar first-order accurate extrapolations, but still reasonable time step sizes for BDF2 with linear extrapolations. These stability constraints are well known also from *coupled*, but *linearised* velocity-pressure formulations, as we considered in Sec. 2.2.

Starting from a linearised convective velocity in a semi-implicit or a fully explicit way, the admissible time step size is already limited to a certain size. Some of the newly introduced linearisations further restrict the time step size, but the maximum reached element CFL numbers of $O(10)$ when combining a semi-implicit convective term with any time discretisation, linearisation and decoupling variant are promising. The different variants of linearising the convective term, however, turn out to perform rather different, where the most popular variant $(\nabla \mathbf{u}) \mathbf{u}^*$ was superior to the other semi-implicit option $(\nabla \mathbf{u}^*) \mathbf{u}$. A fully explicit convective term, $(\nabla \mathbf{u}^*) \mathbf{u}^*$, restricts the time step size quite drastically, but might end up being competitive if the non-symmetric viscous contribution is considered explicit as well. Doing so, a CG instead of FGMRES solver can be employed, reducing the iteration counts in our tests.

The explicit treatment of the non-symmetric and component-coupling viscous term, $(\nabla^\top \mathbf{u}) \nabla \mu^*$ allows decomposing the vector-valued advection-diffusion-reaction equation, that is, the momentum balance equation into d components. This reduces iteration counts while not introducing an additional time step constraint and might thus be a worthwhile strategy depending on the problem setup. The stress-divergence form tends to perform similar to the generalised Laplacian one, while of course the resulting natural boundary condition differs. In the present tests, the decoupled variant of the stress-divergence formulation lead to earlier divergence compared to its generalised Laplacian counterpart—an observation which might be problem dependent and is judged rather minor at this point.

Most importantly, both Q_2Q_1 and Q_1Q_1 elements delivered good results and were found stable over a wide range of time step sizes and CFL numbers. Judging from our tests, equal-order interpolation does not exhibit any spurious oscillations, indicating that an inf-sup condition does not apply as observed by other authors [154, 178–183]. That being said, it is also clear that we lack mathematical analysis regarding PPE-based split-step schemes, which is indeed a pressing issue lying beyond the scope of this work.

3.5.3 Flow through an idealised aneurysm

In this numerical test, we consider a three-dimensional problem in the cardiovascular context, namely, pulsatile flow through an idealised abdominal aortic aneurysm (AAA). All parameters considered are in the physiological range, such that the behaviour of the split-step scheme in realistic haemodynamic scenarios can be investigated. A Carreau fluid, obeying to the rheological law as given in Eqn. (3.40) with upper and lower viscosity limits $\eta_0 = 50$ mPa s and $\eta_\infty = 5$ mPa s, fitting parameters $b = 0.25$ and $\lambda = 10$ s and a density of 1000 kg/m³ resembling blood is employed.

A hexahedral mesh is constructed in a cylindrical domain from $x = 0$ to $L = 20$ cm, with its axis identical to the x -axis and a radius of $R = 1$ cm. Owing to the simple geometric description, a boundary layer mesh is easily designed in this initial configuration, which is then transformed by adapting the radial coordinate $\hat{r} := \sqrt{\hat{y}^2 + \hat{z}^2}$ of each grid node according to

$$r \leftarrow \hat{r} \left(1 + 2 \sin^{10} \frac{\pi x}{L} \right)$$

to obtain the final geometry. The resulting mesh features circular in- and outlet sections and an expanded region with a three-fold increase in the radius, as can be seen from Fig. 3.10. This medium spatial resolution together with boundary layers already leads to roughly 4.9×10^5 finite elements with $\approx 5 \times 10^5$ nodes, which are distributed to 14 subdomains. With the split-step scheme in its variant decoupling the 3 velocity components, the number of nodes is identical to the size of the systems that are solved, if Q_1Q_1 interpolation is employed. Using the coupled (monolithic) velocity-pressure form presented in Sec. 2.2, the same mesh also using Q_1Q_1 interpolation leads to a velocity-pressure system with $\approx 2 \times 10^6$ unknowns (not taking the decoupled viscosity into account), whereas Q_2Q_1 elements lead to $\approx 12.4 \times 10^6$ DoFs. Of course, the solution quality is likely to

change drastically considering these different schemes, but when forced to a specific element size due to geometric constraints in mesh generation, running simulations with such a given fixed spatial discretisation is considerably cheaper following the split-step approach.




Fig. 3.10: Blood flow in idealised aneurysm: mesh of $\approx 4.9 \times 10^5$ hexahedra and $\approx 5 \times 10^5$ nodes distributed to 14 subdomains.

The time step size is adapted in order to achieve a target CFL number of $\text{CFL}_{\max} = 0.3$ in Eqn. (3.37), where no maximum time step size is enforced, i.e., $\Delta t_{\max} = \infty$. A parabolic inflow profile is prescribed at $x = 0$, ramping the inflow smoothly in time via

$$\mathbf{u} = (\hat{u}_1, 0, 0)^\top, \text{ where } \hat{u}_1 = \eta_r \left(1 - 0.3 \cos \frac{\pi t}{0.5} \right) \left(1 - \frac{r^2}{R^2} \right),$$

and η_r defined as in (3.39) with $T_r = 0.5$ s starting from the quiescent state, $\mathbf{u} = \mathbf{0}$, taken as initial condition. At the vessel's walls (i.e., at $\hat{r} = R$ before mapping the mesh), no-slip conditions are enforced, while a zero Neumann condition in the generalised Laplacian formulation is employed to prescribe $\tilde{\mathbf{t}} = \mathbf{0}$ at the outlet. A time step size in the range of 0.5 ms results, while the Reynolds number defined via inlet diameter and spatially mean flow velocity is $\text{Re} = \rho \hat{u}_1 R / \eta_{\infty} \approx 2000$.

The pulsatile flow is computed in the interval from $t = 0$ to $t = 6$ s, recording the maximum element CFL number CFL_e as defined in (3.38), the maximum inlet velocity scaling \hat{u}_1 , the time step size Δt , the maximum element Reynolds number Re_e according to (3.41) and iteration counts in the pressure Dirichlet data projection step (N_{ζ}), PPE solve (N_p), viscosity projection step (N_{μ}) and mean iteration counts over the 3 scalar problems governing the velocity components (\bar{N}_{u_i}).

Results are shown in Fig. 3.11, where one observes nicely bounded iteration counts for all involved linear solvers. The projection steps need an almost constant number of iterations throughout the entire time interval, hinting at the good performance of the AMG-preconditioned mass matrix solve via a CG method also for grids with high aspect ratios as the present one. The projection of pressure boundary data is from a computational cost perspective completely negligible, though, since less than 1% of the overall computing time is spent for assembly and solve of the projection variable ζ . Given that a boundary projection is the only step in the scheme which might be non-trivial, the only motivation for lumping this mass matrix solve might stem from easier implementation. Due to the small time step size resulting from $\text{CFL}_e < 0.3$, the momentum balance steps show remarkably low iteration counts. Despite the convective term considered semi-implicit as $\langle \mathbf{w}, \rho (\nabla \mathbf{u}) \mathbf{u}^* \rangle$, only very mild effects linked to a change in the inlet profile and corresponding Reynolds number are observed. The number of iterations needed to solve the PPE step (N_p) lies around ≈ 40 for the whole simulation period and varies mildly with the inflow, which is caused by the variable right-hand side of the corresponding linear system.

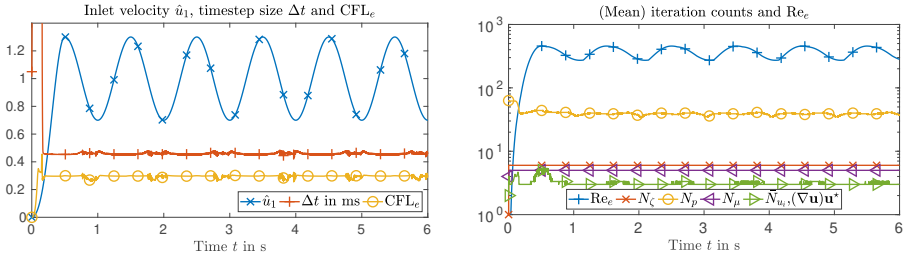


Fig. 3.11: Blood flow in idealised aneurysm: inlet velocity scale \hat{u}_1 and time step size Δt for $\text{CFL}_e < 0.3$, element Reynolds number Re_e and iterations in boundary projection step N_c , PPE solve N_p , viscosity projection N_μ and mean over velocity components solves \bar{N}_{u_i} using BDF2 and Q_1Q_1 elements.

In addition to the obtained iteration counts and element-local CFL and Reynolds numbers, we present the pressure and viscosity at three distinct points $A = (L/2, 3R, 0)^\top$ lying at the apex point in the pipe's center, point $B = (L/2, 3R/\sqrt{2}, 3R/\sqrt{2})^\top$, also at the apex but rotated by 45° into the first quadrant and $C = (L, 0, 0)^\top$ located at the outlet's center. In Fig. 3.12, we can clearly see that a zero pressure is enforced on the outlet via $\bar{\mathbf{t}} = \mathbf{0}$ and the two points A and B , both lying at $x = L/2$, show identical pressure values, smoothly varying in time with the inlet velocity scale. The viscosity values in A and B are identical once the periodic state is reached, after which they hit the upper viscosity limit η_0 periodically. The viscosity in the apex points only differs during an initial period of first intense recirculations, when the fully ramped inflow is reduced rapidly. At the outlet center point C , the apparent viscosity also fluctuates in time given the variable shear rate.

As already indicated by these point values, the solution for this Reynolds number is rotationally symmetric and reaches a periodic state (if the solver tolerances are low enough and the spatial and temporal errors are small). Due to the periodic inflow,

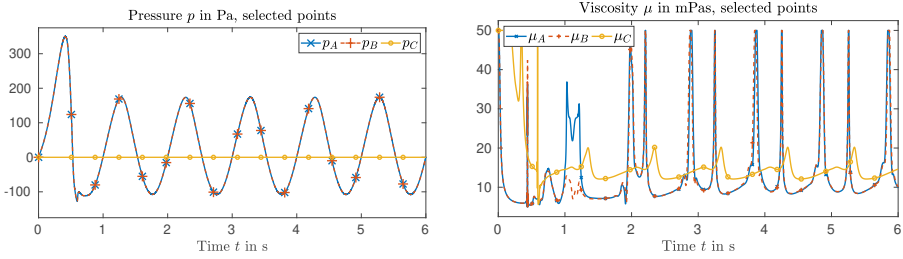


Fig. 3.12: Blood flow in idealised aneurysm: time history of pressure (left) and viscosity (right) at selected points A , B , and C using BDF2 and Q_1Q_1 elements.

recirculating flow is observed in the aneurysm, leading to high velocity gradients, strong differences in the local shear rate and thus to large variations in viscosity. Snapshots of the solution in the systolic phase of the sixth considered cycle at $t = 5.37$ s are depicted in Fig. 3.13, showing selected streamlines colored by the velocity norm and the pressure in half of the domain.

While a more sophisticated in-depth analysis regarding the importance of rheological modelling in such a haemodynamic scenario lies beyond the scope of this work, we restrict ourselves to sheer observations of changes in viscosity throughout the pulse cycles in space. Fig. 3.14 shows a slice through the domain at $y = 0$ at $t = 5.37$ s, that is, at the same time as the snapshots presented in Fig. 3.13. Again, strong recirculations can be observed, leading to large variations in the viscosity.

These patterns can be compared to cut contours of the viscosity in the diastolic phase of the final cycle at $t = 5.83$ s, where the viscosity spans the whole range of admissible values, $\eta_\infty \leq \mu \leq \eta_0$, with the highest values present at the apex of the aneurysm and the inviscid viscosity limit reached in the bulk of the domain. In the region where the aneurysm expands, the flow (from right to left in Fig. 3.15) recirculates, such that a zone with lower shear rates and thus an increased viscosity persists.

3.5.4 Computational performance study

At this point, it is interesting to compare the split-step method with the BVS scheme from Sec. 2.2 and standard Taylor–Hood discretisation. To reduce the computing times, we settle for a slightly simpler setup considering only two pulses within the interval $I_t = (0, 1]$ with a fixed time step size of $\Delta t = 1$ ms while the remaining problem parameters are unchanged from the preceding example in Sec. 3.5.3. Also, a smaller mesh of 2.1×10^5 for Q_1Q_1 interpolation and half as much elements when using the inf-sup stable Q_2Q_1 pairing is considered. Employing these spatial discretisations, the number of velocity DoFs is the same, whereas the coarser mesh with Q_2Q_1 elements features much less nodes to construct the linear space compared to the former case with equal-order interpolation.

The coupled velocity-pressure formulations split the viscosity projection step from the overall system via extrapolation, linearise the convective term via $(\nabla \mathbf{u}) \mathbf{u}^*$ as well, but




Fig. 3.13: Blood flow in idealised aneurysm: streamlines (top) and pressure in half of the domain (bottom) at $t = 5.37$ s using BDF2 and Q_1Q_1 elements.




Fig. 3.14: Blood flow in idealised aneurysm: velocity vectors colored by $\|\mathbf{u}\|$ and viscosity in $y = 0$ at time $t = 5.37$ s using BDF2 and Q_1Q_1 elements with the main flow direction from right to left.




Fig. 3.15: Blood flow in idealised aneurysm: selected viscosity contours in systole at $t = 5.83$ s using BDF2 and Q_1Q_1 elements with the main flow direction from right to left.

solve for the velocity and pressure fields in a monolithic system of the form

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \mathbf{u}^{k+1} \\ \mathbf{p}^{k+1} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix},$$

employing an FGMRES method [135] with a block-triangular right preconditioner [106]

$$\mathcal{P}^{-1} := \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{0} & \mathbf{S} \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & -\mathbf{B} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{pmatrix},$$

with the pressure Schur complement $\mathbf{S} := \mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$, exactly as detailed in Sec. 2.5. Regarding the approximation of the inverse Schur complement, however, we consider now two distinctly different preconditioning options. While the BVS formulation uses an adapted variant of the PCD preconditioner (see Sec. 2.5 for a detailed discussion), using a Taylor–Hood pair, on the contrary, allows consistently choosing a grad-div stabilisation parameter $\gamma_e = \mathcal{O}(1)$ in Eqn. (2.66) [27], such that a grad-div preconditioner accounting for a variable viscosity combining ideas of [129, 134] can be used. So, in the latter case, we approximate the pressure Schur complement via

$$\mathbf{S}^{-1} \approx \mathbf{M}_{\mu,\gamma}^{-1} + \mathbf{L}_p^{-1},$$

with matrices defined in the pressure space (2.89) and (2.5). For details on the preconditioner design, we refer to Sec. 2.2 and [104, 106, 124, 129, 132, 134, 136].

In the preconditioner, all the inverses are approximated via single AMG V-cycles, again, exactly as in Sec. 2.2. The main goal is to compare these standard approaches and the performance that can be achieved via open-source linear algebra packages given a fixed discretisation. All implementations are yet again based on `deal.II` [96–99] and interface Trilinos’ ML package [111] to precondition linear solves reducing the residual by a factor of 10^5 from the last time step’s solution taken as initial guess.

In Fig. 3.16, we report the maximum element Reynolds number reaching a value of almost $\text{Re}_e \approx 600$, the maximum element CFL number smaller than 0.5 and the resulting iteration counts in the velocity-pressure formulations ($N_{u,p}$). For the PPE-based split-step scheme, the iteration counts in the PPE step (N_p) and the *sum* of all iterations spent in the individual velocity component solves ($\sum_i N_{u_i}$) are reported.

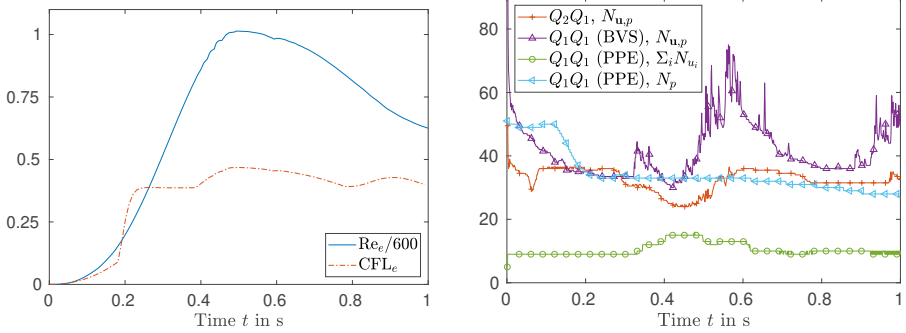


Fig. 3.16: Scaled element Reynolds and CFL numbers (left) and iteration counts $N_{u,p}$ for the velocity-pressure system with Q_2Q_1 or Q_1Q_1 elements compared to the PPE-based approach (right).

The recorded iteration counts are initially higher for all solvers/preconditioners used due to the relative convergence criterion and show a mild dependence on the Reynolds number for the velocity-pressure formulations. The grad-div preconditioner outperforms the PCD preconditioner, showing lower iteration counts and seemingly reduced sensitivity with respect to the flow regime. The split-step scheme shows very mild dependence on the Reynolds number in the momentum balance steps, which features low iteration counts around 10-15, whereas the pressure Poisson solve is in this numerical test unaffected by an increase in the volumetric flow rate, but results in ≈ 35 preconditioned AMG steps, once the initial ramp-up phase is completed.

In an attempt to compare the different solver paradigms, we count individual AMG V-cycles needed during the simulation, neglecting assembly times, preconditioner setup and actual linear solver used, as those are approximately equal or turn out to favour the monolithic solvers. Assuming that all mass matrix problems may be lumped in the approximations of \mathbf{S}^{-1} gives two AMG cycles per preconditioner application (\mathbf{A}^{-1} and \mathbf{L}_p^{-1}) for the monolithic approach with grad-div preconditioner, and three (\mathbf{A}^{-1} and two times \mathbf{L}_p^{-1}) for the BVS Q_1Q_1 approach—replacing the matrix product in Eqn. (2.85), $\mathbf{C} \text{diag}(\rho \mathbf{M}_u)^{-1} \mathbf{B}$, simply by \mathbf{L}_p given large enough, uniform time steps. Each step in the split-step scheme uses an AMG-preconditioned CG or FGMRES solve, such that each iteration on the d velocity components and the pressure Poisson problem account for 1 AMG V-cycle, while the viscosity projection step’s mass matrix solve is lumped.

Tab. 3.2 lists the number of cycles in absolute and relative terms, showing slightly (Q_2Q_1) to significantly (Q_1Q_1) higher counts for the monolithic schemes compared to the *summed* individual steps in the split-step approach. The number of AMG V-cycles on any block, however, does not take into account, how complex solving the velocity-velocity block

with a denser sparsity pattern due to a higher polynomial degree is. Despite the fact that the PCD preconditioner needs more linear solver iterations and more AMG V-cycles per preconditioner application, it still turns out to be faster than its grad-div preconditioned alternative. Also, the increased ratio between velocity to pressure DoFs plays a role, as an optimal number of cores for a fixed problem size (number of elements) differs for scalar problems and the coupled, vector-valued and possibly second-order velocity field. Optimal core counts were used for the comparison, that is in the current implementation 12 cores for Q_2Q_1 interpolation and 10 cores otherwise, which lead to the smallest computing times listed. Additionally, the recorded timings rather show the influence of actual system sizes than overall DoF count on the achieved timings, which does not come as a surprise.

Tab. 3.2: Comparison of monolithic Q_2Q_1 and BVS Q_1Q_1 approaches with Q_1Q_1 ansatz in the split-step scheme: mean AMG V-cycles per time step (top row) and time spent in linear solver routines (bottom row).

	Q_2Q_1 , grad-div	Q_1Q_1 (BVS), PCD	Q_1Q_1 , split-step
mean AMG V-cycles	65.06 (144%)	125.29 (278%)	45.12 (100%)
linear system solve time	9186 s (253%)	8196 s (226%)	3633 s (100%)

Summing up, given small enough time steps, all schemes perform decently well, with the split-step scheme being significantly faster. Note here also that for high Reynolds number flows, the construction of a suitable preconditioner, especially the Schur complement, is more involved and might alter the obtained results in favour of the split-step scheme. When suitably tailored multigrid methods are used, or in different flow regimes, however, monolithic formulations are competitive nonetheless. Comparing the split-step scheme to a nonlinear solver as presented in Sec. 2.2, an additional speed-up of around 2 to 3 can be expected. An advantage inherent to the split-step scheme is the fact that all involved steps are standard problems, for which tailored tools are easily accessible and are continuously improved. On the other hand, fully implicit, coupled formulations do not suffer from severe time step restrictions, such that a general statement is not easily formulated. Depending on the target application and (especially) on the implementation, any of the considered schemes might yield the best results in terms of computing time. In terms of accuracy, however, inf-sup stable, coupled and fully implicit formulations (possibly with discontinuous pressure approximation) are most often preferred, if they are not found too costly.

3.6 Summary and conclusion

The PPE-based split-step scheme presented within this chapter completely replaces the continuity equation with a consistent PPE. It incorporates a suitable consistent Dirichlet condition for the pressure, rendering the PPE uniquely solveable given velocity. Extrapolating in (pseudo-)time, the equations governing viscosity, velocity and pressure are

now decoupled, which allows recovering them independently. Similar to the coupled approach, all fields are interpolated using standard C^0 -continuous finite elements. Also, equal-order interpolation is admissible, which distinguishes the current scheme from related projection methods. To further improve mass conservation, we employ divergence damping, that is, a Helmholtz–Leray decomposition of the past time step’s velocities only, sparing us an additional projection of the divergence-free fluid velocity onto a continuous space. With that, a single scalar-valued Poisson problem per time step needs to be solved additionally, but combining the PPE and projection variable into a single pressure-like quantity leads to a scheme with no added cost. Compared to the Newtonian base variant, we thus merely add a scalar-valued mass-matrix problem projecting the viscosity. An iteration-free method results, involving only three computationally relevant linear systems: the momentum balance, the PPE and viscosity projection steps.

Several numerical experiments of academic nature demonstrate accuracy and robustness of the presented algorithms, showcasing the expected convergence rates. A test applying the split-step scheme to an idealised abdominal aortic aneurysm highlights the scheme’s potential for practical application to haemodynamics.

A comparative study of the coupled velocity-pressure formulation and the PPE-based split-step scheme then further revealed that the latter has greater potential for application in large-scale problems given the fact that the momentum balance and PPE steps are easier completed individually than obtaining velocity and pressure through a coupled system. This is of course heavily dependent on the implementation and preconditioner choice and thus might not hold true applying more elaborate techniques. However, especially when having FSI in mind, the PPE-based split-step scheme has an additional advantage: owing to the decoupling, velocity and pressure can be recovered individually, such that added-mass stable, semi-implicit coupling is enabled! Within such semi-implicit FSI schemes, only the fluid pressure and structural displacement are coupled iteratively at each time step, while the remaining subproblems are treated explicitly and hence only solved once per time step. Consequently, the computational costs associated with a single iteration in the FSI coupling loop decrease tremendously. This work aims at exploiting this enormous potential when coupling the split-step scheme to an elastic structure (see Ch. 7). The monolithic velocity-pressure formulation is not considered within an FSI context given that such approaches are frequently found in literature.

With this, the basic forms of the coupled velocity-pressure formulation as presented in the previous chapter and the related PPE-based split-step scheme are introduced. We thus turn our attention to further including relevant modelling aspects and numerical techniques to solve practical problems in the cardiovascular and general biomedical context in the following chapter.

4 Incompressible viscous flows: extensions to biomedical applications

The previous chapters already demonstrated the basic properties and indicated applicability to real-world applications in the haemodynamic context of the BVS formulation and the split-step scheme. When tackling more involved geometries than the idealised ones already considered, possibly derived from medical image data, several aspects need to be covered in order to refer to simulations as “patient-specific”. This chapter aims to summarise the needed numerical techniques and modelling aspects concerning the fluid solver to tackle problems of clinical interest. First, the extension of the Navier–Stokes system to the moving domain case, adopting an Arbitrary Lagrangian–Eulerian (ALE) formulation, is discussed in Sec. 4.1. Then, Robin boundary conditions—of great interest in haemodynamic FSI—are introduced in Sec. 4.2. Third, the stabilisation of dominant convective effects via suitable residual-based stabilisations is introduced in Sec. 4.3, before we finally cover suitable in- and outlet conditions in Secs. 4.4–4.6 by constructing inflow profiles on non-circular cross-sections, stabilising instabilities triggered by re-entrant flow and modelling the neglected downstream vasculature.

The ALE formulation of the PPE-based split-step scheme follows the presentation in the author’s work [5], while herein, we give a more rigorous introduction to ALE methods. The remainder of this section concerning further extensions (of the split-step scheme in particular) towards practical application in the biomedical context can be found in Schussnig et al. [6]. Within this work, these mentioned facets are introduced in general form, applicable to both the PPE-based split-step scheme and coupled velocity-pressure formulation with BVS. However, we refrain from presenting results within this section, since the bulk of the methods presented here aim at FSI formulations to be considered in Ch. 7.

4.1 Flows in moving domains

By far the most popular application of flows in moving domains is FSI, since a central part thereof is *always* a formulation of the fluid problem able to account for the changing fluid domain. Thus, a large portion of the literature regarding such schemes has been developed in direct or indirect relation to FSI problems. Some of the most popular techniques to cope with moving domains are ALE formulations [187–194], immersed boundary techniques [195–200] and fictitious-domain methods [201–204], either tracking or capturing the motion of the fluid–structure interface, or, from a more general viewpoint: accounting for the moving fluid domain. One might also refer to these

techniques depending on how the boundary of the fluid domain is considered, resulting in a distinction of fitted and unfitted methods.

ALE methods date back to the work by Hirt et al. [205] and were later on employed in a finite element context of the Navier–Stokes equations for compressible [194, 206] and incompressible [188] flow. For a comprehensive overview, the reader is referred to the recent survey by Donea et al. [207], while a summary of the early days may be found in [194] and an introduction can be found in the monograph by Donea and Huerta [39].

The fundamental concept of the ALE method is to neither take the Lagrangian frame with coordinate $\hat{\mathbf{x}}$, nor the spatial configuration with $\mathbf{x}(\hat{\mathbf{x}}, t)$ as reference, but a third, the so-called ALE configuration in terms of $\tilde{\mathbf{x}}$. Our starting point is the momentum balance equation in Eulerian form as given in Eqn. (2.1), repeated here for the convenience of the reader,

$$\rho \partial_t \mathbf{u} + \rho (\nabla \mathbf{u}) \mathbf{u} - \nabla \cdot \boldsymbol{\sigma} = \mathbf{b}, \quad (4.1)$$

where the time derivative $\partial_t \mathbf{u}$ is a shorthand notation for $\partial_t \mathbf{u}|_{\mathbf{x}}$, the spatial time derivative, meaning the time derivative at a point “holding \mathbf{x} fixed”. An interpretation that allows formulating balance equations in a fixed reference frame, which is the natural viewpoint in fluid mechanics problems, as one is typically interested in quantities at fixed points in space, that is, \mathbf{x} in Ω . On the contrary, the material time derivative,

$$d_t \mathbf{u} := \partial_t \mathbf{u}|_{\hat{\mathbf{x}}}, \quad (4.2)$$

describes the time derivative as seen by an individual particle, located at $\hat{\mathbf{x}}$ in the Lagrangian frame $\hat{\Omega}$. The particle-centered viewpoint in Eqn. (4.2) is natural in solid mechanics—with the aim to identify stresses in certain points in the Lagrangian frame, i.e., the “undeformed” configuration—and thus more often applied in this context, as it additionally renders the momentum balance equation free of any convective term,

$$\rho d_t \mathbf{u} - \nabla \cdot \boldsymbol{\sigma} = \mathbf{b}, \quad (4.3)$$

since we have

$$\frac{\partial \mathbf{u}(\mathbf{x}(\hat{\mathbf{x}}, t), t)}{\partial(\hat{\mathbf{x}}, t)} = \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial(\mathbf{x}, t)} \frac{\partial(\mathbf{x}, t)}{\partial(\hat{\mathbf{x}}, t)} = \left(\frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial \mathbf{x}} \quad \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} \Big|_{\mathbf{x}} \right) \begin{pmatrix} \frac{\partial \mathbf{x}}{\partial \hat{\mathbf{x}}} & \frac{\partial \mathbf{x}}{\partial t} \Big|_{\hat{\mathbf{x}}} \\ \mathbf{0} & 1 \end{pmatrix},$$

where multiplication gives the standard relation

$$\frac{\partial \mathbf{u}(\mathbf{x}(\hat{\mathbf{x}}, t), t)}{\partial \hat{\mathbf{x}}} = \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \hat{\mathbf{x}}},$$

and the sought one, connecting material and spatial time derivatives

$$\frac{\partial \mathbf{u}(\mathbf{x}(\hat{\mathbf{x}}, t), t)}{\partial t} \Big|_{\hat{\mathbf{x}}} = \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial t} \Big|_{\hat{\mathbf{x}}} + \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} \Big|_{\mathbf{x}}.$$

Therein, $\partial_t \mathbf{x}|_{\tilde{\mathbf{x}}}$ is the particle velocity, such that we can rewrite in the typical notation

$$\partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} = \mathbf{d}_t \mathbf{u} = (\nabla \mathbf{u}) \mathbf{u} + \partial_t \mathbf{u}|_{\mathbf{x}},$$

such that the connection between Eqns. (4.1) and (4.3) becomes obvious.

In an ALE framework, we might go a similar route, expressing the spatial coordinate in terms of the ALE coordinate, that is $\mathbf{x} = \mathbf{x}(\tilde{\mathbf{x}}, t)$ and go through the same steps once again, yielding

$$\frac{\partial \mathbf{u}(\mathbf{x}(\tilde{\mathbf{x}}, t), t)}{\partial(\tilde{\mathbf{x}}, t)} = \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial(\mathbf{x}, t)} \frac{\partial(\mathbf{x}, t)}{\partial(\tilde{\mathbf{x}}, t)} = \left(\frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial \mathbf{x}} \quad \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} \Big|_{\mathbf{x}} \right) \begin{pmatrix} \frac{\partial \mathbf{x}}{\partial \tilde{\mathbf{x}}} & \frac{\partial \mathbf{x}}{\partial t} \Big|_{\tilde{\mathbf{x}}} \\ \mathbf{0} & 1 \end{pmatrix}.$$

In analogy to the connection between Eulerian and Lagrangian reference frames, we now obtain a relation between the ALE time derivative $\partial_t \mathbf{u}|_{\tilde{\mathbf{x}}}$ and the spatial time derivative $\partial_t \mathbf{u}|_{\mathbf{x}}$ as

$$\frac{\partial \mathbf{u}(\mathbf{x}(\tilde{\mathbf{x}}, t), t)}{\partial t} \Big|_{\tilde{\mathbf{x}}} = \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial t} \Big|_{\tilde{\mathbf{x}}} + \frac{\partial \mathbf{u}(\mathbf{x}, t)}{\partial t} \Big|_{\mathbf{x}}.$$

We can rearrange the above relation to give

$$\partial_t \mathbf{u}|_{\mathbf{x}} = \partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} - (\nabla \mathbf{u}) \partial_t \mathbf{x}|_{\tilde{\mathbf{x}}}, \quad (4.4)$$

with the term $\partial_t \mathbf{x}|_{\tilde{\mathbf{x}}}$, which we can interpret as follows: For a given point in the ALE configuration, $\tilde{\mathbf{x}}$ fixed, the time derivative of the position $\mathbf{x}(\tilde{\mathbf{x}}, t)$ is nothing more than the velocity of the ALE configuration itself. For ease of notation, we thus define the mesh velocity,

$$\mathbf{u}_m(\tilde{\mathbf{x}}, t) := \partial_t \mathbf{x}|_{\tilde{\mathbf{x}}}, \quad (4.5)$$

and finally rewrite the momentum balance equation starting from the Eulerian form as given in Eqn. (4.1) by inserting (4.4). This yields

$$\rho \partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} + \rho (\nabla \mathbf{u}) \mathbf{c} - \nabla \cdot \boldsymbol{\sigma} = \mathbf{b} \quad \text{with} \quad \mathbf{c} := \mathbf{u} - \mathbf{u}_m. \quad (4.6)$$

Inspecting this relation closer, it becomes obvious where the term ‘‘arbitrary’’ in ALE originates from: choosing $\tilde{\mathbf{x}} \equiv \mathbf{x}$, one ends up with the Eulerian formulation, as $\mathbf{u}_m = \mathbf{0}$. Taking $\tilde{\mathbf{x}} \equiv \hat{\mathbf{x}}$ on the other hand, we end up with the material time derivative $\mathbf{d}_t \mathbf{u}$ and the convective term vanishes due to $\mathbf{u} \equiv \mathbf{u}_m$. Note here that the presented derivations do not alter spatial operators, as would be common in a fully Lagrangian setting typical for solid mechanics.

In this setup, the ALE configuration can be chosen freely, not following the particle motion, but also not fixed as in an Eulerian setting. Denoting the time-dependent fluid domain by Ω^t , e.g., transformed from the initial configuration $\Omega^{t=0}$ by a given map, we choose to follow exactly this (prescribed) motion.

Then, we can construct a mapping $\phi(\hat{\mathbf{x}}, t)$ from the initial configuration $\Omega^{t=0}$, which is identical to the Lagrangian reference frame $\hat{\Omega}$, with the help of the displacement $\mathbf{d}(\hat{\mathbf{x}}, t)$ in $\hat{\Omega}$ as

$$\phi : \hat{\mathbf{x}} \in \hat{\Omega} \rightarrow \mathbf{x} \in \Omega^t \quad , \quad \phi(\hat{\mathbf{x}}, t) := \hat{\mathbf{x}} + \mathbf{d}(\hat{\mathbf{x}}, t),$$

as is typical in solid mechanics. Recasting this in terms of $\tilde{\mathbf{x}}$, we see that

$$\begin{aligned} \mathbf{u}_m &:= \partial_t \mathbf{x}(\hat{\mathbf{x}}(\tilde{\mathbf{x}}, t), t)|_{\tilde{\mathbf{x}}} \\ &= \partial_t \hat{\mathbf{x}}(\tilde{\mathbf{x}}, t)|_{\tilde{\mathbf{x}}} + \partial_t \mathbf{d}(\hat{\mathbf{x}}(\tilde{\mathbf{x}}, t), t)|_{\tilde{\mathbf{x}}} \\ &= \partial_t \mathbf{d}(\hat{\mathbf{x}}(\tilde{\mathbf{x}}, t), t)|_{\tilde{\mathbf{x}}}, \end{aligned}$$

since $\partial_t \hat{\mathbf{x}}(\tilde{\mathbf{x}}, t)|_{\tilde{\mathbf{x}}} = 0$, meaning that the Lagrangian coordinate to which a moving grid point corresponds to does not change in time. By design, we also have $\tilde{\mathbf{x}}|_{t=0} = \hat{\mathbf{x}}$ as we start with $\mathbf{d}|_{t=0} = \mathbf{0}$, and hence, the displacement is naturally given in a grid point $\tilde{\mathbf{x}}$. Then, the time derivative $\partial_t \mathbf{d}|_{\tilde{\mathbf{x}}}$ is computed using standard time integration schemes, *without* any further transformations needed, just taking values in grid points.

Discretising Eqn. (4.6) is then carried out almost identical to the standard Eulerian setting, since we can approximate $\partial_t \mathbf{u}|_{\tilde{\mathbf{x}}}$ using grid data as usual, with the only difference being that the grid position is updated each time step. The discretisation in space follows exactly the standard procedure, redirecting the reference element map from the Eulerian configuration to an updated element position in the current domain Ω^t .

Before presenting the ALE formulations of monolithic and split-step schemes, two important aspects need to be addressed: first, one has to note that we assumed all the involved transformations being one-to-one, continuously invertible maps. The coordinate $\mathbf{x}(\tilde{\mathbf{x}}, t)$ is typically constructed by means of a displacement $\mathbf{d}(\hat{\mathbf{x}}(\tilde{\mathbf{x}}, t), t)$ from the initial ALE configuration identical to the Lagrangian frame. Thus, one is typically interested in prescribing $\mathbf{d}|_{\partial\hat{\Omega}}$ such that extending the prescribed boundary displacement into the interior of $\hat{\Omega}$ is a *crucial* aspect of *any* ALE method. There is, however, quite some freedom in constructing a suitable extension leading to a variety of ways available from literature as shall be seen in Sec. 7.5. For now, we simply assume a suitable, i.e., continuous and invertible displacement field $\mathbf{d}(\hat{\mathbf{x}}, t)$ given, such that the balance equations in ALE form are well defined.

A second pressing issue regarding ALE formulations as already identified by Thomas and Lombard [208] is the preservation of a constant solution on a moving grid down to the discrete level, a property which is of such fundamental importance that it has found great attention since then (see, e.g., [209–215]). According to Förster et al. [214], we can fulfil the so-called geometric conservation law (GCL) and achieve the expected temporal order of accuracy by (i) taking the convective ALE form of the Navier–Stokes equations as presented in Eqn. (4.6), (ii) integrating all terms over $\Omega^{t=t^{n+1}}$, denoted as Ω^{n+1} , and (iii) using an identical time integration scheme for the fluid velocity \mathbf{u} and the mesh velocity $\mathbf{u}_m := \partial_t \mathbf{x}|_{\tilde{\mathbf{x}}}$. The conservative formulation of the Navier–Stokes equations involves time derivatives of time-dependent integrals and leads to spatial derivatives of the mesh velocity being present in the weak form. Hence, fulfilling the GCL exactly is rather complex in this scenario. Altogether, ALE formulations based on the convective

form of the Navier-Stokes equations thus only need to modify the convective velocity, merely replacing \mathbf{u} by $\mathbf{u} - \mathbf{u}_m$, update the current domain Ω^t based on a displacement \mathbf{d} (either fully prescribed or extended from the boundary) and respect points (i)–(iii).

In analogy to the Eulerian setting, we arrive at the Navier–Stokes equations for incompressible flow in convective ALE form

$$\rho \partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} + \rho (\nabla \mathbf{u}) \mathbf{c} - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2 \nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b} \quad \text{in } \Omega^t \times (0, T], \quad (4.7)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \Omega^t \times (0, T], \quad (4.8)$$

in the time-dependent domain $\Omega^t \subset \mathbb{R}^d$, $d = 2$ or 3 , with the fluid's density ρ and viscosity μ , the velocity \mathbf{u} and the convective velocity $\mathbf{c} := \mathbf{u} - \mathbf{u}_m$, where \mathbf{u}_m denotes the mesh velocity as given in Eqn. (4.5), pressure p , body force \mathbf{b} and a parameter χ to yield the generalised Laplacian ($\chi = 0$) or stress-divergence forms ($\chi = 1$). We will further decompose the boundary of the domain $\Gamma^t := \partial \Omega^t$ into non-overlapping Dirichlet and Neumann sections denoted by Γ_D^t and Γ_N^t , respectively, such that $\Gamma_D^t \cup \Gamma_N^t = \Gamma^t$ and $\Gamma_D^t \cap \Gamma_N^t = \emptyset$. On these parts of the boundary, we have the familiar conditions

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D^t \times (0, T], \quad (4.9)$$

$$\left(-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \right) \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N^t \times (0, T], \quad (4.10)$$

where $\bar{\mathbf{t}}$ is a given pseudo traction $\bar{\mathbf{t}}$ or real traction \mathbf{t} (see Ch. 2 for details). The problem setup is then completed with the initial condition on the velocity

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0 \quad \text{and} \quad \nabla \cdot \mathbf{u}_0 = 0 \quad \text{in } \Omega^0. \quad (4.11)$$

Extending the BVS velocity-pressure formulation and the split-step scheme to the ALE setting, we follow an identical procedure as in Chs. 2 and 3 for discretisation in space and time, while simultaneously respecting the conditions posed by Förster et al. [214] to fulfil the GCL.

4.1.1 BVS method in ALE form

Similar to Eulerian setting, we combine a suitable PPE and momentum equation, both in ALE form, with the scaled continuity equation, $\beta \nabla \cdot \mathbf{u}$, to get

$$\rho \partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} + \rho (\nabla \mathbf{u}) \mathbf{c} - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2 \nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b} \quad \text{in } \Omega^t \times (0, T], \quad (4.12)$$

$$\beta \nabla \cdot \mathbf{u} + \nabla \cdot [2 \nabla^s \mathbf{u} \nabla \mu - \rho (\nabla \mathbf{u}) \mathbf{c} + \mathbf{b}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = \Delta p \quad \text{in } \Omega^t \times [0, T], \quad (4.13)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D^t \times (0, T], \quad (4.14)$$

$$\left[-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \right] \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N^t \times (0, T], \quad (4.15)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0, \quad (4.16)$$

$$\mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} - \rho (\nabla \mathbf{u}) \mathbf{c} + 2 \nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] = \mathbf{n} \cdot \nabla p \quad \text{on } \Gamma^t \times [0, T], \quad (4.17)$$

which is equivalent to the classical convective ALE form of the Navier–Stokes equations (4.7)–(4.11) as shown next.

Corollary 4.1.1. *For sufficiently regular $p, \mathbf{u}, \mathbf{u}_m, \mathbf{b}, \mathbf{g}, \bar{\mathbf{t}}$, systems (4.7)–(4.11) and (4.12)–(4.17) are equivalent.*

Proof. We can show that system (4.7)–(4.11) implies (4.12)–(4.17) in complete analogy to Thm. 2.3.1 (and hence Thm. 2.1.1) by following the exact same sequence of steps and simply replacing the convective velocity \mathbf{u} by its ALE equivalent $\mathbf{u} - \mathbf{u}_m$ and the Eulerian time derivative $\partial_t \mathbf{u}$ by its ALE counterpart, $\partial_t \mathbf{u}|_{\bar{\mathbf{x}}}$.

The second part of the proof consists of showing that (4.12)–(4.17) imply (4.7)–(4.11). Starting from the BVS method in ALE form (4.12)–(4.17). Taking the divergence of the rewritten momentum balance equation (4.12) and adding it to the PPE (4.13) yields

$$\nabla \cdot (\rho \partial_t \mathbf{u}|_{\bar{\mathbf{x}}}) + \beta \nabla \cdot \mathbf{u} - \nabla \cdot (\mu \Delta \mathbf{u}) - \chi \nabla \cdot [\mu \nabla (\nabla \cdot \mathbf{u})] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = 0,$$

which can be rewritten in a similar manner as in the Eulerian case, yielding

$$\rho \partial_t \nabla \cdot \mathbf{u}|_{\bar{\mathbf{x}}} + \beta \nabla \cdot \mathbf{u} - (1 + \chi) \mu \Delta (\nabla \cdot \mathbf{u}) - (1 + \chi) \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u})] = 0.$$

This is a diffusion-reaction equation in $\Phi := \nabla \cdot \mathbf{u}$, but this time in ALE form,

$$\rho \partial_t \Phi|_{\bar{\mathbf{x}}} + \beta \Phi - (1 + \chi) \nabla \cdot (\mu \nabla \Phi) = 0, \quad (4.18)$$

for which we construct initial and boundary conditions in a similar manner as before,

$$\begin{aligned} \Phi = \nabla \cdot \mathbf{u} = \nabla \cdot \mathbf{u}_0 = 0 & \quad \text{at } t = 0, \\ \mathbf{n} \cdot \nabla \Phi = 0 & \quad \text{on } \Gamma^t, \end{aligned}$$

Thus, $\Phi \equiv 0$ is the only admissible solution, and the alternative system (2.37)–(2.42) also ensures $\nabla \cdot \mathbf{u} = 0$, as requested. All that remains is to show that the modified momentum balance equation implies the original form, which directly follows from Eqn. (2.25), thereby completing the proof. \square

Here, we see that contrary to the standard ALE Navier–Stokes system, which enforces $\nabla \cdot \mathbf{u}$ directly, in the alternative system, the GCL influences mass conservation as will become clear introducing the fully discrete problem in the following. So, we discretise in time by decomposing the time interval $I_t = (0, T]$ into N_t possibly non-uniform time steps, $\Delta t^n = t^{n+1} - t^n$, $n = 0, \dots, N_t$, and employ a generalised θ -scheme, approximating the time derivative as

$$\partial_t \mathbf{u}|_{\bar{\mathbf{x}}} \approx \frac{1}{\Delta t^n} (\mathbf{u}^{n+1} - \mathbf{u}^n) \Big|_{\bar{\mathbf{x}}}.$$

A suitable weak form is found as described in great detail in Secs. 2.3.1 and 2.3.2, where we again just replace the convective velocity by $\mathbf{u} - \mathbf{u}_m$ and the time derivative $\partial_t \mathbf{u}$ by $\partial_t \mathbf{u}|_{\bar{\mathbf{x}}}$. As opposed to a fixed domain Ω , we integrate over the current domain Ω^t and its

respective boundary segments. This leads to a slight modification of the function space definitions, namely

$$X^t := H^1(\Omega^t) \quad \text{and} \quad X_h^t \subset H^1(\Omega_h^t),$$

and consequently updated definitions

$$\begin{aligned} X_h^t(\Omega_h^t) &:= \left\{ q_h \in C^0(\overline{\Omega_h^t}) : q_h \circ \chi_e^t(\boldsymbol{\xi})|_{\Omega_e^t} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{P}^k(\Omega_{\text{ref}}) \quad \forall \Omega_e^t \in \Omega_h^t \right\}, \\ \text{or } X_h^t(\Omega_h^t) &:= \left\{ q_h \in C^0(\overline{\Omega_h^t}) : q_h \circ \chi_e^t(\boldsymbol{\xi})|_{\Omega_e^t} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{Q}^k(\Omega_{\text{ref}}) \quad \forall \Omega_e^t \in \Omega_h^t \right\}, \end{aligned} \quad (4.19)$$

where $\chi_e^t(\boldsymbol{\xi})$ denotes the isoparametric mapping from the reference element Ω_{ref} to a specific element $\Omega_e^t \subset \Omega_h^t$. Furthermore, we will frequently replace the superscript by a time step index i to highlight the specific time, at which the domain $\Omega^{t=t^i}$ is considered. Denoting by $\langle \cdot, \cdot \rangle$ the L^2 scalar product over the current domain $\Omega^{t=t^i}$, we define

$$\begin{aligned} a(\mathbf{w}, \mathbf{u}, \mathbf{c}, p, \mu, t^i) &:= \langle \mathbf{w}, \rho(\nabla \mathbf{u})\mathbf{c} - \mathbf{b}|_{t=t^i} \rangle + (1 - \chi/2) \langle \nabla \mathbf{w} + \chi \nabla^\top \mathbf{w}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle \\ &\quad - (1 - \chi) \langle \mathbf{w}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle - \langle \nabla \cdot \mathbf{w}, p \rangle - \langle \mathbf{w}, \bar{\mathbf{t}}|_{t=t^i} \rangle_{\Gamma_N^{t=t^i}}, \end{aligned}$$

$$\begin{aligned} b(q, \mathbf{u}, \mathbf{c}, p, \mu, t^i) &:= \langle \nabla q, \nabla p + \rho(\nabla \mathbf{u})\mathbf{c} - 2(\nabla^\top \mathbf{u}) \nabla \mu - \mathbf{b}|_{t=t^i} \rangle \\ &\quad + \langle \nabla q \times \mathbf{n}, \mu \nabla \times \mathbf{u} \rangle_{\Gamma^{t=t^i}}, \end{aligned}$$

$$c(q, \mathbf{u}) := \sum_{e=1}^{N_e} \langle q, \frac{1}{\tau_e} \nabla \cdot \mathbf{u} \rangle_{\Omega_e^{t=t^i}}, \quad d(r, \mathbf{u}, \mu) := \langle r, \mu - \eta(\dot{\gamma}(\nabla^s \mathbf{u})) \rangle$$

$$e(\mathbf{w}, \mathbf{u}) := \sum_{e=1}^{N_e} \langle \gamma_e \nabla \cdot \mathbf{w}, \nabla \cdot \mathbf{u} \rangle_{\Omega_e^{t=t^i}},$$

very much similar to the Eulerian case. The stabilisation parameter for the grad-div term γ_e remains unchanged (2.66), while τ_e incorporates the convective velocity \mathbf{c}_h ,

$$\tau_e = \frac{1}{\rho} \left[\left(\frac{4}{\rho h_e^2} \mu_h \right)^2 + \left(\frac{2}{h_e} \|\mathbf{c}_h\| \right)^2 + \left(\frac{2}{\Delta t} \right)^2 \right]^{-1/2}. \quad (4.20)$$

With this, we finally arrive at the fully discrete problem of finding $(\mathbf{u}_h^{n+1}, p_h^{n+1}, \mu_h^{n+1}) \in [X_h^{n+1}]^d \times X_h^{n+1} \times X_h^{n+1}$ with X_h^{n+1} (4.19), given data (\mathbf{u}_h^n, p_h^n) from the previous time step, node positions \mathbf{x}_h^i , $i = n+1, n$ and $\mathbf{u}_h^{n+1}|_{\Gamma_D^{n+1}} = \mathbf{g}_h^{n+1}$, such that

$$\begin{aligned} \langle \mathbf{w}_h, \frac{\rho}{\Delta t^n} (\mathbf{u}_h^{n+1} - \mathbf{u}_h^n) \rangle + \theta a(\mathbf{w}_h, \mathbf{u}_h^{n+1}, \mathbf{c}_h^*, p_h^{n+1}, \mu_h^{n+1}, t^{n+1}) \\ + e(\mathbf{w}_h, \mathbf{u}_h^{n+1}) + \theta' a(\mathbf{w}_h, \mathbf{u}_h^n, \mathbf{c}_h^n, p_h^n, \mu_h^n, t^n) = 0, \end{aligned} \quad (4.21)$$

$$\begin{aligned} \langle \nabla q, \frac{\rho}{\Delta t^n} (\mathbf{u}_h^{n+1} - \mathbf{u}_h^n) \rangle + \theta b(q_h, \mathbf{u}_h^{n+1}, \mathbf{c}_h^*, p_h^{n+1}, \mu_h^{n+1}, t^{n+1}) \\ + \theta' b(q_h, \mathbf{u}_h^n, \mathbf{c}_h^n, p_h^n, \mu_h^n, t^n) + c(q_h, \mathbf{u}_h^{n+1}) = 0, \end{aligned} \quad (4.22)$$

$$d(r_h, \mathbf{u}_h^*, \mu_h^{n+1}) = 0, \quad (4.23)$$

for all $(\mathbf{w}_h, q_h, r_h) \in [X_h^{n+1}]^d \times X_h^{n+1} \times X_h^{n+1}$, where $\mathbf{w}_h|_{\Gamma_D^{n+1}} = \mathbf{0}$ and all integrals are evaluated in Ω_h^{n+1} and its respective boundary segments. The convective velocity \mathbf{c}^i at t^i is computed as $\mathbf{c}^i = \mathbf{u}^i - \mathbf{u}_m^i$, where care must be taken constructing the mesh velocity in a consistent manner, i.e., following [214]

$$\frac{1}{\Delta t^n} (\mathbf{x}^{n+1} - \mathbf{x}^n) \Big|_{\bar{\mathbf{x}}} = \theta \mathbf{u}_m^{n+1} + \theta' \mathbf{u}_m^n. \quad (4.24)$$

Similar to the Eulerian case, we use a linearised convective velocity \mathbf{c}^* instead of \mathbf{c}^{n+1} and decouple the viscosity projection step from the velocity-pressure system by employing an extrapolation of the velocity \mathbf{u}^* ,

$$\mathbf{c}^{n+1} \approx \mathbf{c}^* := \mathbf{u}^* - \mathbf{u}_m^{n+1} \quad \text{and} \quad \mathbf{u}^* := \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{u}^{n-j} \Big|_{\bar{\mathbf{x}}} \approx \mathbf{u}^{n+1}, \quad (4.25)$$

with coefficients β_j^m given in Tab. 2.1.

With this, all terms of the fully discrete BVS scheme in ALE form (4.21)–(4.23) are given. Thus, we are in position to prove that it preserves a constant solution provided the conditions formulated by Förster et al. [214] hold.

Theorem 4.1.2. *The BVS scheme in convective ALE form (4.21)–(4.23) fulfils the GCL given that all integrals are evaluated considering Ω_h^{n+1} and its respective boundary segments. Constant solutions \mathbf{u}_0 and p_0 with corresponding $\mu = \eta(0)$ are thus preserved at the discrete level.*

Proof. Inserting a constant solution $\mathbf{u}^i = \mathbf{u}_0$ and $p^i = p_0$ for $i = n+1, n$ and assuming zero body forces $\mathbf{b} = \mathbf{0}$, we can show that no spurious terms remain. All terms including spatial derivatives vanish and system (4.21)–(4.23) reduces to

$$\begin{aligned} \langle \mathbf{w}_h, \frac{\rho}{\Delta t^n} (\mathbf{u}_0^{n+1} - \mathbf{u}_0^n) \rangle - \theta \langle \nabla \cdot \mathbf{w}_h, p_0^{n+1} \rangle - \theta \langle \mathbf{w}_h, \bar{\mathbf{t}}^{n+1} \rangle \\ - \theta' \langle \nabla \cdot \mathbf{w}_h, p_0^n \rangle - \theta' \langle \mathbf{w}_h, \bar{\mathbf{t}}^n \rangle = 0, \\ \langle \nabla q, \frac{\rho}{\Delta t^n} (\mathbf{u}_0^{n+1} - \mathbf{u}_0^n) \rangle = 0, \\ \langle r, \mu - \eta(0) \rangle = 0, \end{aligned}$$

where temporal indices are kept to indicate origin of the individual terms. Evaluating all integrals over Ω_h^{n+1} and given consistent boundary conditions on \mathbf{u}_0^{n+1} , the time derivative evaluates to zero. Due to consistent Neumann data, $\bar{\mathbf{t}}^{n+1} = \bar{\mathbf{t}}^n$, a suitable pressure scaling is introduced and hence, no spurious terms remain spoiling the conservation of a constant solution. Additionally, $\mu \equiv \eta(0)$ remains as the constant viscosity. \square

Note here, that using \mathbf{c}_h^{n+1} or \mathbf{c}_h^* does not influence this proof of the GCL since we are using the convective ALE form, but of course the temporal accuracy of \mathbf{c}_h has a natural implication on the overall temporal accuracy of the scheme [214]. Thus, adherence to the GCL does not guarantee higher-order convergence rates, and in fact, is not even necessary in this regard [213, 216].

As demonstrated, the BVS formulation in the Eulerian setting is easily adapted to account for domain motion in the ALE case. A given implementation of a fixed-grid CFD solver thus only requires minor adaptations. The timestepping method as presented in Sec. 2.5 for the Eulerian setting is straight-forwardly extended to a BVS scheme in ALE form. Including an initialisation phase, the algorithm then reads

-
1. *Viscosity initialisation:* Given the initial velocity field \mathbf{u}_h^0 in Ω_h^0 , compute the viscosity μ_h^0 , by solving Eqn. (4.23).
 2. *Domain update:* Compute new nodal positions \mathbf{x}^1 of Ω_h^1 and $\mathbf{u}_{m,h}^1$ via Eqn. (4.24).
 3. *Implicit Euler step:* Solve the velocity-pressure system (4.21)–(4.22) for \mathbf{u}_h^1 and p_h^1 with $\theta = 1$ and $\theta' = 0$ and a first order extrapolation $\mathbf{c}_h^* = \mathbf{u}_h^0 - \mathbf{u}_{m,h}^1$.
 4. *Timestepping:*
FOR $n = 1, \dots, N_t$
 - a) Update node coordinates \mathbf{x}^{n+1} of Ω_h^{n+1} and $\mathbf{u}_{m,h}^{n+1}$ via Eqn. (4.24).
 - b) Extrapolate the (convective) velocity via Eqn. (4.25), giving \mathbf{c}_h^* and \mathbf{u}_h^* .
 - c) Compute μ_h^{n+1} via L^2 -projection (4.23) with \mathbf{u}_h^* .
 - d) Solve the linearised velocity-pressure system (4.21)–(4.22) with time integration parameters θ and θ' as desired to obtain \mathbf{u}_h^{n+1} and p_h^{n+1} .

END FOR

Thus, only slight changes of the Eulerian setting are needed: (i) the update of the N_n node coordinates \mathbf{x}_i^t , $i = 1, \dots, N_n$ of the moving Ω_h^t , and (ii) adopting the convective velocity \mathbf{c}_h^* . The remaining parts of the solver are unchanged, making the extension to the ALE case trivial and allowing reuse of the linear solver and preconditioner as presented in Sec. 2.5.

4.1.2 ALE split-step scheme

An extension of the Eulerian split-step scheme to the ALE setting is based on a modification of Eqns. (2.17)–(2.24), incorporating the ALE time derivative $\partial_t \mathbf{u}|_{\tilde{\mathbf{x}}}$ and convective velocity $\mathbf{c} := \mathbf{u} - \mathbf{u}_m$, thus yielding

$$\rho[\partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} + (\nabla \mathbf{u})\mathbf{c}] - \mu[\Delta \mathbf{u} + \chi \nabla(\nabla \cdot \mathbf{u})] - 2\nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b} \quad \text{in } \Omega^t \times (0, T], \quad (4.26)$$

$$-\Delta p + \nabla \cdot [2\nabla^s \mathbf{u} \nabla \mu - \rho(\nabla \mathbf{u})\mathbf{c}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = -\nabla \cdot \mathbf{b} \quad \text{in } \Omega^t \times [0, T], \quad (4.27)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D^t \times (0, T], \quad (4.28)$$

$$\left(-p\mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \right) \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N^t \times (0, T], \quad (4.29)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0, \quad (4.30)$$

$$\nabla \cdot \mathbf{u}_0 = 0 \quad \text{in } \Omega^0, \quad (4.31)$$

$$-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot \left[\mu \left(\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u} \right) \mathbf{n} - \bar{\mathbf{t}} \right] = p \quad \text{on } \Gamma_N^t \times [0, T], \quad (4.32)$$

$$\mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} - \rho(\nabla \mathbf{u})\mathbf{c} + 2\nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] = \mathbf{n} \cdot \nabla p \quad \text{on } \Gamma_D^t \times [0, T]. \quad (4.33)$$

In comparison to Thm. 2.1.1, subtle differences arise, which are summarised in the following Cor. 4.1.3.

Corollary 4.1.3. *For sufficiently regular $p, \mathbf{u}, \mathbf{u}_m, \mathbf{b}, \mathbf{g}, \bar{\mathbf{t}}$, systems (4.7)–(4.11) and (4.26)–(4.33) are equivalent.*

Proof. We can show that system (4.7)–(4.11) implies (4.26)–(4.33) in complete analogy to Thm. 2.1.1 by simply replacing the convective velocity \mathbf{u} by $\mathbf{u} - \mathbf{u}_m$ and the time derivative $\partial_t \mathbf{u}$ by $\partial_t \mathbf{u}|_{\bar{\mathbf{x}}}$ and following the exact same sequence of steps.

Proving the other way round, we aim to show that (4.26)–(4.33) imply (4.7)–(4.11), starting off by taking the divergence of the rewritten momentum balance equation (4.7) and adding the result to the PPE (4.27), yielding

$$\rho \nabla \cdot (\partial_t \mathbf{u}|_{\bar{\mathbf{x}}}) - \nabla \cdot (\mu \Delta \mathbf{u}) - \chi \nabla \cdot [\mu \nabla (\nabla \cdot \mathbf{u})] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = 0,$$

which we can rewrite using

$$\begin{aligned} \nabla \cdot (\mu \Delta \mathbf{u}) &= \nabla \mu \cdot \Delta \mathbf{u} + \mu \nabla \cdot (\Delta \mathbf{u}) \\ &= \nabla \mu \cdot \Delta \mathbf{u} + \mu \Delta (\nabla \cdot \mathbf{u}), \end{aligned}$$

and

$$\nabla \cdot [\mu \nabla (\nabla \cdot \mathbf{u})] = \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u})] + \mu \nabla \cdot [\nabla (\nabla \cdot \mathbf{u})],$$

such that we end up with

$$\rho \partial_t (\nabla \cdot \mathbf{u})|_{\bar{\mathbf{x}}} - (1 + \chi) \mu \Delta (\nabla \cdot \mathbf{u}) - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu - \nabla \mu \cdot [\chi \nabla (\nabla \cdot \mathbf{u})] = 0.$$

Then, with Eqn. (2.28), this is equivalent to

$$\rho \partial_t (\nabla \cdot \mathbf{u})|_{\bar{\mathbf{x}}} - (1 + \chi) \mu \Delta (\nabla \cdot \mathbf{u}) - (1 + \chi) \nabla \mu \cdot [\nabla (\nabla \cdot \mathbf{u})] = 0,$$

which is again a heat equation in the auxiliary variable $\Phi := \nabla \cdot \mathbf{u}$ in ALE form,

$$\rho \partial_t \Phi|_{\bar{\mathbf{x}}} - (1 + \chi) \nabla \cdot (\mu \nabla \Phi) = 0, \quad (4.34)$$

with variable diffusion coefficient $(1 + \chi) \mu / \rho$. Due to Eqn. (4.31), i.e., $\nabla \cdot \mathbf{u}_0 = 0$, we have $\Phi = 0$ as initial condition. Boundary conditions for Eqn. (4.34) on Γ_D^t are derived subtracting the modified momentum balance (4.26) dotted with the unit outward normal \mathbf{n} from minus the pressure Neumann condition (4.33), yielding

$$\mathbf{n} \cdot [\mu \Delta \mathbf{u} + \chi \mu \nabla (\nabla \cdot \mathbf{u}) + \mu \nabla \times (\nabla \times \mathbf{u})] = 0,$$

which, again using Eqn. (2.28), reduces to

$$\mu(1 + \chi) \mathbf{n} \cdot [\nabla (\nabla \cdot \mathbf{u})] = \mu(1 + \chi) \mathbf{n} \cdot \nabla \Phi = 0 \quad \implies \quad \mathbf{n} \cdot \nabla \Phi = 0 \quad \text{on } \Gamma_D^t.$$

On Γ_N^t , we dot the traction condition (4.29) with \mathbf{n} and add the Dirichlet condition on the pressure (4.32) to directly obtain

$$\mu \nabla \cdot \mathbf{u} = 0 \quad \implies \quad \Phi = 0 \quad \text{on } \Gamma_N^t.$$

Summing up, we have a heat equation in $\Phi := \nabla \cdot \mathbf{u}$ with zero initial-, Dirichlet- and Neumann conditions. Therefore, $\Phi \equiv 0$ follows as the only admissible solution and thus, the continuity condition is inherently fulfilled. All that remains is to show that the modified momentum equation (4.26) implies the original one (4.7), which trivially results from using Eqn. (2.25) again. The alternative system thus fulfils the initial and boundary conditions and enforces mass and momentum balance, concluding the proof for both the stress-divergence and Laplacian form of the viscous stress tensor. \square

Opposed to a standard ALE formulation of the Navier–Stokes equations, where the continuity equation remains unaltered from the Eulerian setting, now in the PPE-based system, the GCL might also influence mass conservation via the PPE. To investigate the conservation properties of the alternative system, we thus introduce a suitable fully discrete formulation. Similar to the derivations in Sec. 3, we aim to linearise and decouple the equations for velocity, viscosity and pressure. Including again divergence damping via the auxiliary variable ψ acting only on the past time steps' velocity, we begin with approximating the ALE time derivative via a BDF,

$$\partial_t \mathbf{u}|_{\tilde{\mathbf{x}}} \approx \alpha_0^m \mathbf{u}^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m \left(\mathbf{u}^{n-j} - \nabla \psi^{n-j} \right) \Big|_{\tilde{\mathbf{x}}},$$

and higher-order accurate extrapolations

$$\begin{aligned} \mathbf{u}^{n+1} \approx \mathbf{u}^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{u}^{n-j} \Big|_{\tilde{\mathbf{x}}}, \quad p^{n+1} \approx p^* = \sum_{j=0}^{m-1} \beta_{j+1}^m p^{n-j} \Big|_{\tilde{\mathbf{x}}}, \\ \mu^{n+1} \approx \mu^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mu^{n-j} \Big|_{\tilde{\mathbf{x}}}, \end{aligned} \quad (4.35)$$

with coefficients defined in Tab. 2.1 and specifically noting the fixed $\tilde{\mathbf{x}}$ for clarity. The convective velocity is replaced by $\mathbf{c} := \mathbf{u} - \mathbf{u}_m$, where the mesh velocity $\mathbf{u}_m := \partial_t \mathbf{x}|_{\tilde{\mathbf{x}}}$ is computed via

$$\mathbf{u}_m^{n+1} \approx \sum_{j=0}^m \alpha_j^m \mathbf{x}^{n+1-j} \Big|_{\tilde{\mathbf{x}}}. \quad (4.36)$$

Spatial discretisation is carried out based on standard continuous isoparametric Lagrangian finite elements, using $X_h^t \subset H^1(\Omega_h^t)$ as defined in Eqn. (4.19), while all the weak forms can be extended from the Eulerian setting in a straight-forward manner. Denoting by $\langle \cdot, \cdot \rangle$ the standard L^2 scalar product over Ω^{n+1} (and its respective boundary

segments when indicated by a subscript), the velocity $\mathbf{u}_h^{n+1} \in [X_h^{n+1}]^d$ is sought, such that $\mathbf{u}_h^{n+1}|_{\Gamma_D^{n+1}} = \mathbf{g}^{n+1}$ and

$$\begin{aligned} & \langle \rho \mathbf{w}_h, \alpha_0^m \mathbf{u}_h^{n+1} + (\nabla \mathbf{u}_h^{n+1}) \mathbf{c}_h^* \rangle - (1 - \chi) \langle \mathbf{w}_h, (\nabla^\top \mathbf{u}_h^{n+1}) \nabla \mu_h^* \rangle \\ & + (1 - \chi/2) \langle \nabla \mathbf{w}_h + \chi \nabla^\top \mathbf{w}_h, \mu_h^* \nabla \mathbf{u}_h^{n+1} + \chi \mu_h^* \nabla^\top \mathbf{u}_h^{n+1} \rangle \\ & = \langle \mathbf{w}_h, \mathbf{b}^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \mathbf{u}_h^{n-j} \rangle + \langle \nabla \cdot \mathbf{w}_h, \hat{p}_h^* \rangle + \langle \mathbf{w}_h, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N^{n+1}}, \end{aligned} \quad (4.37)$$

for all $\mathbf{w}_h \in [X_h^{n+1}]^d$ with $\mathbf{w}_h|_{\Gamma_D^{n+1}} = \mathbf{0}$. Employing the generalised Laplacian form with semi-implicit viscous terms, we can alternatively search for the velocity vector components individually, seeking for $w_{i,h} \in X_h^{n+1}$, $i = 1, \dots, d$, such that $w_{i,h}^{n+1}|_{\Gamma_D^{n+1}} = g_i^{n+1}$ and

$$\begin{aligned} & \langle w_{i,h}, \rho \alpha_0^m u_{i,h}^{n+1} + \rho \mathbf{c}_h^* \cdot \nabla u_{i,h}^{n+1} \rangle + \langle \nabla w_{i,h}, \mu_h^* \nabla u_{i,h}^{n+1} \rangle \\ & = \langle w_{i,h}, b_i^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m u_{i,h}^{n-j} + \nabla \mu_h^* \cdot \partial_{x_i} \mathbf{u}_h^* \rangle + \langle \partial_{x_i} w_{i,h}, \hat{p}_h^* \rangle + \langle w_{i,h}, \bar{t}_i^{n+1} \rangle_{\Gamma_N^{n+1}}, \end{aligned} \quad (4.38)$$

for all $w_{i,h} \in X_h^{n+1}$, where $w_{i,h}|_{\Gamma_D^{n+1}} = 0$.

To spare us the additional Poisson problem to compute ψ_h and the vector-valued L^2 -projection of \mathbf{u}_h to recover a weakly divergence-free $\hat{\mathbf{u}}_h^{n+1}$ with $\hat{\mathbf{u}}_h^{n+1}|_{\Gamma_D^{n+1}} \neq \mathbf{g}^{n+1}$, we again gather the fluid's pressure and the Leray projection variable in a joint pressure-like quantity,

$$\hat{p}_h^* := p_h^* + \sum_{j=0}^{m-1} \alpha_{j+1}^m \psi_h^{n-j},$$

applying divergence damping. That is, we consider Leray projections of the past time steps' velocity fields only (cf. Sec. 3.3). Hence, we seek for the pressure-like variable $\hat{p}_h^* \in X_h^{n+1}$ with $\hat{p}_h^*|_{\Gamma_N^{n+1}} = \zeta_h^*$, such that

$$\begin{aligned} \langle \nabla q_h, \nabla \hat{p}_h^* \rangle & = - \langle q_h, \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \nabla \cdot \mathbf{u}_h^{n-j} \rangle + \sum_{j=0}^{m-1} \beta_{j+1}^m l(q_h, \mathbf{u}_h^{n-j}, \mathbf{c}_h^{n-j}, \mu_h^{n-j}, t^{n-j}) \\ & - \langle q_h \mathbf{n}, \rho \sum_{j=0}^m \alpha_j^m \mathbf{g}^{n+1-j} \rangle_{\Gamma_D^{n+1}} \end{aligned} \quad (4.39)$$

for all $q_h \in X_h^{n+1}$, with $q_h|_{\Gamma_N^{n+1}} = 0$ and the form l evaluated at $t = t^i$ given by

$$l(q, \mathbf{u}, \mathbf{c}, \mu, t^i) := \langle \nabla q, \mathbf{b}^i - \rho (\nabla \mathbf{u}) \mathbf{c} + 2 (\nabla^\top \mathbf{u}) \nabla \mu \rangle + \langle \mathbf{n} \times \nabla q, \mu \nabla \times \mathbf{u} \rangle_{\Gamma_D^{n+1}}.$$

Suitable boundary data for the pressure-like quantity \hat{p}_h^* is constructed through a boundary projection, seeking $\zeta_h^* \in Y_h^{n+1} \subset H^{1/2}(\Gamma_N^{n+1})$, such that

$$\begin{aligned} \langle s_h, \zeta_h^* \rangle_{\Gamma_N^{n+1}} = & - \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h, \mu_h^{n+1} \nabla \cdot \mathbf{u}_h^{n-j} \rangle_{\Gamma_N^{n+1}} \\ & + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h \mathbf{n}, \mu_h^{n-j} \left(\nabla \mathbf{u}_h^{n-j} + \chi \nabla^\top \mathbf{u}_h^{n-j} \right) \mathbf{n} - \bar{\mathbf{t}}^{n-j} \rangle_{\Gamma_N^{n+1}}, \end{aligned} \quad (4.40)$$

holds for all $s_h \in Y_h^{n+1} \subset H^{1/2}(\Gamma_N^{n+1})$. Given the velocity \mathbf{u}_h^{n+1} , one may then obtain the viscosity $\mu_h^{n+1} \in X_h^{n+1}$, such that

$$\langle r_h, \mu_h^{n+1} \rangle = \langle r_h, \eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^{n+1})) \rangle \quad \forall r_h \in X_h^{n+1}. \quad (4.41)$$

As in the Eulerian case with divergence damping included, we first solve for the boundary data ζ_h^* of the pressure-like variable \hat{p}_h^* , followed by \hat{p}_h^* itself, then compute the fluid velocity (or its individual components) and lastly the viscosity at each time step. Adding merely the domain and mesh velocity update as additional steps compared to the Eulerian setting, the ALE split-step scheme finally reads

-
1. *Viscosity initialisation:* Compute the viscosity μ_h^0 by solving Eqn. (4.41) with \mathbf{u}_0 in the initial configuration Ω_h^0 .
 2. *Initial lower-order steps:* Until enough time step data is gathered, solve the split-step scheme with lower order $\hat{m} = 1, \dots, m-1$.
 3. *Timestepping:*
FOR $n = m-1, \dots, N_t$
 - a) Update node coordinates \mathbf{x}^{n+1} of Ω_h^{n+1} and $\mathbf{u}_{m,h}^{n+1}$ via Eqn. (4.36).
 - b) Compute continuous pressure boundary data ζ_h^* via (4.40).
 - c) Solve the modified Poisson equation (4.39) for \hat{p}_h^* .
 - d) Extrapolate in time via (4.35), yielding \mathbf{u}_h^* , μ_h^* and $\mathbf{c}_h^* := \mathbf{u}_h^* - \mathbf{u}_{m,h}^{n+1}$.
 - e) Solve the momentum balance equation for \mathbf{u}_h^{n+1} directly via (4.37), or for individual components through (4.38).
 - f) Recover the viscosity μ_h^{n+1} via (4.41), given \mathbf{u}_h^{n+1} .
 - g) Update the time step size Δt^n and corresponding coefficients α_j^m and β_j^m .

END FOR

Next, in analogy to the ALE extension of the BVS formulation, the following Thm. 4.1.4 states GCL-fulfillment for the split-step scheme in ALE form.

Theorem 4.1.4. *The split-step scheme in ALE form (4.37)–(4.41) fulfils the GCL given that all integrals are evaluated considering Ω_h^{n+1} and its respective boundary segments. Constant solutions \mathbf{u}_0 and p_0 with corresponding $\mu = \eta(0)$ are thus preserved at the discrete level.*

Proof. Inserting a constant solution $\mathbf{u}^i = \mathbf{u}_0$ and $p^i = p_0$ for $i = n + 1, n, n - 1, \dots, 0$ and assuming zero body forces $\mathbf{b} = \mathbf{0}$, we can show that no spurious terms remain. Following the sequence of steps when including divergence-damping, we first project suitable boundary data via Eqn. (4.40), which reduces to

$$\langle s_h, \zeta_h^* \rangle_{\Gamma_N^{n+1}} = \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h \mathbf{n}, -\bar{\mathbf{t}}^{n-j} \rangle_{\Gamma_N^{n+1}},$$

since all spatial derivatives of $\mathbf{u} = \mathbf{u}_0$ are zero. With consistent boundary data given by $\bar{\mathbf{t}} := \left(-p\mathbf{I} + m\nabla\mathbf{u} + \chi\nabla^\top\mathbf{u} \right) \mathbf{n}$, this induces a suitable pressure scaling, as

$$\langle s_h, \zeta_h^* \rangle_{\Gamma_N^{n+1}} = \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h \mathbf{n}, (p^{n-j}\mathbf{I}) \mathbf{n} \rangle_{\Gamma_N^{n+1}} = \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h, p^{n-j} \rangle_{\Gamma_N^{n+1}}.$$

Given that $\sum_{j=0}^{m-1} \beta_{j+1}^m = 1$, a constant solution $\zeta_h^* = p_0$ is recovered on Γ_N^{n+1} . The Poisson equation for the pressure-like quantity \hat{p}_h^* (4.39) boils down to

$$\langle \nabla q_h, \nabla \hat{p}_h^* \rangle = -\langle q\mathbf{n}, \rho \sum_{j=0}^m \alpha_j^m \mathbf{g}^{n+1-j} \rangle_{\Gamma_D^{n+1}},$$

where the right-hand side equals 0 due to $\sum_{j=0}^m \alpha_j^m = 0$. Then, $\hat{p}_h^*|_{\Gamma_N^{n+1}} = \zeta_h^* = p_0$ leads to $\hat{p}_h^* = p_0$. Plugging p_0 into the momentum balance step in its vector-valued (4.37) or component-wise form (4.38), we obtain

$$\langle \rho \mathbf{w}_h, \alpha_0^m \mathbf{u}_h^{n+1} \rangle = -\langle \rho \mathbf{w}_h, \sum_{j=1}^m \alpha_j^m \mathbf{u}_h^{n+1-j} \rangle + \langle \nabla \cdot \mathbf{w}_h, p_0 \rangle + \langle \mathbf{w}_h, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N^{n+1}},$$

where a consistent $\bar{\mathbf{t}}$ yields $-\langle \mathbf{w}_h, p_0 \mathbf{n} \rangle_{\Gamma_N^{n+1}}$, cancelling the second to last term since

$$\langle \nabla \cdot \mathbf{w}_h, p_0 \rangle = \langle \mathbf{w}_h, p_0 \mathbf{n} \rangle_{\Gamma_N^{n+1}} - \langle \mathbf{w}_h, \nabla p_0 \rangle = \langle \mathbf{w}_h, p_0 \mathbf{n} \rangle_{\Gamma_N^{n+1}},$$

and $\alpha_0^m = \sum_{j=1}^m \alpha_j^m$ leads to $\mathbf{u}_h^{n+1} = \mathbf{u}_0$. The corresponding viscosity is recovered from Eqn. (4.41), which for $\mathbf{u}_h^{n+1} = \mathbf{u}_0$ simply gives

$$\langle r_h, \mu_h^{n+1} \rangle = \langle r_h, \eta(0) \rangle,$$

and a constant viscosity, $\mu \equiv \eta(0)$, in the entirety of the domain. Thereby, we complete the proof having shown that inserting a constant solution into the ALE split-step scheme, no spurious terms spoil their conservation at the discrete level. \square

In summary, both the BVS formulation and the PPE-based split-step scheme are easily extended to the ALE setting. We fulfil the GCL following the conditions formulated by Förster et al. [214], that is, we (i) integrate all terms over the configuration at time t^{n+1} and its respective boundary segments, (ii) employ the convective form of the

Navier–Stokes equations, and (iii) compute the mesh velocity using a suitable higher-order timestepping scheme. Additionally, linearising the convective velocity \mathbf{c} does not have an influence on fulfilling the GCL, but is considered in a higher-order accurate fashion to not spoil the temporal accuracy of the overall scheme.

The resulting solvers and their implementations are thus in large parts identical to the ones presented in Secs. 2.4 and 3.4. Merely the linearised convective velocity, defined as $\mathbf{c}^{n+1} := \mathbf{u}^* - \mathbf{u}_m^{n+1}$, and the nodal coordinates of the finite element mesh have to be updated in each time step. The convective velocity is already considered as an extrapolation of the fluid velocity (or its last iterate) in the Eulerian setting, thus rendering this adaptation effortless. Further discussion and corresponding numerical investigations are postponed to the FSI setting introduced in Secs. 7.1 and practical application is shown in Secs. 8.5–8.8, where the ALE split-step scheme is considered.

4.2 Robin boundary conditions in the fluid problem

In the context of fluid–structure interaction, introducing Robin interface conditions is found to improve convergence behaviour of the FSI coupling scheme, and hence has lately received great attention in the FSI community (see, e.g., [217–224], or [5, 6, 225] for projection/split-step fluid solvers). In FSI, the Robin interface condition reads

$$\eta^R \mathbf{u}_f + \boldsymbol{\sigma} \mathbf{n} = \mathbf{h} := \eta^R \mathbf{d}_s + \boldsymbol{\sigma}_s \mathbf{n} \quad \text{on } \Sigma^t, \quad (4.42)$$

where the vector \mathbf{h} typically depends on the structure’s displacement \mathbf{d}_s and Cauchy stress $\boldsymbol{\sigma}_s$ on the interface Σ^t . The parameter η^R is used to scale from a Dirichlet condition, enforcing continuity of velocities as $\eta^R \rightarrow \infty$, to a Neumann condition, which equates tractions as $\eta^R \rightarrow 0$. Since a large body of research suggests those conditions to improve performance of FSI coupling schemes, the effects of introducing a corresponding Robin condition to the fluid problem is worth investigating. Thus, as a starting point for later investigations in the FSI context, we introduce a Robin condition to the fluid problem in the form of

$$\eta^R \mathbf{u} + \boldsymbol{\sigma} \mathbf{n} = \mathbf{h} \quad \text{on } \Gamma_R^t, \quad (4.43)$$

dropping the subscripts indicating fluid and solid domain, as we only consider a fluid phase in Ω^t for now. The domain boundary is decomposed into non-overlapping Robin-, Dirichlet- and Neumann segments, i.e., $\partial\Omega^t = \Gamma_D^t \cup \Gamma_N^t \cup \Gamma_R^t$ and $\Gamma_R^t \cap \Gamma_D^t = \emptyset$, $\Gamma_R^t \cap \Gamma_N^t = \emptyset$, $\Gamma_D^t \cap \Gamma_N^t = \emptyset$, where the respective conditions are enforced.

Starting from the standard convective ALE formulation of the Navier–Stokes equations as given in Eqns. (4.7)–(4.11), we add the Robin boundary condition on the time-dependent Γ_R^t , resulting in an alternative system,

$$\rho \partial_t \mathbf{u}|_{\bar{\mathbf{x}}} + \rho (\nabla \mathbf{u}) \mathbf{c} - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2\nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b} \quad \text{in } \Omega^t \times (0, T], \quad (4.44)$$

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in } \Omega^t \times (0, T], \quad (4.45)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D^t \times (0, T], \quad (4.46)$$

$$\left(-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \right) \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N^t \times (0, T], \quad (4.47)$$

$$\eta^R \mathbf{u} + \boldsymbol{\sigma} \mathbf{n} = \mathbf{h} \quad \text{on } \Gamma_R^t \times (0, T], \quad (4.48)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0 \quad (4.49)$$

$$\nabla \cdot \mathbf{u}_0 = 0 \quad \text{in } \Omega^0. \quad (4.50)$$

Given that the Robin condition is formulated in terms of Cauchy tractions $\mathbf{t} = \boldsymbol{\sigma} \mathbf{n}$, we can modify the Robin condition (4.48) to give

$$\mathbf{t} := \boldsymbol{\sigma} \mathbf{n} = \mathbf{h} - \eta^R \mathbf{u}.$$

We can insert this term into the pressure Dirichlet condition on Γ_N^t as used for the stress-divergence form ($\chi = 1$), which reads

$$-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot [2\mu (\nabla^s \mathbf{u}) \mathbf{n} - \mathbf{t}] = p,$$

as the Robin boundary condition *always* incorporates *real* tractions, yielding

$$-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot [2\mu (\nabla^s \mathbf{u}) \mathbf{n} - \mathbf{h} + \eta^R \mathbf{u}] = p$$

on the Robin boundary Γ_R^t . With this, one may formulate an alternative system to the standard ALE form of the Navier–Stokes equations, involving the familiar parts (see Ch. 3 or Sec. 4.1.2), that is

$$\rho [\partial_t \mathbf{u}|_{\bar{\mathbf{x}}} + (\nabla \mathbf{u}) \mathbf{c}] - \mu [\Delta \mathbf{u} + \chi \nabla (\nabla \cdot \mathbf{u})] - 2\nabla^s \mathbf{u} \nabla \mu + \nabla p = \mathbf{b} \quad \text{in } \Omega^t \times (0, T], \quad (4.51)$$

$$-\Delta p + \nabla \cdot [2\nabla^s \mathbf{u} \nabla \mu - \rho (\nabla \mathbf{u}) \mathbf{c}] - [\nabla \times (\nabla \times \mathbf{u})] \cdot \nabla \mu = -\nabla \cdot \mathbf{b} \quad \text{in } \Omega^t \times [0, T], \quad (4.52)$$

$$\mathbf{u} = \mathbf{g} \quad \text{on } \Gamma_D^t \times (0, T], \quad (4.53)$$

$$\left(-p \mathbf{I} + \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \right) \mathbf{n} = \bar{\mathbf{t}} \quad \text{on } \Gamma_N^t \times (0, T], \quad (4.54)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0, \quad (4.55)$$

$$\nabla \cdot \mathbf{u}_0 = 0 \quad \text{in } \Omega^0, \quad (4.56)$$

$$-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot [\mu (\nabla \mathbf{u} + \chi \nabla^\top \mathbf{u}) \mathbf{n} - \bar{\mathbf{t}}] = p \quad \text{on } \Gamma_N^t \times [0, T], \quad (4.57)$$

$$-\mu \nabla \cdot \mathbf{u} + \mathbf{n} \cdot [2\mu (\nabla^s \mathbf{u}) \mathbf{n} - \mathbf{h} + \eta^R \mathbf{u}] = p \quad \text{on } \Gamma_R^t \times [0, T], \quad (4.58)$$

$$\mathbf{n} \cdot [\mathbf{b} - \rho \partial_t \mathbf{u}|_{\bar{\mathbf{x}}} - \rho (\nabla \mathbf{u}) \mathbf{c} + 2\nabla^s \mathbf{u} \nabla \mu - \mu \nabla \times (\nabla \times \mathbf{u})] = \mathbf{n} \cdot \nabla p \quad \text{on } \Gamma_D^t \times [0, T]. \quad (4.59)$$

This system is identical to the ALE setup (4.26)–(4.33) up to the Robin boundary condition, allowing for an equivalence proof along the lines of Cor. 4.1.3.

Corollary 4.2.1. *For sufficiently regular $p, \mathbf{u}, \mathbf{u}_m, \mathbf{b}, \mathbf{g}, \bar{\mathbf{t}}, \mathbf{h}$, systems (4.44)–(4.50) and (4.51)–(4.59) are equivalent.*

Proof. We can show that system (4.44)–(4.50) implies (4.51)–(4.59) in complete analogy to Cor. 4.1.3 and following the exact same sequence of steps. The pressure Dirichlet condition on Γ_R^t (as the only differing part to the previously presented proof) is constructed by dotting the Robin boundary condition Eqn. (4.48) with the normal vector \mathbf{n} and adding a multiple of the continuity equation $-\mu\nabla \cdot \mathbf{u} = 0$, which yields

$$-\mu\nabla \cdot \mathbf{u} + \mathbf{n} \cdot [\eta^R \mathbf{u} + \boldsymbol{\sigma} \mathbf{n} - \mathbf{h}] = 0,$$

where we can factor out the pressure from the Cauchy stress tensor $\boldsymbol{\sigma}$ to recover a Dirichlet condition for p on Γ_R^t , that is

$$p = -\mu\nabla \cdot \mathbf{u} + \mathbf{n} \cdot [2\mu(\nabla^s \mathbf{u}) \mathbf{n} - \mathbf{h} + \eta^R \mathbf{u}].$$

To show that (4.51)–(4.59) imply (4.44)–(4.50), we follow the same steps presented in Cor. 4.1.3 to once more obtain a heat equation in ALE form in the auxiliary variable $\Phi := \nabla \cdot \mathbf{u}$, reading

$$\rho \partial_t \Phi|_{\bar{x}} - (1 + \chi) \nabla \cdot (\mu \nabla \Phi) = 0, \quad (4.60)$$

Since $\nabla \cdot \mathbf{u}_0 = 0$ (4.56), $\Phi = 0$ holds at $t = 0$. Boundary conditions for Eqn. (4.60) on Γ_D^t and Γ_N^t are obtained just as before (see Cor. 4.1.3), i.e.,

$$\mu(1 + \chi) \mathbf{n} \cdot [\nabla(\nabla \cdot \mathbf{u})] = \mu(1 + \chi) \mathbf{n} \cdot \nabla \Phi = 0 \quad \implies \quad \mathbf{n} \cdot \nabla \Phi = 0 \quad \text{on } \Gamma_D^t,$$

as well as

$$\mu \nabla \cdot \mathbf{u} = 0 \quad \implies \quad \Phi = 0 \quad \text{on } \Gamma_N^t.$$

Now, we derive a boundary condition for the heat equation (4.60) on Γ_R^t in a similar fashion as on Γ_N^t . We dot the Robin condition (4.48) with \mathbf{n} and subtract it from the pressure Dirichlet condition (4.58) on Γ_R^t , which directly gives

$$\mu \nabla \cdot \mathbf{u} = 0 \quad \implies \quad \Phi = 0 \quad \text{on } \Gamma_R^t.$$

Altogether, Φ is subject to zero Dirichlet conditions on $\Gamma_R^t \cup \Gamma_N^t$, zero Neumann conditions on Γ_D^t and zero initial conditions. Consequently, $\Phi := \nabla \cdot \mathbf{u} \equiv 0$ is the only admissible solution to the heat equation (4.60) and incompressibility is implicitly enforced. The alternative system thus fulfils the initial and boundary conditions and enforces mass and momentum balance, concluding the proof for both the stress-divergence and Laplacian form of the viscous stress tensor. \square

The main difference to the previously considered problems lies in the momentum balance equation, where the Robin boundary condition is typically introduced rewriting the boundary term arising from integration by parts,

$$\begin{aligned}
\bar{\mathbf{t}} &= \mathbf{t} - (1 - \chi)\mu (\nabla^\top \mathbf{u}) \mathbf{n} \\
&= \boldsymbol{\sigma} \mathbf{n} - (1 - \chi)\mu (\nabla^\top \mathbf{u}) \mathbf{n} \\
&= \mathbf{h} - \eta^R \mathbf{u} - (1 - \chi)\mu (\nabla^\top \mathbf{u}) \mathbf{n}
\end{aligned}$$

such that

$$\langle \mathbf{w}, \eta^R \mathbf{u} + (1 - \chi)\mu (\nabla^\top \mathbf{u}) \mathbf{n} \rangle_{\Gamma_R^t}$$

is added to the left-hand side (if not linearised) and $\langle \mathbf{w}, \mathbf{h} \rangle_{\Gamma_R^t}$ enters the right-hand side of the weak form of the momentum balance equation. Note here, that this term or parts of it might enter the velocity-velocity block and the pseudo traction in the boundary term over Γ_R^t has to be corrected if the generalised Laplacian form is employed. As discussed in Sec. 2.1, choosing $\chi = 0$ results in terms involving the pseudo tractions after integration by parts, which do not coincide with the Cauchy traction vector being a component of the Robin boundary condition.

A second adaptation is necessary in the PPE-based split-step scheme, as the pressure Dirichlet condition on Γ_R^t has to be projected onto a continuous space, similar to a Neumann boundary. Thus, the variable ζ is a continuous projection of the Neumann and Robin boundary data on $\Gamma_\zeta^t := \Gamma_N^t \cup \Gamma_R^t$, which leads to ζ being introduced on both of these boundary segments. From an algorithmic point of view, Robin and Neumann boundaries are treated similarly, where only minor differences in the right-hand side of the boundary data projection step arise. A single step of the ALE split-step scheme including Robin boundary conditions consequently reads

Timestepping: Given enough data at times t^n, \dots, t^{n+1-m} for BDF m time integration, execute the following steps to obtain the solution at t^{n+1} :

1. Update node coordinates \mathbf{x}^{n+1} of Ω_h^{n+1} and \mathbf{u}_m^{n+1} via Eqn. (4.36).
2. Compute continuous pressure boundary data on the Robin and Neumann boundaries, seeking $\zeta_h^* \in Y_h^{n+1} \subset H^{1/2}(\Gamma_\zeta^{n+1})$, such that

$$\begin{aligned}
\langle s_h, \zeta_h^* \rangle_{\Gamma_\zeta^{n+1}} &= - \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h, \mu_h^{n+1} \nabla \cdot \mathbf{u}_h^{n-j} \rangle_{\Gamma_N^{n+1} \cup \Gamma_R^{n+1}} \\
&\quad + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h \mathbf{n}, \mu_h^{n-j} (\nabla \mathbf{u}_h^{n-j} + \chi \nabla^\top \mathbf{u}_h^{n-j}) \mathbf{n} - \bar{\mathbf{t}}^{n-j} \rangle_{\Gamma_N^{n+1}} \\
&\quad + \sum_{j=0}^{m-1} \beta_{j+1}^m \langle s_h \mathbf{n}, 2\mu_h^{n-j} (\nabla^s \mathbf{u}_h^{n-j}) \mathbf{n} - \mathbf{h}^{n-j} + \eta^R \mathbf{u}_h^{n-j} \rangle_{\Gamma_R^{n+1}},
\end{aligned}$$

holds for all $s_h \in Y_h^{n+1} \subset H^{1/2}(\Gamma_\zeta^{n+1})$.

3. Solve the modified Poisson equation for the pressure-like variable $\hat{p}_h^* \in X_h^{n+1}$ with $\hat{p}_h^*|_{\Gamma_\zeta^{n+1}} = \zeta_h^*$, such that

$$\begin{aligned} \langle \nabla q_h, \nabla \hat{p}_h^* \rangle &= - \langle q_h, \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \nabla \cdot \mathbf{u}_h^{n-j} \rangle + \sum_{j=0}^{m-1} \beta_{j+1}^m l(q_h, \mathbf{u}_h^{n-j}, \mathbf{c}_h^{n-j}, \mu_h^{n-j}, t^{n-j}) \\ &\quad - \langle q_h, \rho \sum_{j=0}^m \alpha_j^m \mathbf{g}^{n+1-j} \rangle_{\Gamma_D^{n+1}} \end{aligned}$$

for all $q_h \in X_h^{n+1}$, with $q_h|_{\Gamma_C^{n+1}} = 0$ and the form l evaluated at $t = t^i$ given by

$$l(q, \mathbf{u}, \mathbf{c}, \mu, t^i) := \langle \nabla q, \mathbf{b}^i - \rho(\nabla \mathbf{u})\mathbf{c} + 2(\nabla^\top \mathbf{u}) \nabla \mu \rangle + \langle \mathbf{n} \times \nabla q, \mu \nabla \times \mathbf{u} \rangle_{\Gamma_D^{n+1}}.$$

4. Extrapolate in time via (4.35), yielding \mathbf{u}_h^* , μ_h^* and $\mathbf{c}_h^* := \mathbf{u}_h^* - \mathbf{u}_{m,h}^{n+1}$.
5. Solve the momentum balance equation for the velocity $\mathbf{u}_h^{n+1} \in [X_h^{n+1}]^d$, such that $\mathbf{u}_h^{n+1}|_{\Gamma_D^{n+1}} = \mathbf{g}^{n+1}$ and

$$\begin{aligned} &\langle \rho \mathbf{w}_h, \alpha_0^m \mathbf{u}_h^{n+1} + (\nabla \mathbf{u}_h^{n+1}) \mathbf{c}_h^* \rangle + \langle \mathbf{w}_h, \eta^R \mathbf{u}_h^{n+1} + (1 - \chi) \mu_h^* (\nabla^\top \mathbf{u}_h^{n+1}) \mathbf{n} \rangle_{\Gamma_R^{n+1}} \\ &+ (1 - \chi/2) \langle \nabla \mathbf{w}_h + \chi \nabla^\top \mathbf{w}_h, \mu_h^* \nabla \mathbf{u}_h^{n+1} + \chi \mu_h^* \nabla^\top \mathbf{u}_h^{n+1} \rangle - (1 - \chi) \langle \mathbf{w}_h, (\nabla^\top \mathbf{u}_h^{n+1}) \nabla \mu_h^* \rangle \\ &= \langle \mathbf{w}_h, \mathbf{b}^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \mathbf{u}_h^{n-j} \rangle + \langle \nabla \cdot \mathbf{w}_h, \hat{p}_h^* \rangle + \langle \mathbf{w}_h, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N^{n+1}} + \langle \mathbf{w}_h, \mathbf{h}^{n+1} \rangle_{\Gamma_R^{n+1}}, \end{aligned}$$

for all $\mathbf{w}_h \in [X_h^{n+1}]^d$ with $\mathbf{w}_h|_{\Gamma_D^{n+1}} = \mathbf{0}$, solving a vector-valued problem. Alternatively, apply linearisations of all component-coupling terms in the generalised Laplacian approach ($\chi = 0$) to compute individual $u_{i,h} \in X_h^{n+1}$, $i = 1, \dots, d$, such that $u_{i,h}^{n+1}|_{\Gamma_D^{n+1}} = g_i^{n+1}$ and

$$\begin{aligned} &\langle w_{i,h}, \rho \alpha_0^m u_{i,h}^{n+1} + \rho \mathbf{c}_h^* \cdot \nabla u_{i,h}^{n+1} \rangle + \langle \nabla w_{i,h}, \mu_h^* \nabla u_{i,h}^{n+1} \rangle + \langle w_{i,h}, \eta^R u_{i,h}^{n+1} \rangle_{\Gamma_R^{n+1}} \\ &= \langle w_{i,h}, b_i^{n+1} - \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m u_{i,h}^{n-j} + \nabla \mu_h^* \cdot \partial_{x_i} \mathbf{u}_h^* \rangle + \langle \partial_{x_i} w_{i,h}, \hat{p}_h^* \rangle \\ &\quad + \langle w_{i,h}, \bar{t}_i^{n+1} \rangle_{\Gamma_N^{n+1}} + \langle w_{i,h}, h_i^{n+1} - \mu_h^* \mathbf{n} \cdot \partial_{x_i} \mathbf{u}_h^* \rangle_{\Gamma_R^{n+1}}, \end{aligned}$$

for all $w_{i,h} \in X_h^{n+1}$, where $w_{i,h}|_{\Gamma_D^{n+1}} = 0$.

6. Recover the viscosity via projection of the rheological law given \mathbf{u}_h^{n+1} , i.e., find $\mu_h^{n+1} \in X_h^{n+1}$, such that

$$\langle r_h, \mu_h^{n+1} \rangle = \langle r_h, \eta(\dot{\gamma}(\nabla^s \mathbf{u}_h^{n+1})) \rangle \quad \forall r_h \in X_h^{n+1}.$$

7. Update the time step size Δt^n and corresponding coefficients α_j^m and β_j^m .

Similar to Thm. 4.1.4, we can proof the following Cor. 4.2.2.

Corollary 4.2.2. *The fully discretised convective ALE split-step scheme (4.51)–(4.59) fulfils the GCL given that all integrals are evaluated considering Ω_h^{n+1} and its respective*

boundary segments. Constant solutions \mathbf{u}_0 and p_0 with corresponding $\mu = \eta(0)$ are thus preserved at the discrete level.

Proof. In complete analogy to the proof of Thm. 4.1.4, no spurious terms spoil the conservation of constant solution components at the discrete level, where the only difference is that consistent Robin boundary data \mathbf{h} on Γ_R^t is assumed, similar to $\bar{\mathbf{t}}$ on Γ_N^t . \square

A detailed discussion of the algorithms laid out here are presented together with numerical results in the FSI context in Sec. 7.2, postponing further elaborations for the moment.

4.3 Stabilising dominant convection

The standard Galerkin finite element formulation of the fluid momentum balance equation as presented, e.g., in Chs. 2 or 3 is prone to instabilities, caused by steep gradients in the velocity field. In the convection-dominant regime, spurious oscillations arise, whenever the spatial discretisation cannot resolve the complex flow patterns with high enough resolution. In such regions, the required local element size to obtain a non-oscillatory solution decreases for higher Reynolds numbers. To counteract this phenomenon, residual-based schemes such as the Streamline upwind/Petrov–Galerkin formulation [73] (SUPG) or Galerkin Least-Squares stabilisation [33] (GLS) are among the most prominent techniques resorted to. Similar to residual-based inf-sup stabilisation, the weak form of the balance equations—in this case the momentum equation—is modified with a consistent term, which vanishes as the element size is decreased. Using projection or split-step schemes, the velocity step including convective terms can be stabilised employing these well-established techniques in a straight-forward manner if the splitting is performed on the continuous level [2, 6, 225]. Thus, in the present context, it suffices to add the stabilising terms

$$+\langle \tau_e \mathbf{L}_h(\mathbf{w}_h), \mathbf{R}_h(\mathbf{u}_h^{n+1}, p_h^{n+1}) \rangle_{\Omega_e^t} \quad (4.61)$$

in an element-wise manner, as indicated by Ω_e^t , to the weak form of the fluid momentum equation in the velocity-pressure or split-step scheme. Here, the fully discretised strong form residual $\mathbf{R}_h(\mathbf{u}_h^{n+1}, p_h^{n+1})$ reflects the specific form of the momentum balance equation, which might be similar to

$$\begin{aligned} \mathbf{R}_h(\mathbf{u}_h^{n+1}, p_h^{n+1}) &:= \rho \alpha_0^m \mathbf{u}_h^{n+1} + \rho \sum_{j=0}^{m-1} \alpha_{j+1}^m \left(\mathbf{u}_h^{n-j} - \nabla \psi_h^{n-j} \right) + \nabla p_h^{n+1} \\ &+ \rho \left(\nabla \mathbf{u}_h^{n+1} \right) \mathbf{c}_h^{n+1} - \mu_h^{n+1} \Delta \mathbf{u}_h^{n+1} - \chi \mu_h^{n+1} \nabla (\nabla \cdot \mathbf{u}_h) - 2 \nabla^s \mathbf{u}_h^{n+1} \nabla \mu_h^{n+1}, \end{aligned}$$

while $\mathbf{L}_h(\mathbf{w}_h)$ in Eqn. (4.61) is the corresponding discrete operator acting on the fluid velocity \mathbf{u} , which might read

$$\mathbf{L}_h(\mathbf{w}_h) := \rho \alpha_0^m \mathbf{w}_h + \rho (\nabla \mathbf{w}_h) \mathbf{c}_h^{n+1} - \mu_h^{n+1} \Delta \mathbf{w}_h - \chi \mu_h^{n+1} \nabla (\nabla \cdot \mathbf{w}_h) - 2 \nabla^s \mathbf{w}_h \nabla \mu_h^{n+1}.$$

In the SUPG method, the operator L_h is reduced to the convective term only, which gives (see, e.g., [27, 39, 226])

$$L_h(\mathbf{w}_h) := \rho (\nabla \mathbf{w}_h) \mathbf{c}_h^{n+1}.$$

Note here that in the split-step scheme, p_h^{n+1} is either computed prior to the momentum step, i.e., known or extrapolated as $p_h^{n+1} \approx p_h^*$, such that the pressure term in the discrete strong form residual $R_h(\mathbf{u}_h^{n+1}, p_h^{n+1})$ needs to be readjusted. Similar considerations apply regarding the sequence of steps, further linearisations used or in case the fluid pressure and Leray projection variable are combined into \hat{p}_h^* as demonstrated in Sec. 3.3. The resulting strong form residual $R(\mathbf{u}, p)$ and corresponding operator $L(\mathbf{w})$ are to be adapted accordingly. The stabilisation parameter τ_e in Eqn. (4.61) is defined similar to Eqns. (2.35) and (4.20) taken from [37]

$$\tau_e := \frac{\xi(\iota)}{\rho} \left[\left(\frac{4\mu^{n+1}}{\rho h_e^2} \right)^2 + \left(\frac{2}{\Delta t^n} \right)^2 + \left(\frac{2}{h_e} \|\mathbf{c}_h^{n+1}\| \right)^2 \right]^{-\frac{1}{2}}, \quad (4.62)$$

but this time includes an additional scaling parameter $\xi(\iota)$ as introduced by Hübner and Dinkler [38],

$$\xi(\iota) := \begin{cases} \frac{\iota}{3} & \text{for } \iota < 3, \\ 1 & \text{otherwise,} \end{cases} \quad \text{and} \quad \iota = \frac{\rho \|\nabla \mathbf{u}_h^{n+1}\| \|\mathbf{c}_h^{n+1}\| h_e^2}{4\mu_f^{n+1} \|\mathbf{u}_h^{n+1}\|}.$$

The additional scaling term $\xi(\iota)$ reduces the contribution of the convective stabilisation if the convective term $(\nabla \mathbf{u}) \mathbf{c}$ is small compared to the viscous term $2\mu \nabla^s \mathbf{u}$. Following this strategy, numerical artifacts caused by the inviscid residual when using lower-order elements are mitigated. To arrive at a linear problem, the stabilisation parameter is evaluated with \mathbf{u}_h^* and an appropriately linearised convective velocity \mathbf{c}_h^{n+1} . For the sake of brevity, we refrain from repeating the weak forms with different linearisations or formulations of the viscous term, only noting that the residual $R(\mathbf{u}, p)$ and the corresponding operator $L(\mathbf{w})$ have to reflect the employed strong form, multiplied by a potentially linearised τ_e . Numerical results employing GLS stabilisation are postponed to Secs. 8.5–8.8, closing the discussion on the stabilisations for convective effects for the moment.

4.4 Inflow over non-trivial inlets

In practical applications of flow problems in complex geometries or when the inflow boundary $\Gamma_{\text{in}}^t \subset \Gamma_D^t$ is not fixed, the construction of velocity profiles demands some further attention. Moreover, availability of data plays a role, as the imposition of either Dirichlet or Neumann data is in principle possible, but—taking simulations of flow through the cardiovascular system as an example—blood pressure and volumetric flow rate might not be equally accessible. That being said, the discussion in Sec. 4.5 on Neumann boundaries featuring re-entrant flow already hints at additional measures needed, if the (mean) pressure as part of the Cauchy traction is to be prescribed over Γ_{in} (see,

e.g., the work by Porpora et al. [227] for one possibility). Within this work, however, we restrict ourselves to volumetric flow rates enforced on plane, non-primitive and smooth inlets. Therefore, a technique by Takizawa et al. [228] is adopted, and shortly summarised here for the convenience of the reader.

So, the aim is to map a rotationally symmetric velocity profile defined on a plane circular inlet Γ_{circ} with radial coordinate r onto a more general, but still plane cross-section, yielding \mathbf{u}_{in}^t , while keeping the volumetric flow rate, defined as

$$Q_{\text{circ}}(t) := -\langle \mathbf{u}_{\text{circ}}(r, t), \mathbf{n}_{\text{circ}} \rangle_{\Gamma_{\text{circ}}}$$

unaltered, where \mathbf{n}_{circ} is the unit outward normal of a cylindrical domain with inlet Γ_{circ} . To achieve this, an equivalent radius \tilde{r} is defined for each point on Γ_{in} ,

$$\tilde{r}(\mathbf{x}) = r_M \frac{\|\mathbf{x} - \mathbf{x}_C\|}{\|\mathbf{x} - \mathbf{x}_C\| + \|\mathbf{x} - \mathbf{x}_B\|},$$

with \mathbf{x}_C denoting the inlet center of Γ_{in}^t , \mathbf{x}_B being the point on the inlet face's boundary $\partial\Gamma_{\text{in}}^t$ closest to the evaluation point \mathbf{x} and r_M , which is the mean radius of the inlet computed as the arithmetic mean of $\|\mathbf{x}_C - \mathbf{x}_B\|$. This equivalent radius \tilde{r} allows us to evaluate the given velocity profile \mathbf{u}_{circ} and set

$$\mathbf{u}_{\text{circ}}(r = \tilde{r}(\mathbf{x}), t) = \tilde{\mathbf{u}}_{\text{in}}(\tilde{r}(\mathbf{x}), t).$$

Now, this yields a reasonable velocity profile $\tilde{\mathbf{u}}_{\text{in}}$ on Γ_{in}^t , assuming a plane and smooth enough Γ_{in}^t only mildly deviating from a circular inlet. However, the volumetric flow rate has changed upon mapping the inlet profile, since now we have

$$Q_{\text{circ}}(t) \neq \tilde{Q}_{\text{in}}(t) := \langle \tilde{\mathbf{u}}_{\text{in}}(\tilde{r}(\mathbf{x}), t), \mathbf{n} \rangle_{\Gamma_{\text{in}}^t},$$

which we mend in a final step by rescaling the velocity,

$$\mathbf{u}_{\text{in}}(\mathbf{x}, t) := \frac{Q_{\text{circ}}}{\tilde{Q}_{\text{in}}} \tilde{\mathbf{u}}_{\text{in}}(\tilde{r}(\mathbf{x}, t)), \quad (4.63)$$

such that we obtain

$$Q_{\text{circ}}(t) = Q_{\text{in}}(t) := \langle \mathbf{u}_{\text{in}}(\mathbf{x}, t), \mathbf{n} \rangle_{\Gamma_{\text{in}}^t},$$

that is, the desired flow rate Q_{circ} , easily tuned via $\mathbf{u}_{\text{circ}}(r, t)$, and a reasonable velocity profile on the plane, but non-circular inlet Γ_{in}^t .

In general, the inlet Γ_{in}^t is not fixed in space, such that only the node indices in the finite element mesh corresponding to the points on the discretised Γ_{in}^t do not change over time using the same, but possibly mapped grid. If the inlet, however, is fixed in space, the factor scaling the inlet velocity magnitude $Q_{\text{circ}}/\tilde{Q}_{\text{in}}$ is constant and may be computed once up front before starting the timestepping loop. The numerical tests within this work consider parabolic profiles on Γ_{circ} only and we compute the scaling factor $Q_{\text{circ}}/\tilde{Q}_{\text{in}}$

at a time where $Q_{\text{circ}} > 0$ and hence $\tilde{Q}_{\text{in}} > 0$, sparing us from introducing further safety measures in Eqn. (4.63) to guard the case when $\tilde{Q}_{\text{in}} \approx 0$.

4.5 Backflow stabilisation

Many practical flow problems in science and engineering feature Neumann boundaries, where the fluid is expected to exit the computational domain. Depending on the boundary data prescribed, a certain reference pressure is set on such an outflow boundary denoted by Γ_{out} , while the velocity itself enters in terms of its gradient only. Thus, the actual flow direction remains sought, as desired in most such scenarios. Problems arise, if the pressure enforced on the outlet is high enough to cause flow reversal or strong vortices exit over Γ_{out} , such that $\mathbf{u} \cdot \mathbf{n} < 0$ on some part of the outlet. This results in problems estimating the energy balance, such that a-priori estimates are not available, potentially leading to severe instabilities or simulation divergence. Testing the convective ALE form of the Navier–Stokes equations with the solution (\mathbf{u}, p) , we obtain for the unforced case ($\mathbf{b} = \mathbf{0}$, $\mathbf{u}|_{\Gamma_D^t} = \mathbf{0}$, $\bar{\mathbf{t}}|_{\Gamma_N^t} = \mathbf{0}$)

$$0 = \langle \mathbf{u}, \rho \partial_t \mathbf{u} |_{\bar{\mathbf{x}}} + \rho (\nabla \mathbf{u}) \mathbf{c} \rangle + (1 - \chi/2) \langle \nabla \mathbf{u} + \chi \nabla^\top \mathbf{u}, \mu \nabla \mathbf{u} + \chi \mu \nabla^\top \mathbf{u} \rangle - (1 - \chi) \langle \mathbf{u}, (\nabla^\top \mathbf{u}) \nabla \mu \rangle, \quad (4.64)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard L^2 scalar product over the current domain Ω^t . Note also, that the generalised Laplacian ($\chi = 0$) form introduces an additional non-symmetric viscous term. This additional term demands some further attention when deriving a bound for the viscous energy, but is left as an open problem within this work. Regarding our discussions here, the acceleration and convective terms are relevant. With standard arguments (see, e.g., [24, 229–232]), we can rewrite the convective term as

$$\begin{aligned} \langle \mathbf{u}, \rho (\nabla \mathbf{u}) \mathbf{c} \rangle &= \langle \mathbf{c}, \rho (\nabla^\top \mathbf{u}) \mathbf{u} \rangle = \langle \mathbf{c}, \frac{\rho}{2} \nabla (\mathbf{u} \cdot \mathbf{u}) \rangle \\ &= \langle \mathbf{c} \cdot \mathbf{n}, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle_{\Gamma_N^t} - \langle \nabla \cdot \mathbf{c}, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle \\ &= \langle \mathbf{c} \cdot \mathbf{n}, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle_{\Gamma_N^t} + \langle \nabla \cdot \mathbf{u}_m, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle, \end{aligned} \quad (4.65)$$

since we defined the convective velocity as $\mathbf{c} := \mathbf{u} - \mathbf{u}_m$. The kinetic energy of the system shall be conserved, which may be formulated as (see, e.g., [215, 229])

$$\begin{aligned} \frac{\partial}{\partial t} E_{\text{kin}} &:= \frac{\partial}{\partial t} \int_{\Omega^0} \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \det \mathbf{J} \, d\Omega = \int_{\Omega^0} \frac{\partial}{\partial t} \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \det \mathbf{J} \Big|_{\bar{\mathbf{x}}} \, d\Omega \\ &= \int_{\Omega^0} \frac{\partial}{\partial t} \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \Big|_{\bar{\mathbf{x}}} \det \mathbf{J} + \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \frac{\partial}{\partial t} \det \mathbf{J} \Big|_{\bar{\mathbf{x}}} \, d\Omega \end{aligned}$$

with the Jacobian of the ALE mapping $\mathbf{J} := \frac{\partial \mathbf{x}}{\partial \bar{\mathbf{x}}}$, where one may insert (see, e.g., [214])

$$\frac{\partial}{\partial t} \det \mathbf{J} \Big|_{\bar{\mathbf{x}}} = (\nabla \cdot \mathbf{u}_m) \det \mathbf{J},$$

such that we get, changing the domain of integration to Ω^t ,

$$\begin{aligned} \frac{\partial}{\partial t} E_{\text{kin}} &= \int_{\Omega^0} \frac{\partial}{\partial t} \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \Big|_{\bar{\mathbf{x}}} \det \mathbf{J} + \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} (\nabla \cdot \mathbf{u}_m) \det \mathbf{J} \, d\Omega \\ &= \int_{\Omega^t} \frac{\partial}{\partial t} \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \Big|_{\bar{\mathbf{x}}} + \frac{\rho}{2} (\mathbf{u} \cdot \mathbf{u}) \nabla \cdot \mathbf{u}_m \, d\Omega \\ &= \langle \mathbf{u}, \rho \partial_t \mathbf{u} \Big|_{\bar{\mathbf{x}}} \rangle + \langle \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u}, \nabla \cdot \mathbf{u}_m \rangle. \end{aligned} \quad (4.66)$$

Then, we can insert Eqns. (4.65) and (4.66) into Eqn. (4.64), which yields in the inviscid case ($\mu \equiv 0$)

$$\begin{aligned} 0 &= \langle \mathbf{u}, \rho \partial_t \mathbf{u} \Big|_{\bar{\mathbf{x}}} + \rho (\nabla \mathbf{u}) \mathbf{c} \rangle \\ &= \frac{\partial}{\partial t} E_{\text{kin}} - \langle \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u}, \nabla \cdot \mathbf{u}_m \rangle + \langle \nabla \cdot \mathbf{u}_m, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle + \langle \mathbf{c} \cdot \mathbf{n}, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle_{\Gamma_N^t}, \end{aligned}$$

where the terms involving the divergence of the mesh velocity cancel out. Obviously, this presentation is not intended to be a rigorous proof of energy-stability, but it is sufficient to demonstrate that even when only the acceleration and convective terms are considered, one ends up with

$$\frac{\partial}{\partial t} E_{\text{kin}} = - \langle \mathbf{c} \cdot \mathbf{n}, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle_{\Gamma_N^t}.$$

To obtain an energy-stable scheme, we require the flux kinetic energy entering over Γ_N^t be negative, such that $\frac{\partial}{\partial t} E_{\text{kin}} \leq 0$ in the above relation. As flow enters the computational domain in a relative sense, i.e., $\mathbf{c} \cdot \mathbf{n} < 0$, the system's kinetic energy cannot be bound, existence of weak solutions cannot be proven and a-priori estimates are unattainable [229–231, 233].

Several works in literature are devoted to equalising this term by modifying the Neumann condition on Γ_N^t to guarantee stability, while keeping the impact on the solution as little as possible. A first option is to enforce the total pressure $p_{\text{tot}} = p + \frac{1}{2} \mathbf{u} \cdot \mathbf{u}$ on Γ_N^t [234, 235], leading to stable, but unphysical solutions when $\frac{\rho}{2} \mathbf{u} \cdot \mathbf{u}$ is of the same order of magnitude as the pressure itself [24]. Also, this modification alters the boundary condition in the case of outflow, which is certainly not desired in most cases.

Targeting directly the boundary term that spoils energy balance, the popular velocity-penalisation method [236, 237] and its variants add terms of the form

$$\langle \mathbf{w}, \alpha_b \frac{\rho}{2} \mathbf{u} (\mathbf{c} \cdot \mathbf{n})_- \rangle_{\Gamma_N^t} \quad (4.67)$$

with $(\mathbf{c} \cdot \mathbf{n})_- := \frac{1}{2} (|\mathbf{c} \cdot \mathbf{n}| - \mathbf{c} \cdot \mathbf{n}) = \begin{cases} -\mathbf{c} \cdot \mathbf{n} & \text{for } \mathbf{c} \cdot \mathbf{n} < 0, \\ 0 & \text{otherwise.} \end{cases}$

to the weak form of the momentum balance equation. Then, testing the weak form with the solution (\mathbf{u}, p) gives in the inviscid, unforced case

$$\frac{\partial}{\partial t} E_{\text{kin}} = -\langle \mathbf{c} \cdot \mathbf{n}, \frac{\rho}{2} \mathbf{u} \cdot \mathbf{u} \rangle_{\Gamma_N^t} - \langle \mathbf{u}, \alpha_b \frac{\rho}{2} \mathbf{u} (\mathbf{c} \cdot \mathbf{n})_- \rangle_{\Gamma_N^t},$$

such that $\alpha_b \geq 1$ guarantees an energy-stable scheme (see [231, 233] for the zero pseudo tractions case and [229, 230] for pressure-gradient-driven fluid flows). For the sake of completeness, we mention other well-established techniques, such as penalising the velocity gradient [238, 239], enforcing mean flow rates via Lagrange multipliers [240], extending the computational domain artificially and adapting the balance equations in those regions [241, 242] or enforcing an outflow profile [243] or zero tangential velocity components [244]. The interested reader is referred to the excellent works by Bertoglio et al. [245] and Fouchet-Incaux [229] for a more detailed discussion and in-depth analysis.

Herein, we restrict ourselves to backflow stabilisation adding the stabilisation term given in Eqn. (4.67), since the studies presented in [241, 244, 245] and numerous applications in both the haemodynamic [227, 244, 246–248] and respiratory [230, 249, 250] regime show satisfactory results. Adding the term in Eqn. (4.67) on the Neumann boundary Γ_N^t amounts to changing the (pseudo) traction vector weakly enforced on the Neumann boundary segment to

$$\left(-p\mathbf{I} + \mu\nabla\mathbf{u} + \chi\mu\nabla^\top\mathbf{u} \right) \mathbf{n} = \bar{\mathbf{t}} - \alpha_b \frac{\rho}{2} \mathbf{u} (\mathbf{c} \cdot \mathbf{n})_-, \quad (4.68)$$

where χ once again switches between the generalised Laplacian ($\chi = 0$) and stress-divergence ($\chi = 1$) forms. For the velocity-pressure formulation, the corresponding term is easily added to the momentum balance equation, whereas in the split-step scheme, this leads to further complications. As consistency of the PPE and its boundary conditions is key to mass conservation, the change in the Neumann conditions has to be treated with care. Omitting yet another equivalence proof similar to Thm. 2.1.1 or Cors. 4.1.3 and 4.2.1 with a slightly altered Neumann boundary term on Γ_N^t for the sake of brevity, we directly proceed in presenting the modified pressure Dirichlet condition on Γ_N^t , which now reads

$$p|_{\Gamma_N^t} = \zeta := -\mu\nabla \cdot \mathbf{u} + \mathbf{n} \cdot \left[\mu \left(\nabla\mathbf{u} + \chi\nabla^\top\mathbf{u} \right) \mathbf{n} - \bar{\mathbf{t}} + \alpha_b \frac{\rho}{2} \mathbf{u} (\mathbf{c} \cdot \mathbf{n})_- \right],$$

incorporating the stabilisation term in a straight-forward manner. The modified traction thus not only enters the fluid momentum balance step as was intended from the beginning, but also the projection on the Neumann boundary to obtain continuous boundary

data for the PPE. Hence, the projection step in the split-step scheme seeks for continuous pressure boundary data $\zeta_h^{n+1} \in Y_h^{n+1} \subset H^{1/2}(\Gamma_N^{n+1})$, such that for all $s_h \in Y_h^{n+1}$, there holds

$$\begin{aligned} \langle s_h, \zeta_h^{n+1} \rangle_{\Gamma_N^t} &= \langle s_h \mathbf{n}, \mu_h^{n+1} (\nabla \mathbf{u}_h^{n+1} + \chi \nabla^\top \mathbf{u}_h^{n+1}) \mathbf{n} - \bar{\mathbf{t}}^{n+1} + \alpha_b \frac{\rho}{2} \mathbf{u}_h^{n+1} (\mathbf{c}_h^{n+1} \cdot \mathbf{n})_- \rangle_{\Gamma_N^t} \\ &\quad - \langle s_h, \mu_h^{n+1} \nabla \cdot \mathbf{u}_h^{n+1} \rangle_{\Gamma_N^t}. \end{aligned} \quad (4.69)$$

Depending on the specific scheme chosen, that is, the sequence of steps, combining the pressure and Leray projection variables or extrapolating in time, the backflow stabilisation needs to be evaluated with a suitable convective velocity. The momentum balance step in component-coupling form with divergence damping then consists of finding $\mathbf{u}_h^{n+1} \in [X_h^{n+1}]^d \subset [H^1(\Omega^{n+1})]^d$ with $\mathbf{u}_h^{n+1}|_{\Gamma_D^{n+1}} = \mathbf{g}^{n+1}$, such that there holds

$$\begin{aligned} 0 &= \langle \mathbf{w}_h, \rho \alpha_0^m \mathbf{u}_h^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m (\mathbf{u}_h^{n-j} - \nabla \psi_h^{n-j}) \rangle + \langle \mathbf{w}_h, \rho (\nabla \mathbf{u}_h^{n+1}) \mathbf{c}_h^{n+1} - \mathbf{b}^{n+1} \rangle \\ &+ (1 - \chi/2) \langle \nabla \mathbf{w}_h + \chi \nabla^\top \mathbf{w}_h, \mu_h^{n+1} \nabla \mathbf{u}_h^{n+1} + \chi \mu_h^{n+1} \nabla^\top \mathbf{u}_h^{n+1} \rangle - (1 - \chi) \langle \mathbf{w}_h, (\nabla^\top \mathbf{u}_h^{n+1}) \nabla \mu_h^{n+1} \rangle \\ &\quad - \langle \nabla \cdot \mathbf{w}_h, p_h^{n+1} \rangle - \langle \mathbf{w}_h, \bar{\mathbf{t}}^{n+1} \rangle_{\Gamma_N^{n+1}} + \langle \mathbf{w}_h, \alpha_b \frac{\rho}{2} \mathbf{u}_h^{n+1} (\mathbf{c}_h^{n+1} \cdot \mathbf{n})_- \rangle_{\Gamma_N^{n+1}}, \end{aligned}$$

for all $\mathbf{w}_h \in [X_h^{n+1}]^d$ with $\mathbf{w}_h|_{\Gamma_D^{n+1}} = \mathbf{0}$. Then, an adequate strategy is to consider the approximation $(\mathbf{c}_h^{n+1} \cdot \mathbf{n})_- \approx (\mathbf{c}_h^* \cdot \mathbf{n})_-$ in both PPE *and* momentum balance steps. In the latter, the backflow stabilisation is therefore considered semi-implicit, entering the matrix of the linear system corresponding to the variational formulation of the momentum balance equation, but “activating” backflow stabilisation based on \mathbf{c}_h^* . Evaluating the unit outward normal dotted with the convective velocity \mathbf{c}_h^* is enough to prevent divergence due to re-entrant flow as our numerical tests suggest. Note here, however, that the activation of the stabilising term has to be synchronised between the PPE and momentum balance steps as its contribution can become quite large and might otherwise lead to instabilities. This leads to using $(\mathbf{c}_h^* \cdot \mathbf{n})_-$ for the pressure steps, *even if* the fluid velocity might be readily available.

We refrain at this point from presenting every possible combination in terms of sequence of steps and suitable extrapolations, noting, that above remarks are applicable to all of the variants presented in Ch. 3 and their ALE counterparts in Sec. 4.1. Numerical examples are presented in the FSI context and are thus postponed to Ch. 7 for now.

4.6 Outflow conditions for truncated domains

In addition to the problems encountered when fluid enters the computational domain over Neumann boundaries as discussed in Sec. 4.5, prescribing suitable traction data is non-trivial on boundaries where the domain has to be truncated artificially. At those predefined outlets, the otherwise neglected downstream structures have to be taken into account since they can heavily impact the flow split between outlets by setting appropriate pressure levels. In patient-specific scenarios of aortic blood flow for example, neither

pressure nor velocity profiles are routinely acquired in clinic and thus seldomly available. Simple strategies to resolve this issue might prescribe, e.g., mean reference pressures at all times or velocity profiles computed from volumetric flow data, both of which can lead to stability issues or unphysical results depending on the physical parameters and the specific problem considered. In addition to that, the required flow quantities are often times not available, such that population mean values have to be resorted to, further reducing the predictive power of the numerical model.

The usual boundary conditions are thus replaced by transmission conditions to reduced zero-dimensional Windkessel models, one-dimensional pipe flow or networks thereof (see, e.g., [251], [252] and [253–255] and references therein) to incorporate effects of the cut-off flow regions. This leads to several frameworks coupling pressure, flow rates, or cross-sectional area of one-dimensional models [256–260], closed-loop pipe networks [261–265] or zero-dimensional lumped parameter models [229, 241, 244, 247, 248, 266–272] to the three-dimensional Navier–Stokes equations in pure CFD or FSI settings. In the last group of methods, Windkessel (WK) models describing the relation between volumetric flow and pressure are among the most prominent techniques owing to their simplicity and great versatility. The classical three-element WK model [251, 273] as one of the possible options is employed within this work, but might be viewed as a placeholder for more complex differential algebraic systems (see, e.g., [267]).

So, we proceed in decomposing the Neumann boundary Γ_N^t into non-overlapping outflow sections denoted as Γ_i^t , $i = 1, \dots, N_{\text{out}}$,

$$\Gamma_N^t =: \Gamma_{\text{out}}^t = \bigcup_{i=1}^{N_{\text{out}}} \Gamma_i^t,$$

such that $\bar{\Gamma}_i^t \cap \bar{\Gamma}_j^t = \emptyset$ for $i, j = 1, \dots, N_{\text{out}}$ with $i \neq j$. Over each of these outlet faces, define the outward flux $Q_i(\mathbf{u})$ via

$$Q_i(\mathbf{u}) := \langle \mathbf{u}, \mathbf{n} \rangle_{\Gamma_i^t}, \quad (4.70)$$

such that the three-element WK model governing the mean pressure $p_{c,i}$ on outlet Γ_i^t is given as

$$C_i \frac{\partial}{\partial t} p_{p,i} + \frac{p_{p,i} - p_{d,i}}{R_{d,i}} = Q_i(\mathbf{u}), \quad (4.71)$$

$$p_{c,i} - p_{p,i} = R_{p,i} Q_i(\mathbf{u}), \quad (4.72)$$

for $i = 1, \dots, N_{\text{out}}$. Here, the capacitance C_i incorporates elastic effects of the neglected pipe network, $p_{p,i}$ and $p_{d,i}$ denote proximal and distal pressures, whereas the proximal and distal resistances to the flow, $R_{p,i}$ and $R_{d,i}$, control the amount of backflow by regulating the pressure. The mean pressure $p_{c,i}$ is enforced via traction conditions on Γ_i^t ,

$$\left(-p\mathbf{I} + \mu\nabla\mathbf{u} + \chi\mu\nabla^\top\mathbf{u} \right) \mathbf{n} = -p_{c,i} \mathbf{n},$$

or, including terms related to backflow stabilisation similar to Eqn. (4.68) to counter re-entrant flow, one obtains (see Sec. 4.5)

$$\left(-p\mathbf{I} + \mu\nabla\mathbf{u} + \chi\mu\nabla^\top\mathbf{u}\right)\mathbf{n} = -p_{c,i}\mathbf{n} - \alpha_b\frac{\rho}{2}\mathbf{u}(\mathbf{c}\cdot\mathbf{n}), \quad (4.73)$$

for outlet $i = 1, \dots, N_{\text{out}}$. Note here that the viscous contribution to the traction is completely neglected, following the standard approach. In an attempt to reduce effects on the upstream velocity field, one might incorporate appropriate correcting boundary integrals, but given the fact that in most cardiovascular applications active backflow stabilisation and the pressure level dominate the flow split and the observed flow fields at the outlets, this simplification is accepted, simply ignoring the introduced errors due to their minor relative importance.

A consistent time integration of the differential algebraic system of the WK model (4.71)–(4.72) via a BDF scheme of order m reads

$$C_i \sum_{j=0}^m \alpha_j^m p_{p,i}^{n+1-j} + \frac{p_{p,i}^{n+1} - p_{d,i}^{n+1}}{R_{d,i}} = Q_i(\mathbf{u}^{n+1}), \quad (4.74)$$

$$p_{c,i}^{n+1} - p_{p,i}^{n+1} = R_{p,i} Q_i(\mathbf{u}^{n+1}), \quad (4.75)$$

for each of the $i = 1, \dots, N_{\text{out}}$ outlets. However, $Q_i(\mathbf{u}^{n+1})$ is not directly available for the Neumann condition to enforce $p_{c,i}^{n+1}$ implicitly. In a velocity-pressure formulation or in the momentum balance step of the split-step scheme, $Q_i(\mathbf{u}^{n+1})$ can be considered implicitly, expressing Q_i in terms of the unknown nodal velocity DoFs, being the most stable option. Alternatively, linearisations via $Q_i(\mathbf{u}^n)$ or $Q_i(\mathbf{u}^*)$ are other possible options, potentially affecting the temporal stability of the overall scheme [241, 267].

Numerical examples and further discussion are postponed to Ch. 8, where linearisation and extrapolation options are compared. There, the three-element Windkessel model is applied to an idealised abdominal aortic aneurysm (single outlet), and patient-specific geometries constructed from medical image data of an iliac bifurcation (2 outlets) and a case of aortic dissection (12 outlets) in Secs. 8.5–8.7.

4.7 Summary and conclusion

Targeting practical applications in the cardiovascular context, this chapter presents extensions of the coupled velocity-pressure formulation stabilised via BVS and the PPE-based split-step scheme, incorporating well-established models and concepts. First, Arbitrary Lagrangian–Eulerian formulations are considered accounting for instationary domains as encountered, e.g., in fluid–structure interaction or simply when the domain’s boundaries follow a prescribed motion. Motivated by several studies in the field, we include Robin boundary conditions for the coupled velocity-pressure formulation as well as the PPE-based split-step scheme for later use in FSI, which might also be applied to enclosed flow problems, ensuring a unique pressure field by design. Moreover, practically relevant aspects such as stabilising convective effects via the Streamline upwind/Petrov–

Galerkin (SUPG) formulation or Galerkin Least-Squares (GLS) stabilisation, constructing inflow profiles for non-trivial inlets given volumetric flow data, stabilising re-entrant flow regions and accounting for the neglected downstream vasculature via lumped zero-dimensional parameter models are carefully woven into the proposed framework.

The following chapter now shifts the focus to structural dynamics, introducing suitable initial boundary value problems and a corresponding solver. Basic, purely displacement-based and C^0 -continuous formulations are discussed, considering linear elasticity, a St. Venant–Kirchhoff model or standard constitutive relations in Hyperelasticity such as the neo-Hookean, Mooney–Rivlin or Holzapfel–Gasser–Ogden material models, weakly enforcing incompressibility through a penalty term. Afterwards, practical aspects of relevance in cardiovascular applications are covered in Ch. 6, before a PPE-based split-step FSI scheme is then introduced in Ch. 7.

5 Structural dynamics

In many problems of engineering and biomedical interest, elastic continua play a significant role, potentially undergoing large deformations and showing distinct anisotropic and highly nonlinear material behaviour. Solid mechanics problems and their numerical solution have a long tradition, which is closely related to developments of the finite element method. Soft biological tissues, however, are in this regard particularly challenging, as standard solution techniques can lead to numerical difficulties in the relevant parameter range for (nearly) incompressible continua [274–279]. To capture the rubber-like behaviour of such materials, a vast variety of methods have been presented over the past years to completely resolve or at least reduce the incompressibility dilemma and its associated numerical challenges.

One can differentiate mixed displacement–pressure formulations using inf-sup stable finite element pairs (see, e.g., [280–282]) or stabilised, equal-order formulations such as [283–291], enhanced strain methods [287, 292–295], local projection schemes [296–300], orthogonal sub-scale approaches [301–303], multi-field variational formulations [304–306], methods based on non-conforming finite elements such as Crouzoix–Raviart or discontinuous Lagrange elements [307–311] and possibly the most fundamental class of purely displacement-based penalty schemes [275, 312–315]. Some of these approaches might still suffer from locking effects, especially for lower-order interpolation and in the incompressible limit as the bulk modulus enforcing isochoric deformation increases (and the Poisson’s ratio approaches $1/2$).

However, there is still a certain appeal to lower-order interpolations, mainly for reasons of practicality [291, 316, 317]: first, real-world applications in the biomedical context often feature high-resolution, but lower-order accurate geometrical representations, as they are reconstructed from voxelated medical image data. Second, the sought solutions might not be smooth enough for higher-order methods to unleash their full potential without further measures taken. Furthermore, the solid mechanics solver is within this work utilised as a single-field solver in a fluid–structure interaction algorithm, where only matching grids on the interface are considered for ease of implementation, avoiding projections or interpolations on the interface. Therefore, the potential structure formulations are limited to schemes involving continuous displacement approximations, as continuity of displacements and velocities is enforced onto the fluid solver in a standard strong way, i.e., directly woven into the function space. On the solid side, the applied Robin or Neumann conditions on the interface involve the fluid Cauchy traction and velocity, such that no additional constraints are strongly enforced. Also note that in the coupling scheme presented in Sec. 7.2, the solid subproblem and a corresponding solver are easily replaced, such that more advanced approaches can be integrated into the final resulting FSI scheme due to its partitioned design.

Keeping these considerations in mind, we herein employ one of the least complex options using a purely displacement-based formulation with standard continuous interpolation and a modified strain energy function involving a penalty term to weakly enforce incompressibility. Also, we restrict ourselves to hyperelastic material laws within this work, only consider viscoelastic effects in the surrounding tissue support and account for loading history only in terms of the prestress being present in the reference configuration.

Regarding the replication of in-vivo scenarios, the current state of the art in modelling the cardiovascular system further involves: (i) layered geometric discretisations of the aortic wall, resolving the (major) tissue layers being intima, media and adventitia, (ii) (nearly) incompressible, anisotropic material models including fibre-reinforcement, (iii) surrounding tissue providing viscoelastic support and embedding the computational domain, and (iv) accounting for prestress, being the stress state present in geometries constructed from medical image data at the time of image acquisition. All these aspects have not been embedded in a *single* (FSI) framework so far, and thus constitute the next step towards a fast computational tool for clinical support. Within this chapter, the basic solid mechanics solver including the various modelling aspects (i)–(iv) is introduced.

Based on a formulation taken from Simo and Hughes [313], novel contributions of the present work are the rule-based construction of a material orientation suitable for applications featuring branching, bifurcations and even topology changes as encountered, e.g., in aortic dissection (see Sec. 6.1). Also, the prestress algorithm by Hsu and Bazilevs [318] is robustified by a combined continuation and load-driven pseudo-timestepping approach presented in Sec. 6.4. While adopting Robin boundary and interface conditions as summarised in Secs. 6.2 and 6.3 cannot be considered new, their combination into the overall semi-implicit PPE-based FSI framework in Ch. 7 certainly is, motivating its introduction for later reference. The bulk of these methods are introduced in the author’s work [6] and are introduced here in greater detail, complemented by numerous remarks.

5.1 Governing equations

The field of continuum mechanics has a long and rich history with countless groundbreaking contributions by various authors. It thus comes at no surprise that restricting ourselves to a short introduction for brevity inevitably ignores the vast majority of important problems and possible remedies presented in the literature. Nonetheless, we want to point the interested reader towards the excellent monographs by Holzapfel [276], Bonet and Wood [319] and related works [100, 278, 313, 320–322], which might be useful as a starting point, presenting a much more rigorous and in-depth discussion.

With that being said, let us consider a continuum body in the current or spatial domain Ω^t and let $\hat{\Omega}$ denote the corresponding Lagrangian reference configuration. The Lipschitz-continuous boundary Γ^t and its Lagrangian counterpart $\hat{\Gamma}$ are further decomposed into non-overlapping and non-empty Dirichlet and Neumann segments, Γ_D^t and Γ_N^t and respective $\hat{\Gamma}_D$ and $\hat{\Gamma}_N$. The Lagrangian and current configurations are connected via the mapping

$$\phi : \hat{\mathbf{x}} \in \hat{\Omega} \rightarrow \mathbf{x} \in \Omega^t, \quad \phi(\hat{\mathbf{x}}, t) := \hat{\mathbf{x}} + \mathbf{d}(\hat{\mathbf{x}}, t), \quad (5.1)$$

with the material point $\hat{\mathbf{x}}$ mapped to the current configuration Ω^t via the displacement $\mathbf{d}(\hat{\mathbf{x}}, t)$ from the fixed Lagrangian reference frame $\hat{\Omega}$ (see also Sec. 4.1 and, e.g., [276, 319]). Here and in what follows, we purposely refrain from indicating quantities defined in $\hat{\Omega}$ by $\hat{\cdot}$ for any quantities undoubtedly corresponding to the Lagrangian frame to not further clutter the notation, which will become particularly relevant when introducing the FSI problem later in Ch. 7.

At initial time $t = 0$, the Lagrangian configuration and the current (also called spatial) domain coincide, whereas for $t > 0$, the motion of the continuum body is described in terms of the displacement vector $\mathbf{d}(\hat{\mathbf{x}}, t)$. Then, the total Lagrangian framework conveniently allows expressing stresses and strains in Ω^t in terms of the mapping $\phi(\hat{\mathbf{x}}, t)$ —or more precisely—the displacement $\mathbf{d}(\hat{\mathbf{x}}, t)$, defined in the stationary Lagrangian configuration $\hat{\Omega}$. To this end, we further introduce the deformation gradient \mathbf{F} and its determinant, the Jacobian J , as

$$\mathbf{F} := \hat{\nabla} \phi = \hat{\nabla} \hat{\mathbf{x}} + \hat{\nabla} \mathbf{d} = \mathbf{I} + \hat{\nabla} \mathbf{d}, \quad J := \det \mathbf{F}, \quad (5.2)$$

adopting standard notations. Here, $\hat{\nabla}$ denotes the gradient with respect to material coordinates $\hat{\mathbf{x}}$. To ensure that the mapping ϕ is one-to-one and invertible, we require the Jacobian J be strictly positive, such that the deformation gradient \mathbf{F} is non-singular and \mathbf{F}^{-1} exists. The exact displacement field \mathbf{d} is assumed continuous (no cracks or voids forming), which also holds on the discrete level when using continuous finite elements. Physically interpreting these statements, we ensure that the continuum body does not penetrate itself, if no self-contact occurs.

To conveniently express strains in the Lagrangian configuration, we additionally define the right Cauchy–Green tensor \mathbf{C} ,

$$\mathbf{C} := \mathbf{F}^\top \mathbf{F}, \quad (5.3)$$

and the Green–Lagrange strain tensor

$$\mathbf{E} := 1/2(\mathbf{C} - \mathbf{I}) = 1/2(\mathbf{F}^\top \mathbf{F} - \mathbf{I}) = 1/2[\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d} + (\hat{\nabla}^\top \mathbf{d}) \hat{\nabla} \mathbf{d}], \quad (5.4)$$

introducing a shorthand notation for the transpose of the displacement gradient with respect to $\hat{\mathbf{x}}$ by $\hat{\nabla}^\top \mathbf{d} := (\hat{\nabla} \mathbf{d})^\top$.

Cauchy’s stress theorem states, that at every continuum point, there exists a unique second-order tensor, which, when dotted with the unit outward normal of an arbitrary cutting plane, exactly equates the traction acting on that plane. Denoting by \mathbf{n} the unit normal onto an arbitrary cutting plane in the current configuration Ω^t , this can be written as

$$\mathbf{t} = \boldsymbol{\sigma} \mathbf{n}, \quad (5.5)$$

with the Cauchy or real traction vector \mathbf{t} and the Cauchy stress tensor $\boldsymbol{\sigma}$ as already considered in the context of fluid mechanics in Chs. 2–4. In the Lagrangian reference configuration $\hat{\Omega}$, a similar argument relating the referential traction vector $\hat{\mathbf{t}}$ with a second-order tensor and the plane's unit normal $\hat{\mathbf{n}}$ reads

$$\hat{\mathbf{t}} = \mathbf{P} \hat{\mathbf{n}}. \quad (5.6)$$

Herein, \mathbf{P} denotes the first Piola–Kirchhoff stress tensor defined in the Lagrangian reference configuration $\hat{\Omega}$. With the help of Nanson's formula, a relation between the two stress tensors $\boldsymbol{\sigma}$ and \mathbf{P} can be found as (cf. [276, 319])

$$\mathbf{P} = J \boldsymbol{\sigma} \mathbf{F}^{-\top}. \quad (5.7)$$

Furthermore, we finally introduce the pull-back of the first Piola–Kirchhoff stress tensor \mathbf{P} as the second Piola–Kirchhoff stress tensor \mathbf{S} ,

$$\mathbf{S} := \mathbf{F}^{-1} \mathbf{P} = J \mathbf{F}^{-1} \boldsymbol{\sigma} \mathbf{F}^{-\top}, \quad (5.8)$$

directly using Eqn. (5.7) to express \mathbf{S} in terms of $\boldsymbol{\sigma}$. Note that the Cauchy stress $\boldsymbol{\sigma}$ is symmetric as a consequence of angular momentum balance (see, e.g., [276, 319]), while the deformation gradient \mathbf{F} is not. Consequently, \mathbf{S} is symmetric too, but \mathbf{P} is not. The second Piola–Kirchhoff stress \mathbf{S} is at this point a purely artificial quantity without any real physical interpretation. Nonetheless, it has proven to be useful in the definition of constitutive relations, which are the missing link between stress and strain measures. It can thus be used to conveniently describe relations between, e.g., the Green–Lagrange strain tensor \mathbf{E} (5.4) or invariants of the right Cauchy–Green strain tensor \mathbf{C} (5.3) and the second Piola–Kirchhoff stress tensor \mathbf{S} (5.8).

The balance of linear momentum connecting the rate of change of linear momentum and externally applied loads in the Lagrangian reference frame may then be stated as

$$\rho \mathbf{d}_{tt} \mathbf{d} - \hat{\nabla} \cdot \mathbf{P} = \mathbf{b} \quad \text{in } \hat{\Omega},$$

with the structure's density per unit reference volume ρ , the material acceleration $\mathbf{d}_{tt} \mathbf{d}$, the divergence of the first Piola–Kirchhoff stress tensor with respect to the material coordinates $\hat{\nabla} \cdot \mathbf{P}$ and $\mathbf{b}(\hat{\mathbf{x}}, t)$ being a volumetric force per unit reference volume. A standard initial boundary value problem governing the structure's displacement \mathbf{d} in the time interval from $t = 0$ to $t = T$ including Dirichlet and Neumann boundary conditions on $\hat{\Gamma}_D$ and $\hat{\Gamma}_N$, respectively, reads

$$\rho \mathbf{d}_{tt} \mathbf{d} - \hat{\nabla} \cdot \mathbf{P} = \mathbf{b} \quad \text{in } \hat{\Omega} \times (0, T], \quad (5.9)$$

$$\mathbf{d}(\hat{\mathbf{x}}, 0) = \mathbf{d}_0 \quad \text{in } \hat{\Omega} \text{ at } t = 0, \quad (5.10)$$

$$\mathbf{d}_t \mathbf{d}(\hat{\mathbf{x}}, 0) = \dot{\mathbf{d}}_0 \quad \text{in } \hat{\Omega} \text{ at } t = 0, \quad (5.11)$$

$$\mathbf{d} = \mathbf{g} \quad \text{on } \hat{\Gamma}_D \times (0, T], \quad (5.12)$$

$$\mathbf{P} \hat{\mathbf{n}} = \hat{\mathbf{t}} \quad \text{on } \hat{\Gamma}_N \times (0, T], \quad (5.13)$$

where \mathbf{d}_0 and $\dot{\mathbf{d}}_0$ are given initial displacement and velocity fields, \mathbf{g} is a given Dirichlet datum on $\hat{\Gamma}_D$ and the traction vector $\hat{\mathbf{t}}$ given on $\hat{\Gamma}_N$.

Now, a constitutive relation to determine $\mathbf{P}(\mathbf{d})$ needs to be defined to capture the material behaviour. In order to derive such a relation, hyperelastic continua are based on the fundamental assumption that there exists a scalar-valued strain-energy density per unit reference volume $\Psi = \Psi(\mathbf{E}) : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^+$. For purely mechanical problems, the total internal potential energy $\Pi_{\text{int}}(t)$ is then given by

$$\Pi_{\text{int}} = \int_{\hat{\Omega}} \Psi \, d\hat{\Omega}.$$

The constitutive equation for the second Piola–Kirchhoff stress tensor can then be formulated as

$$\mathbf{S} = \frac{\partial \Psi}{\partial \mathbf{E}}, \quad (5.14)$$

such that the system (5.9)–(5.13) is finally closed. A first simple constitutive relation connecting the Green–Lagrange strain tensor \mathbf{E} and the strain-energy density Ψ is the so-called Saint Venant–Kirchhoff model (see, e.g., [276, 320]), specified as

$$\Psi(\mathbf{E}) := \frac{\lambda}{2} (\text{tr } \mathbf{E})^2 + \mu \text{tr}(\mathbf{E}^2), \quad (5.15)$$

with λ and μ being Lamé parameters. For homogeneous isotropic materials these parameters may be expressed in terms of Young’s modulus E and Poisson’s ratio ν , yielding

$$\lambda = \frac{E\nu}{(1-2\nu)(1+\nu)}, \quad \mu = \frac{E}{2(1+\nu)}, \quad (5.16)$$

where one can already see that for $\nu \rightarrow 1/2$, that is, approaching the fully incompressible case, numerical difficulties arise. Then, owing to Eqn (5.14), the second Piola–Kirchhoff stress tensor for Saint Venant–Kirchhoff material is given by

$$\mathbf{S} = \lambda \mathbf{I} \text{tr } \mathbf{E} + 2\mu \mathbf{E}.$$

In the context of fluid mechanics, μ translates to the dynamic viscosity (see, e.g., Chs. 2 and 3), whereas in solid mechanics, one refers to the second Lamé parameter μ as the shear modulus. Similarly, the bulk modulus κ_b can be written for homogeneous isotropic materials as

$$\kappa_b = \frac{E}{3(1-2\nu)}. \quad (5.17)$$

Expressed in terms of Lamé parameters, we have

$$E = \mu \frac{3\lambda + 2\mu}{\lambda + \mu}, \quad \nu = \frac{\lambda}{2(\lambda + \mu)}, \quad \kappa_b = \lambda + \frac{2}{3}\mu, \quad (5.18)$$

which we introduce here for later reference. In the small strain limit, a frequently used simplification is based on further assuming $\|\hat{\nabla} \mathbf{d}\| \ll 1$, which leads to a linearised strain tensor $\boldsymbol{\epsilon}$ being the symmetric gradient of \mathbf{d} with respect to $\hat{\mathbf{x}}$, since

$$\mathbf{E} := 1/2 \left[\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d} + \left(\hat{\nabla}^\top \mathbf{d} \right) \hat{\nabla} \mathbf{d} \right] \approx 1/2 \left(\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d} \right) =: \boldsymbol{\epsilon},$$

deducing a stress tensor linear in the displacement \mathbf{d} . Aiming for governing equations linear in the displacement, the Piola transform is neglected, such that stress tensors in the spatial and material configuration coincide, i.e., Eqn. (5.8) reduces to

$$\mathbf{P} = \mathbf{F} \mathbf{S} \approx \mathbf{S}, \quad \mathbf{P} = J \boldsymbol{\sigma} \mathbf{F}^{-\top} \approx \boldsymbol{\sigma}.$$

The assumptions of linear elasticity thus finally lead to

$$\mathbf{P} \approx \mathbf{S} \approx \lambda \mathbf{I} \operatorname{tr} \boldsymbol{\epsilon} + 2\mu \boldsymbol{\epsilon} = \lambda \mathbf{I} \hat{\nabla} \cdot \mathbf{d} + \mu \left(\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d} \right), \quad (5.19)$$

whereas the Saint Venant–Kirchhoff model (5.15) leads to the nonlinear expression

$$\mathbf{P} = \mathbf{F} \mathbf{S} = \lambda \mathbf{F} \operatorname{tr} \mathbf{E} + 2\mu \mathbf{F} \mathbf{E}. \quad (5.20)$$

Note here that λ enforces a divergence-free displacement field via a penalty term, leading to fully incompressible material behaviour as the Poisson’s ratio $\nu \rightarrow 0.5$ and consequently $\lambda \rightarrow \infty$. In such a scenario—as encountered in the case of rubber-like materials or soft biological tissue—numerical difficulties are inevitable when using the standard formulation just presented. As such materials’ response to shear and compression/extension are substantially different, the deformation gradient \mathbf{F} is split multiplicatively into so-called isochoric and volumetric parts following Flory [323],

$$\begin{aligned} \mathbf{F} &= \mathbf{F}_{\text{vol}} \bar{\mathbf{F}}, \\ \text{where } \mathbf{F}_{\text{vol}} &:= J^{1/3} \mathbf{I} \quad \text{and} \quad \bar{\mathbf{F}} := J^{-1/3} \mathbf{F}. \end{aligned} \quad (5.21)$$

Consequently, \mathbf{F}_{vol} captures all volumetric changes due to $\det \mathbf{F}_{\text{vol}} = J$ and $\bar{\mathbf{F}}$ contains only volume-preserving deformations, since $\det \bar{\mathbf{F}} = 1$. In a similar manner, the right Cauchy–Green tensor is split into

$$\begin{aligned} \mathbf{C} &= \mathbf{C}_{\text{vol}} \bar{\mathbf{C}}, \\ \text{with } \mathbf{C}_{\text{vol}} &:= J^{2/3} \mathbf{I} \quad \text{and} \quad \bar{\mathbf{C}} := J^{-2/3} \mathbf{C}, \end{aligned}$$

and the strain-energy function is postulated in a decoupled form, namely

$$\Psi(\mathbf{C}) = \Psi_{\text{vol}}(J) + \bar{\Psi}(\bar{\mathbf{C}}).$$

This gives rise to an additive split of the stress response into purely volumetric (\mathbf{S}_{vol}) and isochoric ($\bar{\mathbf{S}}$) contributions,

$$\begin{aligned} \mathbf{S} &= 2 \frac{\partial \Psi(\mathbf{C})}{\partial \mathbf{C}} = 2 \frac{\partial \Psi_{\text{vol}}(J)}{\partial J} \frac{\partial J}{\partial \mathbf{C}} + 2 \frac{\partial \bar{\Psi}(\bar{\mathbf{C}})}{\partial \bar{\mathbf{C}}} \\ &= J \frac{\partial \Psi_{\text{vol}}(J)}{\partial J} \mathbf{C}^{-1} + 2 \frac{\partial \bar{\Psi}(\bar{\mathbf{C}})}{\partial \bar{\mathbf{C}}} = \mathbf{S}_{\text{vol}} + \bar{\mathbf{S}}, \end{aligned} \quad (5.22)$$

using the relation for the Jacobian's derivative with respect to the right Cauchy–Green tensor \mathbf{C} (see, e.g. [276])

$$\frac{\partial J}{\partial \mathbf{C}} = \frac{\partial}{\partial \mathbf{C}} \det \mathbf{F} = 1/2 J \mathbf{C}^{-1}. \quad (5.23)$$

The strain-energy function $\Psi(\mathbf{C})$ is an isotropic tensor function, meaning, it is invariant under a rotation by an orthogonal tensor \mathbf{Q} , i.e., $\Psi(\mathbf{C}) = \Psi(\mathbf{Q} \mathbf{C} \mathbf{Q}^\top)$ and can therefore be expressed in terms of the principal invariants of its argument \mathbf{C} (see, e.g., [276] and [321, 322] for a proof of the so-called representation theorem for invariants). Strain-energy functions $\Psi(\mathbf{C})$ are thus frequently formulated in the principal invariants of \mathbf{C} , $I_i(\mathbf{C})$, $i = 1, 2, 3$, which are defined as

$$\begin{aligned} I_1 &:= \text{tr } \mathbf{C} = \mathbf{C} : \mathbf{I}, \\ I_2 &:= 1/2 \left[(\text{tr } \mathbf{C})^2 - \text{tr}(\mathbf{C}^2) \right] = 1/2 (I_1^2 - \mathbf{C} : \mathbf{C}) \\ I_3 &:= \det \mathbf{C}, \end{aligned} \quad (5.24)$$

such that derivatives with respect to \mathbf{C} are given as [276],

$$\frac{\partial I_1}{\partial \mathbf{C}} = \mathbf{I}, \quad \frac{\partial I_2}{\partial \mathbf{C}} = I_1 \mathbf{I} - \mathbf{C}, \quad \frac{\partial I_3}{\partial \mathbf{C}} = I_3 \mathbf{C}^{-1}. \quad (5.25)$$

Analogously, we can define the invariants of $\bar{\mathbf{C}}$ for convenient definitions of strain-energy functions of nearly incompressible hyperelastic continua as

$$\bar{I}_1 := \text{tr } \bar{\mathbf{C}} = J^{-2/3} I_1 \quad (5.26)$$

$$\bar{I}_2 := 1/2 \left[(\text{tr } \bar{\mathbf{C}})^2 - \text{tr}(\bar{\mathbf{C}}^2) \right] = J^{-4/3} I_2$$

$$\bar{I}_3 := \det \bar{\mathbf{C}} \equiv 1, \quad (5.27)$$

such that the derivatives with respect to \mathbf{C} are easily computed using (5.23), (5.25) and the chain rule, yielding

$$\begin{aligned} \frac{\partial \bar{I}_1}{\partial \mathbf{C}} &= J^{-2/3} \frac{\partial I_1}{\partial \mathbf{C}} + I_1 \frac{\partial}{\partial \mathbf{C}} J^{-2/3} = J^{-2/3} \left(\mathbf{I} - 1/3 I_1 \mathbf{C}^{-1} \right) \\ \frac{\partial \bar{I}_2}{\partial \mathbf{C}} &= J^{-4/3} \frac{\partial I_2}{\partial \mathbf{C}} + I_2 \frac{\partial}{\partial \mathbf{C}} J^{-4/3} = J^{-4/3} \left(I_1 \mathbf{I} - \mathbf{C} - 2/3 I_2 \mathbf{C}^{-1} \right) \\ \frac{\partial \bar{I}_3}{\partial \mathbf{C}} &= \mathbf{0}. \end{aligned} \quad (5.28)$$

Now, we are finally in position to introduce the strain-energy functions corresponding to the neo-Hookean ($\bar{\Psi}_{\text{NH}}$, [324]) and Mooney–Rivlin ($\bar{\Psi}_{\text{MR}}$, [325, 326]) material models for incompressible materials as

$$\bar{\Psi}_{\text{NH}} := c_{10}/2 \left(\bar{I}_1 - 3 \right), \quad (5.29)$$

$$\bar{\Psi}_{\text{MR}} := c_{10}/2 \left(\bar{I}_1 - 3 \right) + c_{01}/2 \left(\bar{I}_2 - 3 \right), \quad (5.30)$$

such that the corresponding second Piola–Kirchhoff stress tensors are given as

$$\bar{\mathbf{S}}_{\text{NH}} := 2 \frac{\partial \bar{\Psi}_{\text{NH}}}{\partial \mathbf{C}} = c_{10} \frac{\partial \bar{I}_1}{\partial \mathbf{C}}, \quad (5.31)$$

$$\bar{\mathbf{S}}_{\text{MR}} := 2 \frac{\partial \bar{\Psi}_{\text{MR}}}{\partial \mathbf{C}} = c_{10} \frac{\partial \bar{I}_1}{\partial \mathbf{C}} + c_{01} \frac{\partial \bar{I}_2}{\partial \mathbf{C}}, \quad (5.32)$$

with constants c_{01} and c_{10} relating to the shear modulus as $\mu = c_{10}$ for the neo-Hookean model and $\mu = c_{10} + c_{01}$ for the Mooney–Rivlin model, respectively. For the volumetric contribution to the strain-energy function $\Psi_{\text{vol}}(J)$, almost equally as many different representations exist, popular choices being

$$\Psi_{\text{vol}} := \kappa_b/2 (J - 1)^2, \quad \mathbf{S}_{\text{vol}} := \kappa_b (J - 1) J \mathbf{C}^{-1}, \quad (5.33)$$

$$\Psi_{\text{vol}} := \kappa_b/4 (J^2 - 1 - 2 \ln J), \quad \mathbf{S}_{\text{vol}} := \kappa_b/2 (J^2 - 1) \mathbf{C}^{-1}, \quad (5.34)$$

$$\Psi_{\text{vol}} := \kappa_b/2 \ln(J)^2, \quad \mathbf{S}_{\text{vol}} := \kappa_b \ln(J) \mathbf{C}^{-1}, \quad (5.35)$$

where merely (5.23) and the chain rule are used to derive \mathbf{S}_{vol} (see, e.g., [276, 327–329] and references therein).

5.2 Tissue models accounting for fibre-reinforcement

In the context of cardiovascular solid mechanics and FSI, the load-bearing capacity of collagen fibres plays a crucial role [330–333]. These fibres reinforce the ground material, introducing strong anisotropic effects, significantly stiffening the arterial tissue once an increased stretch is reached. To account for such a complex microstructure in a continuum-mechanical approach, Holzapfel et al. [334] introduced

$$\bar{\Psi}_{\text{fib}} := \frac{k_1}{2k_2} \sum_{i=4,6} \left\{ \exp \left[k_2 \left(\bar{I}_i - 1 \right)^2 \right] - 1 \right\}, \quad (5.36)$$

where k_1 is a stress-like parameter, k_2 is a dimensionless parameter and two additional invariants \bar{I}_i , $i = 4, 6$, defined as

$$\bar{I}_i := J^{-2/3} I_i, \quad \text{with } I_i := \mathbf{C} : \mathbf{A}_i, \quad (5.37)$$

are employed to capture the contribution of two fibre families. The directions of the fibre families are incorporated via the structural tensors \mathbf{A}_i , $i = 4, 6$, given by

$$\mathbf{A}_i := \mathbf{m}_i \otimes \mathbf{m}_i, \quad (5.38)$$

where \mathbf{m}_i denote given mean fibre directions in $\hat{\Omega}$. Based on the circumferential vessel direction \mathbf{e}_1 and the longitudinal direction \mathbf{e}_2 (to be constructed), one can then conveniently define mean fibre directions \mathbf{m}_4 and \mathbf{m}_6 as

$$\mathbf{m}_4 := \frac{\mathbf{e}_1 + \mathbf{e}_2 \tan \alpha_c}{\|\mathbf{e}_1 + \mathbf{e}_2 \tan \alpha_c\|}, \quad \mathbf{m}_6 := \frac{\mathbf{e}_1 - \mathbf{e}_2 \tan \alpha_c}{\|\mathbf{e}_1 - \mathbf{e}_2 \tan \alpha_c\|}, \quad (5.39)$$

for symmetric fibre reinforcements inclined by an angle α_c from the vessel's circumferential direction. For now, we simply assume \mathbf{e}_1 , \mathbf{e}_2 and α_c given, and postpone a more in-depth discussion on how to construct these vector fields appropriately to Sec. 6.1. The strain-energy function (5.36) has been extended to account for (non-symmetric) fibre dispersion [335, 336], exclusion of compressed fibres [337] and fibre degradation [338]. To limit complexity, we will herein focus on the fibre contribution as formulated by Gasser et al. [335], reading

$$\bar{\Psi}_{\text{fib}} := \frac{k_1}{2k_2} \sum_{i=4,6} \left(\exp \left\{ k_2 \left[\kappa_c \bar{I}_1 + (1 - 3\kappa_c) \bar{I}_i - 1 \right]^2 \right\} - 1 \right), \quad (5.40)$$

with a fibre dispersion parameter κ_c . This form corresponds to isotropic dispersion for $\kappa_c = 1/3$ and reduces to the zero dispersion case as shown in Eqn. (5.36) for $\kappa_c = 0$. The second Piola–Kirchhoff stress tensor is then given as [335]

$$\bar{\mathbf{S}}_{\text{fib}} := 2 \frac{\partial \bar{\Psi}_{\text{fib}}}{\partial \mathbf{C}} = \frac{k_1}{k_2} \sum_{i=4,6} \left[\frac{\partial}{\partial \mathbf{C}} \exp \left(k_2 G_i^2 \right) \right] = \sum_{i=4,6} \left[2k_1 G_i \exp \left(k_2 G_i^2 \right) \frac{\partial G_i}{\partial \mathbf{C}} \right], \quad (5.41)$$

where G_i , $i = 4, 6$ is introduced for convenience [268],

$$G_i := J^{-2/3} \left[\kappa_c I_1 + (1 - 3\kappa_c) I_i \right] - 1, \quad i = 4, 6,$$

such that we have

$$\begin{aligned} \frac{\partial G_i}{\partial \mathbf{C}} &= J^{-2/3} \left[\kappa_c \frac{\partial I_1}{\partial \mathbf{C}} + (1 - 3\kappa_c) \frac{\partial I_i}{\partial \mathbf{C}} \right] + \left[\kappa_c I_1 + (1 - 3\kappa_c) I_i \right] \frac{\partial J^{-2/3}}{\partial \mathbf{C}} \\ &= J^{-2/3} \left\{ \kappa_c \mathbf{I} + (1 - 3\kappa_c) \mathbf{A}_i - 1/3 \left[\kappa_c I_1 + (1 - 3\kappa_c) I_i \right] \mathbf{C}^{-1} \right\}, \end{aligned}$$

since $\frac{\partial}{\partial \mathbf{C}} I_i = \frac{\partial}{\partial \mathbf{C}} (\mathbf{C} : \mathbf{A}_i) = \mathbf{A}_i$ in analogy to (5.25).

Considering Eqn. (5.36) or (5.40), an important assumption is introduced by Holzapfel et al. [334], namely, that fibres under compression are *not* capable of withstanding any load and buckle *immediately*, such that the corresponding parts in the strain-energy function vanish. So, a fibre family is considered if $\mathbf{C} : \mathbf{A}_i \geq 1$ and omitted otherwise (see [339] for a detailed discussion). This assumption is not only physically meaningful,

but also needed to ensure material stability. For isotropic reinforcement, i.e., $\kappa_c = 0$, one can show that (5.36) is a convex potential, and for the form incorporating fibre dispersion (5.40), a proof of convexity (under simplifying assumptions) is given in [335]. A proof in the general case is to the best of the author's knowledge not available at this point in time. However, the present anisotropic constitutive model can be seen as a placeholder for a variety of material models and may thus be exchanged with low effort given the partitioned design of the final FSI scheme.

In contrast to the original form of the HGO model incorporating dispersed fibres as given in Eqn. (5.40) [335], modifications have been proposed incorporating \mathbf{C} instead of $\overline{\mathbf{C}}$ for the anisotropic fibre contributions to suppress unphysical results [340, 341]. We adopt this approach for the aortic dissection case presented in Sec. 8.7, but keep the general form of the stress tensors for the following derivations, as only incorporating the non-split \mathbf{C} and its invariants reduces the complexity of what follows and is thus the more general case.

The individual contributions \mathbf{S}_{vol} , $\overline{\mathbf{S}}_{\text{NH}}$ or $\overline{\mathbf{S}}_{\text{MR}}$ and optionally $\overline{\mathbf{S}}_{\text{fib}}$ in sum yield

$$\mathbf{S} = \mathbf{S}_{\text{vol}} + \overline{\mathbf{S}}_{\text{MR}} + \overline{\mathbf{S}}_{\text{fib}},$$

which is easily transformed via $\mathbf{P} = \mathbf{F} \mathbf{S}$ (5.8) to yield the first Piola–Kirchhoff stress tensor used in the momentum balance equation (5.9). Keeping the generic notation \mathbf{P} for the first Piola–Kirchhoff stress tensor irrespective of the material law considered, we can then proceed with the discretisation of the initial boundary value problem (5.9)–(5.13), to be presented next.

5.3 Discretisation in time

To discretise the momentum balance equation (5.9) in time, decompose the time interval $I_t = (0, T]$ from initial time $t = 0$ to end time $t = T$ into N_t possibly non-uniform time steps denoted by $\Delta t = t^{n+1} - t^n$, $n = 0, \dots, N_t$. Adopting standard notation, we abbreviate the material acceleration, velocity and displacement at time $t = t^{n+1}$ by

$$\ddot{\mathbf{d}}^{n+1} := \text{d}_t \mathbf{d}(\hat{\mathbf{x}}, t^{n+1}) \quad , \quad \dot{\mathbf{d}}^{n+1} := \text{d}_t \mathbf{d}(\hat{\mathbf{x}}, t^{n+1}) \quad , \quad \mathbf{d}^{n+1} := \mathbf{d}(\hat{\mathbf{x}}, t),$$

and employ generalised- α time integration [342], which evaluates acceleration terms at the generalised midpoint $t^{n+1-\alpha_m} := (1 - \alpha_m)t^{n+1} + \alpha_m t^n$, while forcing terms, displacement and velocity are evaluated at $t^{n+1-\alpha_f} := (1 - \alpha_f)t^{n+1} + \alpha_f t^n$. Thus, we insert the intermediate values at the generalised midpoints,

$$\begin{aligned} \ddot{\mathbf{d}}^{n+1-\alpha_m} &:= \alpha'_m \ddot{\mathbf{d}}^{n+1} + \alpha_m \ddot{\mathbf{d}}^n, \\ \mathbf{d}^{n+1-\alpha_f} &:= \alpha'_f \mathbf{d}^{n+1} + \alpha_f \mathbf{d}^n, \end{aligned}$$

with shorthand notations $\alpha'_m = 1 - \alpha_m$ and $\alpha'_f = 1 - \alpha_f$, into the structure's momentum balance equation as given in (5.9) to obtain

$$\begin{aligned} \rho \ddot{\mathbf{d}}^{n+1-\alpha_m} - \hat{\nabla} \cdot \mathbf{P}(\mathbf{d}^{n+1-\alpha_f}) &= \mathbf{b}^{n+1-\alpha_f}, \\ \rho (\alpha'_m \ddot{\mathbf{d}}^{n+1} + \alpha_m \ddot{\mathbf{d}}^n) - \hat{\nabla} \cdot \mathbf{P}(\alpha'_f \mathbf{d}^{n+1} + \alpha_f \mathbf{d}^n) &= \mathbf{b}^{n+1-\alpha_f} \end{aligned} \quad (5.42)$$

This specific time integration scheme corresponds to the generalised midpoint rule, which uses $\mathbf{P}^{n+1-\alpha_f} \approx \mathbf{P}(\mathbf{d}^{n+1-\alpha_f})$ for the nonlinear stress term. For ease of presentation and implementation—especially relevant in the FSI setting (see Ch. 7)—we instead employ the generalised trapezoidal rule, that is,

$$\mathbf{P}^{n+1-\alpha_f} \approx \alpha'_f \mathbf{P}^{n+1} + \alpha_f \mathbf{P}^n,$$

with $\mathbf{P}^{n+1} = \mathbf{P}(\mathbf{d}^{n+1})$ and likewise $\mathbf{P}^n = \mathbf{P}(\mathbf{d}^n)$, achieving same orders of accuracy in time [343], to end up with a time-discrete momentum balance equation of the form

$$\rho (\alpha'_m \ddot{\mathbf{d}}^{n+1} + \alpha_m \ddot{\mathbf{d}}^n) - \alpha'_f \hat{\nabla} \cdot \mathbf{P}^{n+1} - \alpha_f \hat{\nabla} \cdot \mathbf{P}^n = \alpha'_f \mathbf{b}^{n+1} + \alpha_f \mathbf{b}^n. \quad (5.43)$$

Further, we employ Newmark formulae [344], linking the material acceleration, velocity and displacement via

$$\begin{aligned} 1/\Delta t (\mathbf{d}^{n+1} - \mathbf{d}^n) &\approx \dot{\mathbf{d}}^n + \Delta t \beta \ddot{\mathbf{d}}^{n+1} + \Delta t (1/2 - \beta) \ddot{\mathbf{d}}^n \\ 1/\Delta t (\dot{\mathbf{d}}^{n+1} - \dot{\mathbf{d}}^n) &\approx \gamma \ddot{\mathbf{d}}^{n+1} + (1 - \gamma) \ddot{\mathbf{d}}^n, \end{aligned}$$

with time integration parameters β and γ . This can be rewritten to obtain relations for material acceleration and velocity at $t = t^{n+1}$, implicit in terms of \mathbf{d}^{n+1} only:

$$\ddot{\mathbf{d}}^{n+1} \approx \frac{1}{\beta \Delta t^2} (\mathbf{d}^{n+1} - \mathbf{d}^n) - \frac{1}{\beta \Delta t} \dot{\mathbf{d}}^n + \left(1 - \frac{1}{2\beta}\right) \ddot{\mathbf{d}}^n, \quad (5.44)$$

$$\dot{\mathbf{d}}^{n+1} \approx \frac{\gamma}{\Delta t \beta} (\mathbf{d}^{n+1} - \mathbf{d}^n) + \left(1 - \frac{\gamma}{\beta}\right) \dot{\mathbf{d}}^n + \Delta t \left(1 - \frac{\gamma}{2\beta}\right) \ddot{\mathbf{d}}^n. \quad (5.45)$$

Specific choices of the time integration parameters γ , β , α_m and α_f determine accuracy and stability properties of the resulting scheme. As Chung and Hulbert [342] have shown, second-order accuracy in the displacement and the velocity, first-order accuracy in the acceleration and unconditional stability are achieved for linear problems using

$$\gamma = \frac{1}{2} - \alpha_m + \alpha_f \quad , \quad \beta = \frac{1}{4} (1 - \alpha_m + \alpha_f)^2.$$

For nonlinear problems, the analysis was extended by Erlicher et al. [343], proving second-order accuracy in the displacement and velocity, first-order accuracy in acceleration and energy stability in the high-frequency range depending on the algorithmic parameter ρ_∞ . This parameter ρ_∞ denotes the spectral radius in the high frequency limit and is utilised to set appropriate α_m and α_f according to Tab. 5.1, resulting in a family of

Tab. 5.1: Parameters α_f and α_m in the generalised- α time integration scheme to obtain the N- β [344], HHT- α [345], WBZ- α [346] or CH- α [342] methods with ρ_∞ being the spectral radius in the high frequency limit. Zero numerical dissipation is added for $\rho_\infty = 1$ and maximum dissipation is achieved with $\rho_\infty = \rho_\infty^{\min}$, $\rho_\infty^{\min} = 1/2$ for HHT- α and 0 else.

	N- β	HHT- α	WBZ- α	CH- α
$\alpha_m(\rho_\infty)$	0	0	$\frac{\rho_\infty - 1}{1 + \rho_\infty}$	$\frac{2\rho_\infty - 1}{1 + \rho_\infty}$
$\alpha_f(\rho_\infty)$	0	$\frac{1 - \rho_\infty}{1 + \rho_\infty}$	0	$\frac{\rho_\infty}{1 + \rho_\infty}$
$\alpha_m(\rho_\infty = 1)$	0	0	0	$1/2$
$\alpha_f(\rho_\infty = 1)$	0	0	0	$1/2$
$\alpha_m(\rho_\infty^{\min})$	0	0	-1	-1
$\alpha_f(\rho_\infty^{\min})$	0	$1/3$	0	0

time integration schemes with specific choices being: (i) the Newmark- β (N- β) method by Newmark [344], (ii) the HHT- α method introduced by Hilber et al. [345], (iii) the WBZ- α method by Wood et al. [346], and (iv) the CH- α method by Chung and Hulbert [342]. Suitable ranges are $\rho_\infty \in [\rho_\infty^{\min}, 1]$, where $\rho_\infty^{\min} = 1/2$ for the HHT- α time integration scheme and $\rho_\infty^{\min} = 0$ for all other methods. Using $\rho_\infty = 1$, no dissipation is introduced and N- β , WBZ- α and HHT- α schemes coincide. Using the maximum amount of numerical dissipation in the WBZ- α and CH- α schemes, $\rho_\infty^{\min} = 0$, both schemes yield identical time integration parameters as can be seen from Tab. 5.1. For this so-called asymptotic annihilation case, the high-frequency response is completely annihilated after a single time step [342]. Maximum numerical dissipation is obtained setting $\rho_\infty^{\min} = 1/2$ in the HHT- α scheme, resulting in the time integration coefficients $\alpha_f = 1/3$ and $\alpha_m = 0$.

The generalised- α time integration scheme was just recently extended to third- and higher-order accuracy for first-order ordinary differential equations [347], hyperbolic problems [348] and explicit structural dynamics integrators [349]. Herein, however, the discussion is limited to second-order accurate, single-step schemes, which allow easy control of the time step size based on explicit-implicit predictor-corrector approaches (see, e.g. [151, 350]), or error estimation by Taylor series expansion [137].

So, the time-discrete weak form of linear momentum balance in the displacement \mathbf{d}^{n+1} is found by inserting (5.44) into (5.43), yielding

$$\begin{aligned}
 0 = \rho \alpha'_m \left[\frac{1}{\beta \Delta t^2} (\mathbf{d}^{n+1} - \mathbf{d}^n) - \frac{1}{\beta \Delta t} \dot{\mathbf{d}}^n + \left(1 - \frac{1}{2\beta} \right) \ddot{\mathbf{d}}^n \right] + \rho \alpha_m \ddot{\mathbf{d}}^n \\
 - \alpha'_f (\hat{\nabla} \cdot \mathbf{P}^{n+1} + \mathbf{b}^{n+1}) - \alpha_f (\hat{\nabla} \cdot \mathbf{P}^n + \mathbf{b}^n).
 \end{aligned} \tag{5.46}$$

5.4 Fully discrete, linearised problem

Based on the time-discrete problem (5.46) and introducing $[X(\hat{\Omega})]^d \subset [H^1(\hat{\Omega})]^d$, a weak form is found multiplying with $\boldsymbol{\varphi} \in [X(\hat{\Omega})]^d$ with $\boldsymbol{\varphi}|_{\hat{\Gamma}_D} = \mathbf{0}$, giving

$$0 = \langle \rho \boldsymbol{\varphi}, \alpha'_m \ddot{\mathbf{d}}^{n+1} (\mathbf{d}^{n+1}) + \alpha_m \ddot{\mathbf{d}}^n \rangle - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{b}^{n+1} + \alpha_f \mathbf{b}^n \rangle \\ - \langle \boldsymbol{\varphi}, \alpha'_f \hat{\nabla} \cdot \mathbf{P}^{n+1} + \alpha_f \hat{\nabla} \cdot \mathbf{P}^n \rangle,$$

where we denote by $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle_{\hat{\Gamma}_N}$ the standard L^2 scalar products in $\hat{\Omega}$ and $\hat{\Gamma}_N$ similar to previous sections. Applying the divergence theorem, one can then obtain a suitable weak form in the structure's displacements \mathbf{d} only, rewriting the weak form as

$$0 = \langle \rho \boldsymbol{\varphi}, \alpha'_m \ddot{\mathbf{d}}^{n+1} (\mathbf{d}^{n+1}) + \alpha_m \ddot{\mathbf{d}}^n \rangle - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{b}^{n+1} + \alpha_f \mathbf{b}^n \rangle \\ + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}^{n+1} + \alpha_f \mathbf{P}^n \rangle - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{P}^{n+1} \hat{\mathbf{n}} + \alpha_f \mathbf{P}^n \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_N},$$

where $\boldsymbol{\varphi}|_{\Gamma_D} = \mathbf{0}$ is used, since the Dirichlet boundary conditions on \mathbf{d}^{n+1} (5.12) or their homogeneous counterparts are incorporated in the ansatz space as shall soon become clear. Plugging in the Neumann data, i.e., first Piola–Kirchhoff traction vectors $\hat{\mathbf{t}}$ at times t^{n+1} and t^n given on $\hat{\Gamma}_N$, we then finally obtain a semilinear form $R(\mathbf{d})(\boldsymbol{\varphi})$ resembling the residual as

$$R(\mathbf{d})(\boldsymbol{\varphi}) := \langle \rho \alpha'_m \boldsymbol{\varphi}, \frac{1}{\beta \Delta t^2} (\mathbf{d} - \mathbf{d}^n) - \frac{1}{\beta \Delta t} \dot{\mathbf{d}}^n + \left(1 - \frac{1}{2\beta} + \frac{\alpha_m}{\alpha'_m}\right) \ddot{\mathbf{d}}^n \rangle \\ - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{b}^{n+1} + \alpha_f \mathbf{b}^n \rangle + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}(\mathbf{d}) + \alpha_f \mathbf{P}^n \rangle \\ - \langle \boldsymbol{\varphi}, \alpha'_f \hat{\mathbf{t}}^{n+1} + \alpha_f \hat{\mathbf{t}}^n \rangle_{\hat{\Gamma}_N}. \quad (5.47)$$

Based on this variational form of the residual, we can thus define the problem seeking $\mathbf{d}^{n+1} \in [X(\hat{\Omega})]^d$ with inhomogeneous Dirichlet conditions $\mathbf{d}^{n+1}|_{\hat{\Gamma}_D} = \mathbf{g}^{n+1}$, such that

$$R(\mathbf{d}^{n+1})(\boldsymbol{\varphi}) = 0 \quad (5.48)$$

holds for all $\boldsymbol{\varphi} \in [X(\hat{\Omega})]^d$ with $\boldsymbol{\varphi}|_{\hat{\Gamma}_D} = \mathbf{0}$. Note here, however, that $R(\mathbf{d})(\boldsymbol{\varphi})$ as given in Eqn. (5.47) and hence Eqn. (5.48) is potentially nonlinear in the structure's displacement \mathbf{d}^{n+1} depending on the definition of the first Piola–Kirchhoff stress tensor $\mathbf{P}(\mathbf{d})$. To this end, we employ Newton's method to solve the nonlinear quasi-stationary problem at each time step t^{n+1} . Newton's method is particularly prominent for solving nonlinear systems arising from solid mechanics problems due to its locally quadratic convergence rate. In the present context, that is, instationary problems with small time steps, the standard Newton method is thus expected to perform well. This is especially true when the structure is to be coupled to flow problems within FSI coupling algorithms. One of the main reasons for fast convergence of Newton's method in such a scenario is the initial guess—usually extrapolated incorporating data from the last time steps—lies in the vicinity of the nonlinear problem's solution. Moreover, the Newton method can still

be used in the case of linear elasticity, and simply converges in a single step given low enough tolerances in the linear solver.

Newton's method consists of an iterative process repeatedly updating the k -th approximation of \mathbf{d}^{n+1} , denoted by $\mathbf{d}^{n+1,k}$. In several steps starting from an initial guess $\mathbf{d}^{n+1,0}$, the next iterate $\mathbf{d}^{n+1,k+1}$ is obtained via a simple vector update of the form

$$\mathbf{d}^{n+1,k+1} = \mathbf{d}^{n+1,k} + \delta \mathbf{d}^k, \quad (5.49)$$

until suitable absolute and/or relative convergence criteria are reached. Here, $\delta \mathbf{d}^k$ is the increment computed by solving the linear defect-correction problem

$$\mathbf{J}(\mathbf{d}^{n+1,k})(\boldsymbol{\varphi}, \delta \mathbf{d}^k) = -\mathbf{R}(\mathbf{d}^{n+1,k})(\boldsymbol{\varphi}), \quad (5.50)$$

with $\mathbf{J}(\mathbf{d})(\boldsymbol{\varphi}, \delta \mathbf{d})$ being the directional or Gâteaux derivative of the (potentially) non-linear residual $\mathbf{R}(\mathbf{d})(\boldsymbol{\varphi})$ with respect to the argument \mathbf{d} , which is defined as

$$\begin{aligned} \mathbf{J}(\mathbf{d})(\boldsymbol{\varphi}, \delta \mathbf{d}) &:= \mathbf{D}_d \mathbf{R}(\mathbf{d} + \epsilon \delta \mathbf{d})(\boldsymbol{\varphi})|_{\epsilon=0} \\ &= \lim_{\epsilon \rightarrow 0} \left\{ \frac{\mathbf{R}(\mathbf{d} + \epsilon \delta \mathbf{d})(\boldsymbol{\varphi}) - \mathbf{R}(\mathbf{d})(\boldsymbol{\varphi})}{\epsilon} \right\}. \end{aligned}$$

Naturally, the test function $\boldsymbol{\varphi}$, known boundary and body force terms and terms involving the past time step's displacements, velocities or accelerations present in $\mathbf{R}(\mathbf{d})(\boldsymbol{\varphi})$ do not depend on the solution \mathbf{d}^{n+1} at t^{n+1} (simply \mathbf{d} in Eqn. (5.47)), such that the Jacobian $\mathbf{J}(\mathbf{d})(\boldsymbol{\varphi}, \delta \mathbf{d})$ is given by

$$\mathbf{D}_d \langle \rho \alpha'_m \boldsymbol{\varphi}, \frac{1}{\beta \Delta t^2} \mathbf{d} \rangle = \langle \rho \alpha'_m \boldsymbol{\varphi}, \frac{1}{\beta \Delta t^2} \delta \mathbf{d} \rangle,$$

stemming from the acceleration term, and the Jacobian of the remaining stress term at t^{n+1} , which consists only of

$$\mathbf{D}_d \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}(\mathbf{d}) \rangle = \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{D}_d [\mathbf{P}(\mathbf{d})] \rangle,$$

since for now, we consider *given* first Piola–Kirchhoff tractions prescribed on $\hat{\Gamma}_N$. Thus, $\mathbf{D}_d [\mathbf{P}(\mathbf{d})]$ remains to be defined for each constitutive relation introduced and is derived based on the standard relations (see, e.g., [276, 277, 319])

$$D_d [\mathbf{F}(\mathbf{d})] = \hat{\nabla} \delta \mathbf{d}, \quad (5.51)$$

$$D_d [\mathbf{F}^\top(\mathbf{d})] = \{D_d [\mathbf{F}(\mathbf{d})]\}^\top = \hat{\nabla}^\top \delta \mathbf{d}, \quad (5.52)$$

$$D_d [J(\mathbf{d})] = D_d [\det \mathbf{F}(\mathbf{d})] = J \operatorname{tr} (\mathbf{F}^{-1} \hat{\nabla} \delta \mathbf{d}), \quad (5.53)$$

$$D_d [\mathbf{F}^{-1}(\mathbf{d})] = -\mathbf{F}^{-1} (\hat{\nabla} \delta \mathbf{d}) \mathbf{F}^{-1}, \quad (5.54)$$

$$D_d [\mathbf{F}^{-\top}(\mathbf{d})] = \{D_d [\mathbf{F}^{-1}(\mathbf{d})]\}^\top = -\mathbf{F}^{-\top} (\hat{\nabla}^\top \delta \mathbf{d}) \mathbf{F}^{-\top}, \quad (5.55)$$

such that we can directly determine all needed directional derivatives for each of the considered material models. In linear elasticity, we have (5.19)

$$\mathbf{P} = \lambda \mathbf{I} \hat{\nabla} \cdot \mathbf{d} + \mu (\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d}),$$

such that the directional derivative corresponding to the stress term is

$$\begin{aligned} D_d [\mathbf{P}(\mathbf{d})] &= D_d [\lambda \mathbf{I} \hat{\nabla} \cdot \mathbf{d} + \mu (\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d})] \\ &= \lambda \mathbf{I} \hat{\nabla} \cdot \delta \mathbf{d} + \mu (\hat{\nabla} \delta \mathbf{d} + \hat{\nabla}^\top \delta \mathbf{d}). \end{aligned} \quad (5.56)$$

The nonlinear structure models considered within this work all involve the transformation $\mathbf{P} = \mathbf{F} \mathbf{S}$, which leads together with (5.51) to a directional derivative of

$$\begin{aligned} D_d [\mathbf{P}(\mathbf{d})] &= D_d [\mathbf{F}(\mathbf{d}) \mathbf{S}(\mathbf{d})] \\ &= D_d [\mathbf{F}(\mathbf{d})] \mathbf{S}(\mathbf{d}) + \mathbf{F}(\mathbf{d}) D_d [\mathbf{S}(\mathbf{d})] \\ &= (\hat{\nabla} \delta \mathbf{d}) \mathbf{S}(\mathbf{d}) + \mathbf{F}(\mathbf{d}) D_d [\mathbf{S}(\mathbf{d})], \end{aligned}$$

where now merely the second term containing $D_d \mathbf{S}$ depends on the specific constitutive law. The Saint Venant–Kirchhoff model (5.20) leads to

$$D_d [\mathbf{S}(\mathbf{d})] = D_d [\lambda \mathbf{I} \operatorname{tr} \mathbf{E} + 2\mu \mathbf{E}] = \lambda \mathbf{I} \operatorname{tr} (D_d \mathbf{E}) + 2\mu D_d \mathbf{E}, \quad (5.57)$$

$$\text{with } D_d [\mathbf{E}(\mathbf{d})] = D_d [1/2 (\mathbf{F}^\top \mathbf{F} - \mathbf{I})] = 1/2 [D_d (\mathbf{F}^\top) \mathbf{F} + \mathbf{F}^\top D_d \mathbf{F}].$$

For incompressible continua applying a multiplicative split of \mathbf{F} and \mathbf{C} into volumetric and isochoric parts and assuming a decoupled strain-energy function leads to (5.22)

$$\mathbf{S} = \mathbf{S}_{\text{vol}}(J) + \overline{\mathbf{S}}(\mathbf{C}).$$

For the different definitions of \mathbf{S}_{vol} given in Eqns. (5.33)–(5.35) we obtain

$$D_d [\kappa_b (J - 1) J \mathbf{C}^{-1}] = \kappa_b [(2J - 1) D_d (J) \mathbf{C}^{-1} + (J - 1) J D_d (\mathbf{C}^{-1})], \quad (5.58)$$

$$D_d [\kappa_b / 2 (J^2 - 1) \mathbf{C}^{-1}] = \kappa_b [J D_d (J) \mathbf{C}^{-1} + 1/2 (J^2 - 1) D_d (\mathbf{C}^{-1})], \quad (5.59)$$

$$D_d [\kappa_b \ln(J) \mathbf{C}^{-1}] = \kappa_b [J^{-1} D_d (J) \mathbf{C}^{-1} + \ln(J) D_d (\mathbf{C}^{-1})], \quad (5.60)$$

$$D_d [\mathbf{C}^{-1}(\mathbf{d})] = D_d (\mathbf{F}^{-1} \mathbf{F}^{-\top}) = D_d (\mathbf{F}^{-1}) \mathbf{F}^{-\top} + \mathbf{F}^{-1} D_d (\mathbf{F}^{-\top}). \quad (5.61)$$

Now, all that is left is to determine $D_d \bar{\mathbf{S}}$, which might be composed of $\bar{\mathbf{S}}_{\text{NH}}$ (5.31) or, alternatively, $\bar{\mathbf{S}}_{\text{MR}}$ (5.32) and might contain fibre contributions via $\bar{\mathbf{S}}_{\text{fib}}$ (5.41). Since neo-Hookean and Mooney–Rivlin models can be written as

$$\begin{aligned} \bar{\mathbf{S}}_{\text{NH}} &= c_{10} J^{-2/3} (\mathbf{I} - 1/3 I_1 \mathbf{C}^{-1}) \\ \text{and } \bar{\mathbf{S}}_{\text{MR}} &= \bar{\mathbf{S}}_{\text{NH}} + c_{01} J^{-4/3} (I_1 \mathbf{I} - \mathbf{C} - 2/3 I_2 \mathbf{C}^{-1}), \end{aligned}$$

it suffices to further specify $D_d \mathbf{C}$, $D_d I_1$ and $D_d I_2$,

$$D_d [\mathbf{C}(\mathbf{d})] = D_d (\mathbf{F}^\top \mathbf{F}) = D_d (\mathbf{F}^\top) \mathbf{F} + \mathbf{F}^\top D_d \mathbf{F}, \quad (5.62)$$

$$D_d [I_1(\mathbf{d})] = D_d \text{tr} \mathbf{C} = \text{tr}(D_d \mathbf{C}), \quad (5.63)$$

$$D_d [I_2(\mathbf{d})] = D_d [1/2 (I_1^2 - \mathbf{C} : \mathbf{C})] = I_1 D_d I_1 - \mathbf{C} : D_d \mathbf{C}. \quad (5.64)$$

Then, applying the chain rule, one directly obtains

$$D_d [\bar{\mathbf{S}}_{\text{NH}}(\mathbf{d})] = c_{10} \left[(\mathbf{I} - 1/3 I_1 \mathbf{C}^{-1}) D_d (J^{-2/3}) + J^{-2/3} D_d (\mathbf{I} - 1/3 I_1 \mathbf{C}^{-1}) \right], \quad (5.65)$$

$$\begin{aligned} D_d (J^{-2/3}) &= -\frac{2}{3} J^{-5/3} D_d J, \\ D_d (\mathbf{I} - 1/3 I_1 \mathbf{C}^{-1}) &= -1/3 [\mathbf{C}^{-1} D_d I_1 + I_1 D_d (\mathbf{C}^{-1})], \end{aligned}$$

and

$$\begin{aligned} D_d [\bar{\mathbf{S}}_{\text{MR}}(\mathbf{d})] &= D_d \bar{\mathbf{S}}_{\text{NH}} + c_{01} (I_1 \mathbf{I} - \mathbf{C} - 2/3 I_2 \mathbf{C}^{-1}) D_d (J^{-4/3}) \\ &\quad + c_{01} J^{-4/3} D_d (I_1 \mathbf{I} - \mathbf{C} - 2/3 I_2 \mathbf{C}^{-1}), \end{aligned} \quad (5.66)$$

with

$$\begin{aligned} D_d (J^{-4/3}) &= -\frac{4}{3} J^{-7/3} D_d J, \\ D_d (I_1 \mathbf{I} - \mathbf{C} - 2/3 I_2 \mathbf{C}^{-1}) &= I D_d I_1 - D_d \mathbf{C} - 2/3 [\mathbf{C}^{-1} D_d I_2 + I_2 D_d (\mathbf{C}^{-1})]. \end{aligned}$$

For fibre-reinforced continua, $\bar{\mathbf{S}}$ is augmented by (5.41)

$$\bar{\mathbf{S}}_{\text{fib}} = \sum_{i=4,6} \left[2k_1 G_i \exp(k_2 G_i^2) \frac{\partial G_i}{\partial \mathbf{C}} \right],$$

and the chain rule results in

$$D_d [\bar{\mathbf{S}}_{\text{fib}}(\mathbf{d})] = \sum_{i=4,6} \left\{ D_d [2k_1 G_i \exp(k_2 G_i^2)] \frac{\partial G_i}{\partial \mathbf{C}} + [2k_1 G_i \exp(k_2 G_i^2)] D_d \left(\frac{\partial G_i}{\partial \mathbf{C}} \right) \right\}, \quad (5.67)$$

$$\begin{aligned}
\mathbf{D}_d \left[2k_1 G_i \exp \left(k_2 G_i^2 \right) \right] &= 2k_1 \exp \left(k_2 G_i^2 \right) \left(1 + 2k_2 G_i^2 \right) \mathbf{D}_d G_i, \\
\mathbf{D}_d \left[G_i(\mathbf{d}) \right] &= \mathbf{D}_d \left(J^{-2/3} \right) \left[\kappa_c I_1 + (1 - 3\kappa_c) I_i \right] + J^{-2/3} \left[\kappa_c \mathbf{D}_d I_1 + (1 - 3\kappa_c) \mathbf{D}_d I_i \right], \\
\mathbf{D}_d \left[I_i(\mathbf{d}) \right] &= \mathbf{D}_d \left(\mathbf{C} : \mathbf{A}_i \right) = \mathbf{A}_i : \mathbf{D}_d \mathbf{C}, \\
\mathbf{D}_d \left(\frac{\partial G_i}{\partial \mathbf{C}} \right) &= \mathbf{D}_d \left(J^{-2/3} \right) \left\{ \kappa_c \mathbf{I} + (1 - 3\kappa_c) \mathbf{A}_i - 1/3 \left[\kappa_c I_1 + (1 - 3\kappa_c) I_i \right] \mathbf{C}^{-1} \right\} \\
&\quad - \frac{1}{3} J^{-2/3} \left\{ \left[\kappa_c \mathbf{D}_d I_1 + (1 - 3\kappa_c) \mathbf{D}_d I_i \right] \mathbf{C}^{-1} + \left[\kappa_c I_1 + (1 - 3\kappa_c) I_i \right] \mathbf{D}_d \left(\mathbf{C}^{-1} \right) \right\}
\end{aligned}$$

finally completes the Jacobian. Thus, the problem is to find the k -th Newton update $\delta \mathbf{d}^k \in [X(\hat{\Omega})]^d$ with $\delta \mathbf{d}^k|_{\hat{\Gamma}_D} = \mathbf{0}$, such hat

$$\langle \rho \alpha'_m \boldsymbol{\varphi}, \frac{1}{\beta \Delta t^2} \delta \mathbf{d} \rangle + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{D}_d \mathbf{P} \left(\mathbf{d}^{n+1,k} \right) (\delta \mathbf{d}) \rangle = -\mathbf{R} \left(\mathbf{d}^{n+1,k} \right) (\boldsymbol{\varphi}), \quad (5.68)$$

for all $\boldsymbol{\varphi} \in [X(\hat{\Omega})]^d$ with $\boldsymbol{\varphi}|_{\hat{\Gamma}_D} = \mathbf{0}$. The solution $\delta \mathbf{d}^k$ to (5.68) is then directly used to update the last iterate via $\mathbf{d}^{n+1,k+1} = \mathbf{d}^{n+1,k} + \delta \mathbf{d}^k$ and the Dirichlet boundary conditions (5.12), i.e., $\mathbf{d}^{n+1}|_{\hat{\Gamma}_D} = \mathbf{g}^{n+1}$ are imposed on the initial guess $\mathbf{d}^{n+1,0}$ prior to the iterative process of Newton's method.

When incorporating the effects of dispersed fibres via the strain-energy functions (5.36) or (5.40) and excluding fibres whenever $\mathbf{C} : \mathbf{A}_i < 1$ (irrespective of the Flory split considered for the anisotropic contribution, see [339]), the Jacobian given in (5.67) is merely an approximation. This can lead to stagnation of the nonlinear solver, calling for globalisation strategies like Aitken's acceleration, an Armijo scheme or arc-length methods. However, employing the structural solver within an FSI scheme, time step sizes are most often limited by the fluid flow solver and in addition to that, the structural solver, i.e., the Newton method is executed once per FSI coupling iteration. Therefore, rather crude *relative* tolerances can be accepted in the linear and nonlinear solvers and in fact, one might even consider linearising the tension-compression switch, i.e., the displacement iterate used to ensure $\mathbf{C} : \mathbf{A}_i \geq 1$ by a suitable extrapolation in time. Within this work, we thus borrow ideas from Nobile et al. [351, 352], linearising the tension-compression switch with the initial guess, while in a conservative approach, the remainder of the nonlinear terms are adequately captured, reducing the residual by, e.g., a factor of 10^3 or 10^4 depending on the desired accuracy in the FSI coupling scheme. Depending on the problem at hand, i.e., nonlinearities, initial guess, mean needed FSI steps per time step and more, such an approach typically leads to 1 to 3 Newton steps per solver call.

As a noteworthy alternative to the stress term used in (5.68), which is based on $\langle \hat{\nabla} \boldsymbol{\varphi}, \mathbf{P} \rangle$, one can make use of the relation

$$\hat{\nabla} \boldsymbol{\varphi} : \mathbf{P} = \hat{\nabla} \boldsymbol{\varphi} : (\mathbf{F} \mathbf{S}) \equiv 1/2 \left(\hat{\nabla}^\top \boldsymbol{\varphi} \mathbf{F} + \mathbf{F}^\top \hat{\nabla} \boldsymbol{\varphi} \right) : \mathbf{S},$$

which can be shown by trivial component-wise calculation and exploiting symmetry of the second Piola–Kirchhoff stress tensor, $\mathbf{S} = \mathbf{S}^\top$.

The corresponding directional derivative can then be written in a form which separates geometric and material nonlinearities in a slightly different way,

$$\mathrm{D}_d \langle \hat{\nabla}^\top \boldsymbol{\varphi} \mathbf{F} + \mathbf{F}^\top \hat{\nabla} \boldsymbol{\varphi}, \mathbf{S} \rangle = \langle \hat{\nabla}^\top \boldsymbol{\varphi} \hat{\nabla} \delta \mathbf{d} + \hat{\nabla}^\top \delta \mathbf{d} \hat{\nabla} \boldsymbol{\varphi}, \mathbf{S} \rangle + \langle \hat{\nabla}^\top \boldsymbol{\varphi} \mathbf{F} + \mathbf{F}^\top \hat{\nabla} \boldsymbol{\varphi}, \mathrm{D}_d \mathbf{S} \rangle,$$

Nonetheless, the non-symmetric $\langle \hat{\nabla} \boldsymbol{\varphi}, \mathbf{P} \rangle$ is considered in all of the presented computations and is found rather uncritical, given that the mass matrix scaled by $\rho/\beta \Delta t^2$ dominates the linear system solved in each Newton step in most cases.

Regardless of the specific linearisation employed, the spatial discretisation may be carried out in a standard manner, approximating the domain $\hat{\Omega}$ by a decomposition $\hat{\Omega}_h$ into N_e shape-regular, non-overlapping Lagrangian finite elements $\hat{\Omega}_e$, i.e.,

$$\hat{\Omega} \approx \hat{\Omega}_h = \bigcup_{e=1}^{N_e} \hat{\Omega}_e.$$

The reference elements Ω_{ref} are isoparametrically mapped to the Lagrangian configuration $\hat{\Omega}_h$ in an element-wise manner via

$$\boldsymbol{\chi}_e(\boldsymbol{\xi}) := \sum_{i=1}^{N_{\text{ref}}} N_i(\boldsymbol{\xi}) \hat{\boldsymbol{x}}_i,$$

where N_{ref} is the number of nodes per reference element, $N_i(\boldsymbol{\xi})$ is the shape function of node i of the reference element and $\hat{\boldsymbol{x}}_i$ is the corresponding coordinate of the node in the Lagrangian reference frame. Standard continuous Lagrangian polynomials of order k are constructed on the reference element Ω_{ref} , spanning spaces $\mathbb{P}^k(\Omega_{\text{ref}})$ and $\mathbb{Q}^k(\Omega_{\text{ref}})$ on simplicial or tensor-product reference elements, respectively. Mapping this Lagrangian basis to $\hat{\Omega}$, one defines a C^0 -continuous function space $X_h(\hat{\Omega}_h) \subset H^1(\hat{\Omega}_h)$ via

$$\begin{aligned} X_h(\hat{\Omega}_h) &:= \left\{ q_h \in C^0(\hat{\Omega}_h) : q_h \circ \boldsymbol{\chi}_e(\boldsymbol{\xi})|_{\hat{\Omega}_e} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{P}^k(\Omega_{\text{ref}}) \quad \forall \hat{\Omega}_e \in \hat{\Omega}_h \right\}, \\ \text{or } X_h(\hat{\Omega}_h) &:= \left\{ q_h \in C^0(\hat{\Omega}_h) : q_h \circ \boldsymbol{\chi}_e(\boldsymbol{\xi})|_{\hat{\Omega}_e} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{Q}^k(\Omega_{\text{ref}}) \quad \forall \hat{\Omega}_e \in \hat{\Omega}_h \right\}. \end{aligned}$$

Plugging in nodal shape functions $\boldsymbol{\varphi}_h$ and $\delta \mathbf{d}_h \in [X_h(\hat{\Omega})]^d$ into the variational formulation of Newton's update step (5.68) yields the fully discrete problem to be solved in each iteration of Newton's method. Denoting the nodal DoF vectors of the sought new iterate, the last computed solution and the Newton update by $\underline{\mathbf{d}}^{n+1,k+1}$, $\underline{\mathbf{d}}^{n+1,k}$ and $\underline{\delta \mathbf{d}}^k$, respectively, and the residual and Jacobian evaluated with $\mathbf{d}^{n+1,k}$ by $\underline{\mathbf{R}}(\mathbf{d}^{n+1,k})$ and $\underline{\mathbf{J}}(\mathbf{d}^{n+1,k})$, the linear system corresponding to (5.68) reads

$$\underline{\mathbf{J}}(\mathbf{d}^{n+1,k}) \underline{\delta \mathbf{d}}^k = -\underline{\mathbf{R}}(\mathbf{d}^{n+1,k}). \quad (5.69)$$

The absolute and relative convergence criteria on the Newton update, i.e., the difference in structural displacement iterates are defined very much similar to (2.79) as

$$\left\| \underline{\mathbf{d}}^{n+1,k+1} - \underline{\mathbf{d}}^{n+1,k} \right\| \leq \epsilon_{\text{abs}}^{\text{nl}} \quad \text{and} \quad \left\| \underline{\mathbf{d}}^{n+1,k+1} - \underline{\mathbf{d}}^{n+1,k} \right\| \leq \epsilon_{\text{rel}}^{\text{nl}} \|\underline{\mathbf{d}}^n\|. \quad (5.70)$$

Then, constructing the initial guess $\mathbf{d}^{n+1,0}$ by extrapolation based on previous time step's solution vectors, that is,

$$\mathbf{d}^{n+1,0} = \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{d}^{n-j}, \quad (5.71)$$

with coefficients β_j^m listed in Tab. 2.1, the Newton algorithm applied to solve the nonlinear structure problem (5.48) iteratively at time t^{n+1} reads

-
1. *Initial guess:* Extrapolate in time via (5.71) based on \mathbf{d}^{n-j} , $j = 0, \dots, m-1$.
 2. *Apply Dirichlet conditions:* Enforce Dirichlet conditions (5.12) directly on the initial guess $\mathbf{d}^{n+1,0}$, i.e., set $\mathbf{d}^{n+1,0}|_{\Gamma_D} = \mathbf{g}^{n+1}$.
 3. *Newton iterations:*
WHILE not converged according to (5.70) **DO**
 - a) Assemble the Jacobian system (5.69) using the last iterate $\mathbf{d}^{n+1,k}$ and solve for the update $\delta \mathbf{d}^k$ with $\delta \mathbf{d}^k|_{\Gamma_D} = \mathbf{0}$.
 - b) Update the k -th iterate via (5.49), i.e., $\mathbf{d}^{n+1,k+1} = \mathbf{d}^{n+1,k} + \delta \mathbf{d}^k$.

END WHILE

This scheme is the basic form of the nonlinear solver applied within the FSI coupling scheme introduced in Ch. 7, and will be slightly modified with respect to the boundary integral terms to account for viscoelastic tissue support or Robin interface conditions in the FSI setting (see Secs. 6.2 and 6.3). Furthermore, prestress present in the Lagrangian reference frame in geometries reconstructed from medical images is accounted for following Hsu and Bazilevs [318], again slightly altering the residual and corresponding Jacobian in the Newton system as will be discussed in Sec. 6.4. Apart from these modifications affecting the weak form and its linearisation, the following section moreover introduces an algorithm to construct suitable fibre orientations in complex geometries as encountered in the cardiovascular setting.

6 Structural dynamics: extensions to biomedical applications

In the modelling and simulation of biomedical systems, several aspects beyond the material models and respective parameters need to be considered in order to capture the underlying effects in challenging patient-specific scenarios. Similar to the flow solver's extensions towards practical applications in the biomedical context presented in Ch. 4, the structural solver as introduced in the preceding section is enhanced to account for key aspects imperative for the design of digital twins, virtual surgery or other kinds of patient-specific simulations. First, we introduce an automatic algorithm for the creation of local material orientations to make full use of anisotropic material models in complex geometries involving thin tissue layers wetted from both sides in Sec. 6.1. Second, a well-established prestress algorithm is described, giving details on the continuation method, nonlinear solver and pseudo-timestepping procedure to reliably find a stress state in the reference frame counteracting flow conditions present at the time of medical imaging. Third, Robin boundary conditions are covered to account for external tissue support, restricting the motion of the respective boundary segment of the tissue or organ. Fourth and last, the inclusion of suitable Robin boundary terms on the fluid–structure interface relevant for Robin-based FSI coupling as discussed in Ch. 7 is laid out for later reference. This section combines the author's works [5, 6], herein summarising the contributions in a unified presentation.

As these aspects can have a critical impact on the computational results obtained with numerical simulation, their inclusion is vital for any of the mentioned applications. However, within this work we limit ourselves to the said models without aiming for full clinical relevance at this point. The goal is to factor in the needed algorithms to account for essential effects while preserving accuracy, robustness and efficiency of the original methods rather than validating model and parameter choices in clinical applications. This is motivated by the fact that only when suitably incorporated in a first step, which we target herein, and then validated in a second step, such computational tools can be used for clinical support and decision making to compare various treatment options and simulating possible outcomes. Capturing realistic responses and identifying potential risk factors for disease initiation or progression via computational models can thus only be reached when capturing all major effects of the biomedical systems of interest.

For the sake of presentation, we will focus on vessel trees without loss of generality. As the cardiovascular system is essentially a pipe network, we define a single inlet $\hat{\Gamma}_{\text{in}}$ and a total of N_{out} outlets denoted by $\hat{\Gamma}_i$, $i = 1, \dots, N_{\text{out}}$ which compose the outlet boundary

$$\hat{\Gamma}_{\text{out}} := \bigcup_{i=1}^{N_{\text{out}}} \hat{\Gamma}_i, \quad (6.1)$$

which together with the inlet face make up the entire Dirichlet boundary for the structural dynamics problem

$$\hat{\Gamma}_D := \hat{\Gamma}_{\text{in}} \cup \hat{\Gamma}_{\text{out}}. \quad (6.2)$$

Further, we define the fluid–structure interface in the Lagrangian reference configuration $\hat{\Sigma}$ being the solid’s boundary section in contact with the fluid and the exterior surface of the structure $\hat{\Gamma}_R$, where viscoelastic support is accounted for. On the latter two boundary segments, Robin boundary conditions are enforced as detailed in Secs. 6.2 and 6.3. Altogether, we thus consider a decomposition of the domain’s boundary into non-overlapping segments $\hat{\Gamma}_{\text{in}}$, $\hat{\Gamma}_{\text{out}}$, $\hat{\Sigma}$ and $\hat{\Gamma}_R$,

$$\partial\hat{\Omega} = \hat{\Gamma}_{\text{in}} \cup \hat{\Gamma}_{\text{out}} \cup \hat{\Sigma} \cup \hat{\Gamma}_R. \quad (6.3)$$

6.1 Material orientation for cardiovascular settings

For complex solid geometries as often encountered in clinical scenarios, constructing a physiologically meaningful local material orientation is not straightforward. Hence, standard algorithms often fail to deliver satisfactory results for relevant examples in the medical context such as bifurcating, aneurysmatic or dissected vessels, the tracheobronchial tree or the heart (see, e.g., [266, 316, 353]). As the element-local material coordinates enter the constitutive equations, standard tools for mesh generation from medical image data need to be augmented by suitable methods to construct meaningful orientations in given volumetric meshes. The computational algorithms are typically based on a simple set of rules to achieve this task, while remaining fast, reliable and tunable by the user. Common approaches involve solving multiple auxiliary Laplace problems with user-specified boundary data on specific sections of the geometry [354, 355]. A first step in most algorithms for constructing a material orientation applied to a pipe network might then read

$$-\Delta\phi_l = 0 \quad \text{in } \hat{\Omega}, \quad (6.4)$$

$$\phi_l = 0 \quad \text{on } \hat{\Gamma}_{\text{in}}, \quad (6.5)$$

$$\hat{\mathbf{n}} \cdot \nabla\phi_l = h_{l,i} \quad \text{on } \hat{\Gamma}_i, \quad i = 1, \dots, N_{\text{out}}, \quad (6.6)$$

$$\hat{\mathbf{n}} \cdot \nabla\phi_l = 0 \quad \text{on } \hat{\Sigma} \cup \hat{\Gamma}_R, \quad (6.7)$$

to obtain an auxiliary scalar variable ϕ_l dependent on the prescribed fluxes $h_{l,i}$, potentially set individually for each of the N_{out} outlets. Enforcing zero Dirichlet conditions for ϕ_l at the inlet and zero natural conditions on the remaining boundary segments, i.e., on the fluid–structure interface $\hat{\Sigma}$ and the exterior boundary $\hat{\Gamma}_R$, allows constructing a scalar field with fluxes approximating the longitudinal vessel direction. The Neumann

data $h_{l,i}$ is then often fine-tuned manually, such that the longitudinal direction vector \mathbf{e}_2 is well approximated by the scaled gradient of the scalar field ϕ_l , given by

$$\mathbf{e}_2 := \frac{\nabla \phi_l}{\|\nabla \phi_l\|}. \quad (6.8)$$

Standard approaches then proceed by constructing the radial direction, i.e., the tissue-thickness direction based on a second auxiliary Laplace problem,

$$-\Delta \phi_n = 0 \quad \text{in } \hat{\Omega}, \quad (6.9)$$

$$\phi_n = 0 \quad \text{on } \hat{\Gamma}_R, \quad (6.10)$$

$$\hat{\mathbf{n}} \cdot \nabla \phi_n = 0 \quad \text{on } \hat{\Gamma}_D, \quad (6.11)$$

$$\hat{\mathbf{n}} \cdot \nabla \phi_n = h_n \quad \text{on } \hat{\Sigma}, \quad (6.12)$$

to retrieve ϕ_n with a large enough gradient in radial direction \mathbf{e}_3 , again manually tuning the parameter h_n as needed. Note that zero Neumann conditions are enforced at the in- and outlets, summarised as $\hat{\Gamma}_D$, and at the exterior surface $\hat{\Gamma}_R$, zero Dirichlet conditions are prescribed. Setting the gradient at the fluid–structure interface $\hat{\Sigma}$ via h_n thus yields the desired gradient in radial/tissue-thickness direction. Then, we can define

$$\mathbf{e}_3 := \frac{\nabla \phi_n}{\|\nabla \phi_n\|}, \quad (6.13)$$

to finally obtain the circumferential direction \mathbf{e}_1 via

$$\mathbf{e}_1 := \frac{\mathbf{e}_2 \times \mathbf{e}_3}{\|\mathbf{e}_2 \times \mathbf{e}_3\|}. \quad (6.14)$$

The longitudinal and radial directions are not necessarily exactly orthogonal to each other given the method of construction, but the circumferential direction is normal to both the other vectors in the basis by design. One might thus rotate either \mathbf{e}_2 or \mathbf{e}_3 prior to constructing \mathbf{e}_1 in the plane spanned by \mathbf{e}_2 and \mathbf{e}_3 by a minimal angle, such that $\mathbf{e}_2 \cdot \mathbf{e}_3 = 0$, that is, enforcing orthogonality. This, however, does not improve the approximation. A simpler approach in this regard is to construct \mathbf{e}_1 via Eqn. (6.14) as just laid out and then construct mean fibre directions \mathbf{m}_4 and \mathbf{m}_6 based on \mathbf{e}_1 and \mathbf{e}_2 *only*, as shown in Eqn. (5.39),

$$\mathbf{m}_4 := \frac{\mathbf{e}_1 + \mathbf{e}_2 \tan \alpha_c}{\|\mathbf{e}_1 + \mathbf{e}_2 \tan \alpha_c\|}, \quad \mathbf{m}_6 := \frac{\mathbf{e}_1 - \mathbf{e}_2 \tan \alpha_c}{\|\mathbf{e}_1 - \mathbf{e}_2 \tan \alpha_c\|},$$

thereby implicitly rotating \mathbf{e}_3 (for orthogonality) and keeping \mathbf{e}_2 fixed. The mean fibre orientation vectors \mathbf{m}_4 and \mathbf{m}_6 then directly enter the strain-energy (5.40) and thus the second Piola–Kirchhoff stress tensor (5.41).

This rule-based assignment of boundary conditions in the auxiliary Laplace equations including boundary data as tuning parameters is well suited for curved, bulging and even bifurcating vessels. However, changes of the cross-section of the vessel can lead to problematic configurations, if the exterior of the vessel remains largely unchanged.

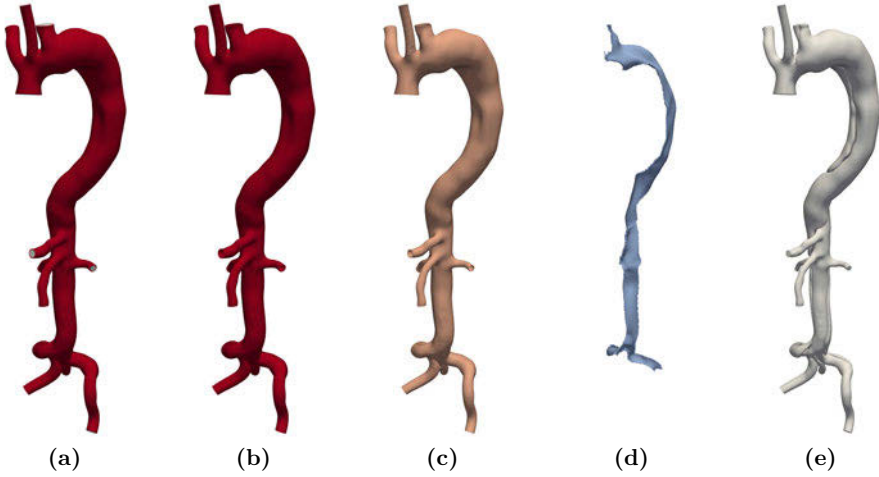


Fig. 6.1: Geometric discretisation of the aortic dissection case: overall geometry (a) composed of adventitia (b) and media (c) tissue layers and the dissection flap (d) separating the true and false lumina composing the fluid domain (e).

In such a scenario, prescribing fluxes \mathbf{h}_n on $\hat{\Gamma}_R$ or any Dirichlet data in the auxiliary Laplace problem with boundary conditions enforced on segments as introduced above does not allow constructing a meaningful radial direction \mathbf{e}_3 via ϕ_n . A typical example in this regard are geometries encountered in the context of aortic dissection (AD), which is a cardiovascular disease where the single original lumen splits into so-called true and false lumina (see, e.g., [332, 356, 357]). To ease discussion and to support our arguments, consider a patient-specific AD reconstructed from medical data provided by the authors of [266]. The model of the aortic tree as depicted in Fig. 6.1 considers the media and adventitia tissue layers and additionally differentiates the so-called dissection flap from the remaining tissue. The dissection flap, shown in light blue in Fig. 6.1, is a thin tissue layer separating the cavities where blood is flowing through. One differentiates the *true* lumen being the original lumen and the *false* lumen, which is created by tissue rupturing in a spiralling pattern along the aorta. The intimal tear, where the dissection flap separates from the media is in the present geometry located at the anterior vessel wall of the ascending aortic arch. The dissection flap stretches from the aortic arch down into the iliac bifurcation, lying at the very bottom of the considered geometry.

Focusing now on the arch region, Fig. 6.2 depicts the finite element discretisation including all tissue layers and the region occupied by fluid, see Fig. 6.2(a), where omitting anterior vessel wall segments of the media and adventitia layers of the structural model, see Fig. 6.2(b), reveals the conforming mesh underneath, which occupies the fluid domain shown in grey. Further leaving the fluid elements out in Fig. 6.2(b), the dissection flap coloured in light blue is visible in Fig. 6.2(c). In the current scenario, the larger false lumen lies anterior (in the front), whereas the smaller true lumen is located posterior (in the back). The blood originally circulated through what is now the true lumen, before the initiation and onset of aortic dissection created the false lumen.




Fig. 6.2: Finite element discretisation of the dissected aortic arch: adventitia (red), media (peach) and dissection flap (light blue) regions of the tissue and lumen (grey). Complete mesh (a), cut aortic arch revealing lumen mesh (b) and cut aortic arch showing tissue elements only (c).

Applying the procedure introduced above with appropriate $h_{l,i}$ $i = 1, \dots, N_{\text{out}}$ and h_n , we obtain the scalar fields ϕ_l and ϕ_n as shown in Figs. 6.3(a) and 6.3(b). The auxiliary scalar ϕ_l is suited well for deriving a longitudinal tissue orientation \mathbf{e}_2 as can be seen from Fig. 6.3(c). Only in the transition regions from the main to the branching vessels and at the intimal tear, \mathbf{e}_2 might not align with the longitudinal direction perfectly. It is, however, clear that the physiological mean fibre orientations vary smoothly in orientation in these transition regions as well. This motivates the use of the longitudinal direction \mathbf{e}_2 as is and accepting potential modelling errors caused by possible discrepancies.

Inspecting the solution to the second auxiliary Laplace problem ϕ_n shown in Fig. 6.3(b), one observes that prescribing the gradient $\mathbf{n} \cdot \phi_n = h_n$ at the fluid–structure interface leads to smaller values of ϕ_n in the dissection flap, whereas ϕ_n remains larger in the media and adventitia layers of the tissue. The circumferential orientation vector \mathbf{e}_1 in Fig. 6.3(d) based on this ϕ_n is a physiologically meaningful approximation in the exterior wall region, but ϕ_n fails to deliver even remotely useful mean fibre directions in the dissection flap region. Considering anisotropic tissue models for the dissection flap is thus unfeasible, but might be important given possibly large strains present especially in these tissue regions.

This might be mended by (i) considering media and adventitia layers and the dissection flap separately or (ii) building the finite element mesh from pre-built and mapped building blocks with intrinsic material orientations pre-defined. Since the first route would require detecting different regions of the tissue to formulate proper boundary conditions for the arising subproblems and therefore necessitate further tuning parameters or manual intervention, such an approach is rather impractical for clinical application. The second option might in fact yield satisfactory results, especially since higher-order

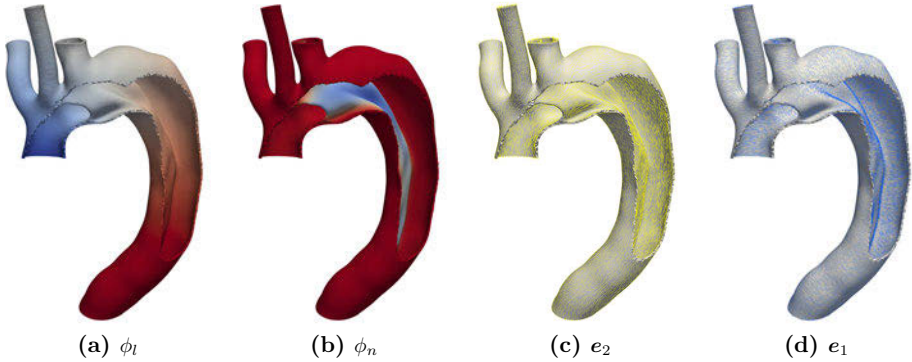


Fig. 6.3: Constructing the material orientation in the aortic arch region: solutions to the auxiliary Laplace problems (a) and (b) with lower values in blue and higher values in red to construct longitudinal (c) and radial directions (d). In the dissection flap, the circumferential \mathbf{e}_1 based on ϕ_n is meaningless.

meshes including boundary layers can be constructed automatically based on predefined coarse meshes capturing all possible topological scenarios—this approach, however, is not considered within this work.

Herein, we resort to a much simpler remedy based on a given spatial discretisation simply exchanging the second step of the orientation algorithm, which gave the faulty radial direction \mathbf{e}_3 for the topologies of interest, i.e., pipes with thin tissue layers wetted from both sides. The idea is to replace the second auxiliary Laplace problem governing ϕ_n (6.9)–(6.12) with an extrapolation of the fluid–structure interface normal into the solid domain $\hat{\Omega}$ as presented in our previous work (see, e.g., [6, 10–12]). In tissue parts being in contact with the fluid from both sides such as the dissection flap, this is still applicable if the orientation vectors are extrapolated and averaged with care. Only data stemming from either the front or the back of the dissection flap can be considered, otherwise the situation is not improved compared to the purely Laplace-based approach.

Thus, a first step determines the normal vectors $\hat{\mathbf{n}}$ on $\hat{\Sigma}$ and stores them for all elements lying on $\hat{\Sigma}$. In a second step, the mean orientation of all elements touching the previously marked element layer is determined, which is repeated until \mathbf{e}_3 is set in all elements. During this step, either consider only data from the first neighbour of each element that has an orientation, or optionally consider all neighbours of a given element with radial direction $\mathbf{e}_{3,j}$ deviating from the first encountered neighbour with orientation $\mathbf{e}_{3,i}$ satisfying the angle condition

$$\alpha_{\text{tol}} \geq \arccos \left(\frac{\mathbf{e}_{3,i} \cdot \mathbf{e}_{3,j}}{\|\mathbf{e}_{3,i}\| \|\mathbf{e}_{3,j}\|} \right). \quad (6.15)$$

Condition (6.15) holds if the smallest angle between the vectors $\mathbf{e}_{3,i}$ and $\mathbf{e}_{3,j}$ is less than some tolerance angle α_{tol} . Afterwards, a few cycles of conditional averaging radial directions of neighbouring elements suffices to obtain a satisfactory vector field.

The final algorithm is consequently only based on detecting the boundary segments $\hat{\Gamma}_{\text{in}}$, $\hat{\Gamma}_i$, $i = 1, \dots, N_{\text{out}}$, $\hat{\Gamma}_R$ and $\hat{\Sigma}$ and prescribing suitable fluxes $\mathbf{h}_{l,i}$ on each of the N_{out} outlets. A single scalar-valued Laplace problem is solved for ϕ_l to compute the longitudinal direction vector \mathbf{e}_2 , after which the unit normal vectors on $\hat{\Sigma}$ are conditionally extrapolated into the structural domain layer by layer (element neighbour by element neighbour) and smoothed afterwards using condition (6.15) again. The overall algorithm to construct physiologically meaningful material orientation vectors \mathbf{e}_1 and \mathbf{e}_2 can thus be summarised as

-
1. *Auxiliary Laplace problem:* Solve (6.4)–(6.7) for ϕ_l with suitable Neumann data $\mathbf{h}_{l,i}$ given on $\hat{\Gamma}_i$, $i = 1, \dots, N_{\text{out}}$.
 2. *Longitudinal direction:* Construct \mathbf{e}_2 from ϕ_l via Eqn. (6.8).
 3. *Radial direction:*
 - a) Set $\mathbf{e}_3 = \hat{\mathbf{n}}$ in elements on $\hat{\Sigma}$ and extrapolate layer by layer into $\hat{\Omega}$:
WHILE \mathbf{e}_3 not set in all elements **DO**
 - i. On each element find neighbouring elements with \mathbf{e}_3 set and store $\mathbf{e}_{3,i}$ being the first encountered neighbour's \mathbf{e}_3 .
 - ii. Average all neighbour's orientations $\mathbf{e}_{3,j}$ for which (6.15) holds and assign the average to \mathbf{e}_3 .**END WHILE**
 - b) Perform conditional averaging cycles:
FOR N_{avg} iterations **DO**
 - i. Average $\mathbf{e}_{3,i}$ on each element, taking neighbour's directions $\mathbf{e}_{3,j}$ into account given (6.15) is met.**END FOR**
 4. *Circumferential direction:* Compute \mathbf{e}_1 given \mathbf{e}_2 and \mathbf{e}_3 via (6.14).
 5. *Mean fibre directions:* Optionally precompute \mathbf{m}_4 and \mathbf{m}_6 via (5.39).
-

This algorithm using $h_{l,i} = 100$ and $\alpha_{\text{tol}} = 120^\circ$ in (6.15) and a total of $N_{\text{avg}} = 5$ averaging rounds yields satisfactory vector fields \mathbf{e}_1 and \mathbf{e}_2 as shown in Figs. 6.4(a) and (c) for a zoom-in on the dissection flap. Fig. 6.4(b) also shows the circumferential direction as computed by the approach solving two Laplace problems. It can clearly be seen that the orientation vectors in the circumferential direction are greatly improved employing the newly proposed algorithm. Note that a smooth transition of the material orientation is seen in the dissection flap attachment regions, but naturally, the leading edge of the dissection flap does still not feature a “proper” material orientation.

The proposed method based on the conditionally-averaged extrapolation of the interface normal into the structural domain is applied to all relevant examples presented in Secs. 8.3 and 8.5–8.7 without any further parameter tuning. In all considered numerical tests, suitable material orientations for fibre-reinforced solids are easily constructed. The




Fig. 6.4: Constructing the material orientation in the aortic arch region: longitudinal direction e_2 computed from ϕ_l (a) and comparison of circumferential directions e_1 computed from ϕ_n via a second Laplace problem (b) and an extrapolated, conditionally-averaged interface normal (c).

algorithm is fairly robust with respect to the parameter choice, as $h_{l,i} = 100$, $\alpha_{\text{tol}} = 120^\circ$ and $N_{\text{avg}} = 5$ yields reasonable results in all our tests including the aortic dissection case. Additionally, it is found completely irrelevant in terms of computing time compared to fluid flow, nonlinear structural mechanics or FSI simulations.

6.2 Accounting for external tissue support

In many practical applications and also, e.g., in the targeted biomedical setting, the geometric discretisation of the computational domain can only incorporate a limited subset of the solid continuum. In the cardiovascular context, supporting structures, e.g., the spine, diaphragm or other organs tethering the abdominal aorta restrain excessive movement [358–360]. However, a rigorous multicontact model and volumetric representation of all the involved organs, soft tissue and bones would tremendously increase the computational complexity. This renders such an approach much too convoluted, almost infeasible judged from the current state of the art, especially when taking the need for segmentation and mesh generation into account. Therefore, the parts in focus—say, a limited set of vessels—are embedded into a larger bulk continuum due to (i) lack of data and/or (ii) to reduce the computational burden. This inevitable domain truncation then considers only regions of major influence, modelling the remainder by means of (non-)linear boundary conditions—geometrically condensing the surrounding structures.

A simple strategy along these lines consists of including springs and dashpots to model the external tissue’s viscoelastic support. Such models have been considered in the context of cardiovascular FSI with great success, see, e.g. [248, 266, 359–364] and other works in literature, enforcing a Robin boundary condition of the form

$$\mathbf{P} \hat{\mathbf{n}} = -k_e \mathbf{d} - c_e \mathbf{d}_t - p_e \hat{\mathbf{n}} \quad \text{on } \hat{\Gamma}_R, \quad (6.16)$$

where k_e and c_e are the parameters of the surrounding viscoelastic tissue and p_e denotes the intrathoracic pressure acting on the exterior surface of the vessel $\hat{\Gamma}_R$. Such a condition with suitable k_e and c_e has important implications on the numerical treatment being that realistic displacements of the vessel wall are obtained tuning k_e and c_e rather than adapting the tissue's stiffness beyond the physiologically justifiable range. Also, the computed displacements are lower when comparing to a zero Neumann condition enforced on $\hat{\Gamma}_R$ and damping is introduced via c_e , such that solutions to structural dynamics or FSI problems are easier found given the increased temporal stability and structural stiffness. Another advantage of introducing Robin boundary conditions is that it does not introduce any additional degrees of freedom, such that the system size remains unchanged. In addition to that, the entries added to the Jacobian matrix are symmetric mass matrix terms stemming from a boundary integral over $\hat{\Gamma}_R$.

Starting again from the generalised trapezoidal rule applied to the first Piola–Kirchhoff stress tensor's divergence in the weak form, one can write

$$\begin{aligned} -\langle \boldsymbol{\varphi}, \hat{\nabla} \cdot \mathbf{P}^{n+1-\alpha_f} \rangle &= -\langle \boldsymbol{\varphi}, \alpha'_f \hat{\nabla} \cdot \mathbf{P}^{n+1} + \alpha_f \hat{\nabla} \cdot \mathbf{P}^n \rangle \\ &= \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}^{n+1} + \alpha_f \mathbf{P}^n \rangle - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{P}^{n+1} \hat{\mathbf{n}} + \alpha_f \mathbf{P}^n \hat{\mathbf{n}} \rangle_{\partial \hat{\Omega}}, \end{aligned}$$

where we can directly insert the Robin condition in the boundary term on $\hat{\Gamma}_R \subset \partial \hat{\Omega}$. There, leaving the stress terms at time t^n in terms of \mathbf{d}^n , i.e., *not* inserting the Robin condition, we can focus on the evaluation and linearisation at $t = t^{n+1}$. The relevant boundary term is thus rewritten as

$$-\langle \alpha'_f \boldsymbol{\varphi}, \mathbf{P}^{n+1} \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_R} = \langle \alpha'_f \boldsymbol{\varphi}, k_e \mathbf{d}^{n+1} + c_e \dot{\mathbf{d}}^{n+1} + p_e \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_R},$$

which is recast in terms of \mathbf{d}^{n+1} using the Newmark formula (5.45). The boundary term added on $\hat{\Gamma}_R$ hence reads

$$\langle \alpha'_f \boldsymbol{\varphi}, k_e \mathbf{d}^{n+1} + c_e \left[\frac{\gamma}{\Delta t \beta} (\mathbf{d}^{n+1} - \mathbf{d}^n) + \left(1 - \frac{\gamma}{\beta}\right) \dot{\mathbf{d}}^n + \Delta t \left(1 - \frac{\gamma}{2\beta}\right) \ddot{\mathbf{d}}^n \right] + p_e \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_R}. \quad (6.17)$$

Naturally, this term contributes to the residual of the momentum balance equation and corresponding Jacobian in Newton's method. To this end, adapt the semilinear form $\mathbf{R}(\mathbf{d})(\boldsymbol{\varphi})$ in its original form defined in Eqn. (5.47) to

$$\begin{aligned} \mathbf{R}(\mathbf{d})(\boldsymbol{\varphi}) &:= \langle \rho \boldsymbol{\varphi}, \alpha'_m \ddot{\mathbf{d}}^{n+1}(\mathbf{d}) + \alpha_m \ddot{\mathbf{d}}^n \rangle \\ &\quad - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{b}^{n+1} + \alpha_f \mathbf{b}^n \rangle + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}(\mathbf{d}) + \alpha_f \mathbf{P}^n \rangle - \langle \boldsymbol{\varphi}, \alpha'_f \hat{\mathbf{t}}^{n+1} + \alpha_f \hat{\mathbf{t}}^n \rangle_{\hat{\Sigma}} \\ &\quad + \langle \boldsymbol{\varphi}, \alpha'_f [k_e \mathbf{d} + c_e \dot{\mathbf{d}}^{n+1}(\mathbf{d}) + p_e \hat{\mathbf{n}}] - \alpha_f \mathbf{P}^n \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_R}, \end{aligned} \quad (6.18)$$

assuming for now tractions $\hat{\mathbf{t}}^{n+1}$ and $\hat{\mathbf{t}}^n$ given on $\hat{\Sigma}$ and with $\dot{\mathbf{d}}^{n+1}(\mathbf{d})$ and $\ddot{\mathbf{d}}^{n+1}(\mathbf{d})$ according to Eqns.(5.44)–(5.45).

Since the newly introduced Robin term depends on the solution \mathbf{d}^{n+1} , the Jacobian $\mathbf{J}(\mathbf{d})(\boldsymbol{\varphi})$ is modified as well, and the final problem is to find the k -th Newton update $\delta \mathbf{d}^k \in [X(\hat{\Omega})]^d$ with $\delta \mathbf{d}|_{\hat{\Gamma}_D} = \mathbf{0}$ to the last iterate $\mathbf{d}^{n+1,k}$ of the displacement \mathbf{d}^{n+1} at time $t = t^{n+1}$, given the last time steps' $\ddot{\mathbf{d}}^n$, $\dot{\mathbf{d}}^n$ and \mathbf{d}^n , such that

$$\begin{aligned} \langle \rho \boldsymbol{\varphi}, \frac{\alpha'_m}{\beta \Delta t^2} \delta \mathbf{d} \rangle + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{D}_d \mathbf{P}(\mathbf{d}^{n+1,k}) (\delta \mathbf{d}) \rangle + \langle \alpha'_f \boldsymbol{\varphi}, k_e \delta \mathbf{d} + \frac{c_e \gamma}{\Delta t \beta} \delta \mathbf{d} \rangle_{\Gamma_R} \\ = -\mathbf{R}(\mathbf{d}^{n+1,k})(\boldsymbol{\varphi}), \end{aligned} \quad (6.19)$$

for all $\boldsymbol{\varphi} \in [X(\hat{\Omega})]^d$ with $\boldsymbol{\varphi}|_{\Gamma_D} = \mathbf{0}$ in analogy to (5.68). Inhomogeneous boundary conditions are enforced on the initial guess, $\mathbf{d}^{n+1,0}$, such that the Newton iterate incorporating homogeneous Dirichlet conditions is updated via $\mathbf{d}^{n+1,k+1} = \mathbf{d}^{n+1,k} + \delta \mathbf{d}^k$.

6.3 Robin interface conditions for fluid–structure interaction

Partitioned FSI approaches applied in the context of biomechanics are prone to instabilities. As the solid and fluid densities approach each other, the computational domain slenderness increases, or the time step size is reduced, the so-called added-mass effect becomes increasingly important [365, 366]. Therefore, the computational modelling of compliant vessels or related problems stemming from the cardiovascular system is a challenging task. Standard partitioned methods such as the serial staggered method [367] and even fully implicit methods suffer under certain physical parameter combinations typical for cardiovascular FSI (see, e.g., [88, 368, 369]).

Among the remedies presented in Sec. 7.4 to reach robust convergence of the coupling scheme and temporal stability is introducing Robin interface conditions, leading to Robin–Robin coupling schemes and variants thereof (see, e.g., [217–224, 362] or [5, 6, 225] for projection/split-step fluid solvers). These schemes are distinctly different from the classical Dirichlet–Neumann approach, as will be discussed in great detail in Ch. 7. Relevant within this section is solely the Robin condition as enforced on the fluid–structure interface $\hat{\Sigma}$ in the solid subproblem, which reads

$$\eta^R \mathbf{d}_t \mathbf{d} + \mathbf{P} \hat{\mathbf{n}} = \hat{\mathbf{h}} := \eta^R \mathbf{u}_f + \mathbf{P}_f \hat{\mathbf{n}} \quad \text{on } \hat{\Sigma}, \quad (6.20)$$

where the traction vector $\hat{\mathbf{h}}$ typically depends on the fluid’s velocity \mathbf{u}_f and the fluid’s first Piola–Kirchhoff stress tensor \mathbf{P}_f on the interface $\hat{\Sigma}$. The parameter η^R scales from enforcing continuity of velocities ($\eta^R \rightarrow \infty$) to a Neumann condition equating tractions ($\eta^R \rightarrow 0$). As a starting point for later investigations in the FSI context, we introduce a Robin condition to the structure problem in the form of

$$\eta^R \mathbf{d}_t \mathbf{d} + \mathbf{P} \hat{\mathbf{n}} = \hat{\mathbf{h}} \quad \text{on } \hat{\Sigma}, \quad (6.21)$$

similar to the fluid Robin condition considered in Eqn. (4.42) in Sec. 4.2 or the Robin condition considered for the structure problem to account for viscoelastic tissue support (6.16) in the preceding section. As already shown, the stress terms after applying the divergence theorem are given as

$$\begin{aligned} -\langle \boldsymbol{\varphi}, \hat{\nabla} \cdot \mathbf{P}^{n+1-\alpha_f} \rangle &= -\langle \boldsymbol{\varphi}, \alpha'_f \hat{\nabla} \cdot \mathbf{P}^{n+1} + \alpha_f \hat{\nabla} \cdot \mathbf{P}^n \rangle \\ &= \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}^{n+1} + \alpha_f \mathbf{P}^n \rangle - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{P}^{n+1} \hat{\mathbf{n}} + \alpha_f \mathbf{P}^n \hat{\mathbf{n}} \rangle_{\partial \hat{\Omega}}, \end{aligned}$$

which allows inserting the Robin condition on $\hat{\Sigma}$, rewritten as

$$\mathbf{P} \hat{\mathbf{n}} = \hat{\mathbf{h}} - \eta^R d_t \mathbf{d},$$

yielding for the implicit term at time $t = t^{n+1}$

$$-\langle \alpha'_f \boldsymbol{\varphi}, \mathbf{P}^{n+1} \hat{\mathbf{n}} \rangle_{\hat{\Sigma}} = \langle \alpha'_f \boldsymbol{\varphi}, \eta^R \dot{\mathbf{d}}^{n+1} - \hat{\mathbf{h}}^{n+1} \rangle_{\hat{\Sigma}},$$

which can once again be recast in terms of \mathbf{d}^{n+1} via the Newmark formula (5.45). The boundary term involving \mathbf{d}^{n+1} added on $\hat{\Sigma}$ then reads

$$\langle \alpha'_f \boldsymbol{\varphi}, \eta^R \left[\frac{\gamma}{\Delta t \beta} (\mathbf{d}^{n+1} - \mathbf{d}^n) + \left(1 - \frac{\gamma}{\beta}\right) \dot{\mathbf{d}}^n + \Delta t \left(1 - \frac{\gamma}{2\beta}\right) \ddot{\mathbf{d}}^n \right] - \hat{\mathbf{h}}^{n+1} \rangle_{\hat{\Sigma}}. \quad (6.22)$$

As before, the momentum balance residual and the Jacobian in Newton's method are subject to changes. The semilinear form $\mathbf{R}(\mathbf{d})(\boldsymbol{\varphi})$ is once again adapted from (5.47) or (6.18) to

$$\begin{aligned} \mathbf{R}(\mathbf{d})(\boldsymbol{\varphi}) &:= \langle \rho \boldsymbol{\varphi}, \alpha'_m \ddot{\mathbf{d}}^{n+1}(\mathbf{d}) + \alpha_m \ddot{\mathbf{d}}^n \rangle \\ &\quad - \langle \boldsymbol{\varphi}, \alpha'_f \mathbf{b}^{n+1} + \alpha_f \mathbf{b}^n \rangle + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}(\mathbf{d}) + \alpha_f \mathbf{P}^n \rangle - \langle \alpha_f \boldsymbol{\varphi}, \mathbf{P}^n \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_R \cup \hat{\Sigma}} \\ &\quad + \langle \alpha'_f \boldsymbol{\varphi}, k_e \mathbf{d} + c_e \dot{\mathbf{d}}^{n+1}(\mathbf{d}) + p_e \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_R} + \langle \alpha'_f \boldsymbol{\varphi}, \eta^R \dot{\mathbf{d}}^{n+1}(\mathbf{d}) - \hat{\mathbf{h}}^{n+1} \rangle_{\hat{\Sigma}}, \end{aligned} \quad (6.23)$$

with $\dot{\mathbf{d}}^{n+1}(\mathbf{d})$ and $\ddot{\mathbf{d}}^{n+1}(\mathbf{d})$ according to Eqns. (5.44)–(5.45) and where Robin boundary conditions on the exterior boundary $\hat{\Gamma}_R$ are included as well. The Jacobian $\mathbf{J}(\mathbf{d})(\boldsymbol{\varphi})$ is adapted accordingly, such that each Newton step consists of finding the Newton update $\delta \mathbf{d}^k \in [X(\hat{\Omega})]^d$ with $\delta \mathbf{d}|_{\hat{\Gamma}_D} = \mathbf{0}$ given the last iterate $\mathbf{d}^{n+1,k}$ of the displacement \mathbf{d}^{n+1} at time $t = t^{n+1}$ and the last time step's solutions $\ddot{\mathbf{d}}^n$, $\dot{\mathbf{d}}^n$ and \mathbf{d}^n , such that

$$\begin{aligned} &\langle \rho \boldsymbol{\varphi}, \frac{\alpha'_m}{\beta \Delta t^2} \delta \mathbf{d} \rangle + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f D_d \mathbf{P}(\mathbf{d}^{n+1,k})(\delta \mathbf{d}) \rangle \\ &\quad + \langle \alpha'_f \boldsymbol{\varphi}, k_e \delta \mathbf{d} + \frac{c_e \gamma}{\Delta t \beta} \delta \mathbf{d} \rangle_{\hat{\Gamma}_R} + \langle \alpha'_f \boldsymbol{\varphi}, \frac{\eta^R \gamma}{\Delta t \beta} \delta \mathbf{d} \rangle_{\hat{\Sigma}} = -\mathbf{R}(\mathbf{d}^{n+1,k})(\delta \mathbf{d}), \end{aligned} \quad (6.24)$$

for all $\boldsymbol{\varphi} \in [X(\hat{\Omega})]^d$, $\boldsymbol{\varphi}|_{\hat{\Gamma}_D} = \mathbf{0}$, in analogy to (5.68) and (6.19). Inhomogeneous boundary conditions are enforced on the initial guess, $\mathbf{d}^{n+1,0}$, and the iterate is updated following $\mathbf{d}^{n+1,k+1} = \mathbf{d}^{n+1,k} + \delta \mathbf{d}^k$ with homogeneous Dirichlet conditions.

6.4 Prestressed reference configurations

In patient-specific haemodynamic simulations, or generally in scenarios where the structure is already loaded during imaging (to construct the computational domain), the solid stress state already present at the time of image acquisition—the prestress—is unknown, but might not have a negligible effect on the overall stress state in the continuum at the time of interest. In the cardiovascular system, computed tomography and magnetic resonance images are taken during the diastolic phase under corresponding flow conditions. As a matter of fact, regardless of when in the cardiac cycle images are taken, the quiescent state would never be reached and a certain pressure level is maintained at any point. Therefore, this initial loading can have a large influence on the overall response, such that prestress effects can not be neglected [266, 318, 370, 371].

From a continuum mechanics point of view, several possibilities exist, the most popular ones being: (i) find a new stress-free reference configuration, from which a newly defined deformation field maps to the deformed configuration (see, e.g., [355, 372]), or (ii) incorporate the prestress tensor in the solid’s balance of linear momentum, equating fluid tractions under diastolic flow conditions [318, 370, 371]. We herein employ the strategy introduced by Hsu and Bazilevs [318], since it can be straightforwardly incorporated into existing structural mechanics solvers, keeping the total Lagrangian viewpoint. So, we replace Eqn. (5.9) by its static counterpart including prestress to solve for a vector field \mathbf{d}_0 . This displacement \mathbf{d}_0 enters the prestress tensor \mathbf{S}_0 using the same material law and the basic relation $\mathbf{P} = \mathbf{F} \mathbf{S}$ as [318]

$$\mathbf{P}^{n+1} = \mathbf{F}(\mathbf{d}^{n+1}) \left[\mathbf{S}(\mathbf{d}^{n+1}) + \mathbf{S}_0(\mathbf{d}_0) \right],$$

where we ignore the obvious inconsistency being the additive split of nonlinear \mathbf{S} , viewing this approach simply as a convenient choice acting as a placeholder in our framework. In a first step, however, static equilibrium has to be determined using the prestress tensor \mathbf{S}_0 only, since in equilibrium, we have $\mathbf{F} = \mathbf{S} = \mathbf{0}$ and \mathbf{S}_0 in $\hat{\Omega}$ counteracts the (diastolic) external loading and fulfils static equilibrium

$$-\hat{\nabla} \cdot \mathbf{S}_0(\mathbf{d}_0) = \mathbf{0} \quad \text{in } \hat{\Omega}. \quad (6.25)$$

Starting from the quiescent state and ramping up the inlet conditions, flow quantities are kept fixed at a user-specified reference time t_0 , which corresponds to the imaging data, and Eqn. (6.25) is solved until the convergence criteria

$$\left\| \mathbf{d}_0^{k+1} - \mathbf{d}_0^k \right\|_{\max} < \epsilon_{\text{abs}}^{d_0} \quad \text{or} \quad \left\| \mathbf{d}_0^{k+1} - \mathbf{d}_0^k \right\| < \epsilon_{\text{rel}}^{d_0} \left\| \mathbf{d}_0^{k+1} \right\| \quad (6.26)$$

are fulfilled employing once again a Newton method. Afterwards, we set $\mathbf{d}_0 = \mathbf{d}_0^{k+1}$ and keep $\mathbf{S}_0(\mathbf{d}_0)$ fixed. Then, we can directly apply the standard timestepping scheme in the structural solver, with the solid momentum equation incorporating \mathbf{S}_0 :

$$\rho_s \mathbf{d}_{tt} \mathbf{d}_s - \hat{\nabla} \cdot [\mathbf{P}(\mathbf{d}_s) + \mathbf{F}(\mathbf{d}_s) \mathbf{S}_0(\mathbf{d}_0)] = \mathbf{0} \quad \text{in } \hat{\Omega}. \quad (6.27)$$

To efficiently tackle Eqn. (6.25), we combine adaptive, load-driven pseudo-timestepping with Aitken relaxation (2.80) (cf. Sec. 2.5) and a continuation technique. That is, we actually solve Eqn. (6.27) with fixed diastolic fluid load and $\eta^R = 0$. Searching for a stationary state, one is free to choose large physical mass- and stiffness-proportional damping combined with less tight tolerance settings. Applying a continuation method, the diastolic fluid load is scaled from 0 to 1 via a parameter η_c , while the density ρ_s is successively lowered from an initially increased value over N_c steps and strong viscous parameter c_e and mass-proportional damping ease convergence towards the stationary state. Simultaneously, the sought \mathbf{d}_0^{k+1} is updated using Aitken's relaxation (2.80) and the pseudo time step size is increased once less than N_{nl} Newton steps per nonlinear solver call are needed. The pseudo time integration is stopped once a steady state is achieved, i.e.,

$$\left\| \ddot{\mathbf{d}}_s^{k+1} \right\| < \epsilon_{\text{rel}}^{\ddot{\mathbf{d}}_0} \left\| \mathbf{d}_s^{k+1} \right\|, \quad (6.28)$$

and the convergence criteria (6.26) are fulfilled. For the convenience of the reader, the prestress procedure employed within this work is summarised as

-
1. *Prestress load*: Determine interface traction $\hat{\mathbf{t}}_{f,0} := \boldsymbol{\sigma}_f \hat{\mathbf{n}}$ acting on $\hat{\Sigma}$ in the undeformed configuration, e.g., by a precursor flow simulation.
 2. *Initialisation*: Set zero initial guess $\mathbf{S}_0 = \mathbf{0}$ via $\mathbf{d}_0^k = \mathbf{0}$ and zero initial conditions for pseudo-timestepping algorithm $\ddot{\mathbf{d}}^n = \dot{\mathbf{d}}^n = \mathbf{d}^n = \mathbf{0}$, $n = 0$, $k = 0$, $l = 0$.
 3. *Newton-steered pseudo-timestepping and continuation method*: Increase load factor η_c from 0 to 1 over N_c steps and adapt the pseudo time step size.

WHILE $l < N_c$ or convergence criteria (6.26) and (6.28) are not met **DO**

- a) Compute scaling parameter in continuation method $\eta_c = \frac{l+1}{N_c}$ and update structure density $\rho = \eta_c \rho^{\text{max}} + (1 - \eta_c) \rho^{\text{min}}$.
- b) Solve a pseudo time step of the structure's momentum balance equation (6.27) for \mathbf{d}^{n+1} with $\mathbf{S}_0 = \mathbf{0}$ and scaled traction $\hat{\mathbf{t}}|_{\hat{\Sigma}} = \eta_c \hat{\mathbf{t}}_{f,0}$.
- c) Relax the update via Aitken's method: $\mathbf{d}^{k+1} = \omega_k \mathbf{d}^{n+1} + (1 - \omega_k) \mathbf{d}^k$ and increment $k = k + 1$ and $n = n + 1$.

IF Newton solver converged in less than N_{nl} steps **THEN**

- i. Increase the pseudo time step size.
- ii. Set $l = l + 1$ if $l < N_c$.

END IF

END WHILE

4. *Start forward simulation*: Set $\mathbf{d}_0 = \mathbf{d}^{k+1}$ and reset the time history data of the structure, e.g., $\ddot{\mathbf{d}}^n = \dot{\mathbf{d}}^n = \mathbf{d}^n = \mathbf{0}$ and consider $\mathbf{S}_0(\mathbf{d}_0)$ in (6.27). The density and applied load are as desired (since $\eta_c = 1$ if l was increased N_c times), but the remaining structure parameters have to be reset.
-

Note that driving the flow with a ramped inlet condition from the quiescent state is often needed, especially when using Windkessel models. This is another purpose for the

fluid stepping included in the above algorithm and it hence does not add much overhead. Splitting the precursor flow simulation from the prestress algorithm is still preferred over performing unidirectional coupling. Main reasons for this design choice are that (i) the flow field can be stored and recovered easily after completing the pure flow phase once and (ii) the precursor flow simulation may include several cardiac cycles and a large amount of time steps not needed to converge the prestress algorithm. In a closing remark, we also want to note here again that the combination of pseudo-timestepping, relaxation and continuation methods is merely a placeholder for more advanced techniques common in nonlinear solid mechanics and is easily replaced in the final partitioned FSI scheme.

6.5 Summary and conclusion

Within this section, we introduce the structure subproblem, covering various constitutive relations, that is, linear elasticity, a St. Venant–Kirchhoff model, and hyperelastic neo-Hookean and Holzapfel–Gasser–Ogden models, weakly enforcing incompressibility via a penalty approach and considering fibre-reinforcement in the last case. A purely displacement-based, C^0 -continuous formulation adopting generalised- α time integration is given and the related nonlinear problem solved via a standard Newton scheme, where we give a detailed presentation of the required directional derivatives. With this setup, enforcing interface coupling conditions in the FSI context, linking the structural solver to the fluid subproblem adopting a (stabilised) coupled velocity-pressure formulation or PPE-based split-step scheme is enabled in a trivial fashion assuming matching spatial discretisations. Exchanging the solid subproblem solver is thus easily possible thanks to the partitioned design, while certain modelling aspects relevant for practical applications in the biomedical context need to be incorporated at the solid solver’s level. Targeting haemodynamic FSI, we then further introduce a novel algorithm to construct suitable material orientations to make use of anisotropic fibre-reinforcement in complex geometries such as encountered in aortic dissection, where the dissection flap, a thin tissue layer wetted from both sides, separates the so-called true and false lumina. Moreover, we account for viscoelastic external tissue support to suppress excessive deformation when considering vessels under realistic flow conditions and adopting tissue parameters in the physiological range, which becomes particularly relevant in FSI problems. Another important aspect covered herein concerns prestress present in the reference geometries obtained from segmented medical image data. The prestress counteracts a certain reference load, most often obtained from precursor flow simulations or prescribing pressure levels, such that this specified load leads to zero deformation *but* a nonzero stress state within the solid.

That being said, it is emphasised that the development of a structural solver does not lie within this work’s main focus, but should rather be seen as a necessity to formulate FSI problems. The presented approach may hence be replaced by a large variety of alternative schemes. Thus, contributions in this regard are (i) the material orientation algorithm and (ii) the overall prestress approach embedded in the FSI scheme. Furthermore, incorporating Robin (interface) conditions for exterior viscoelastic support or coupling in the context of FSI problems are vital ingredients as shown in Chs. 7 and 8.

At this point, the introduction of all the methods and techniques concerning the solution of the structure subproblem considered within this work is completed. Numerical examples of the presented methods and various combinations of algorithmic parameters are presented in the context of FSI applications and for now postponed to later chapters. The fluid and structure subproblems with suitable boundary conditions and accompanying solution strategies are now established, which allows moving on to the final topic, which is the coupling of fluid and structure subproblems on a shared interface following a partitioned FSI paradigm. The following Ch. 7 starts from the basic partitioned algorithm, then introduces more advanced coupling schemes and linearisation approaches and includes first small numerical experiments to verify the computational methods. Afterwards, Secs. 8.5–8.8 present results for practical application to an idealised abdominal aortic aneurysm, patient-specific geometries of an iliac bifurcation, a case of aortic dissection and an idealised setup of human phonation, which demonstrates the scheme's applicability in the context of aeroelasticity.

7 The coupled problem

Fluid–structure interaction (FSI) is a surface-coupled multiphysics problem, characterised by the strong mutual dependence of fluid flow through a deforming domain and structural deformation. Momentum is exchanged at the interface, while continuity of velocities and displacements holds. Fluid forces acting on the structure trigger deformations and hence induce strains within the solid continuum, thereby dictating the stress state in the solid phase. A moving or deforming solid, in turn, alters the domain the fluid is flowing through and therefore has a significant impact on the flow quantities observed. Capturing this bidirectional dependence of fluid flow and corresponding structural response lies at the heart of every FSI scheme.

Countless practical FSI problems of great interest are found in science, engineering and medicine, ranging from airfoils or whole wind turbines [373, 374], wind-induced periodic motion of bridge decks [375], offshore engineering [376, 377], insect flight [378] to blood flow through the circulatory system [11, 353, 364, 379–381], human phonation [382, 383] or respiration [353, 384]. Consequently, the development of suitable models and solution procedures has been an active area of research over the past 40 years, which lead to significant progress in the field.

Simultaneously, rapid advances in computational (bio-)mechanics have and still continue to increase the reliability and significance of patient-specific simulations through ever-increasing computational power, more efficient algorithms on the one side and statistical tools such as uncertainty quantification and sensitivity analysis on the other side. Latter of which live in symbiosis with the traditional numerical approaches based on continuum mechanics, that is, in many cases build on top of high-fidelity methods which we target within this work. In clinical decision making, treatment planning, prototyping, testing on (virtual) cohorts or in other patient-specific scenarios, costs in resources, patient risk and ethical concerns can be mitigated, while traditional in-vivo studies and procedures are more intrusive in multiple ways. As numerical methods continue to advance, the developed computational tools can be employed for increasingly complex tasks, where individual mathematical models coupled with one another may have an even stronger impact on the overall algorithm in terms of efficiency, accuracy, robustness and naturally in terms of computational results obtained.

In medicine and biomedical engineering, many of these applications feature incompressible flow interacting with soft biological tissue, leading to an FSI problem. Computational methods can in this regard be seen as an additional powerful tool to assist clinicians and gain further insights into the initiation, progression or possible treatment of life-threatening diseases. It comes as no surprise that the applicability has to be questioned in many practical scenarios, especially given the high risk at stake. However, reliable,

profound and risk-based recommendations for clinical support are as of now only available for a few selected problems, but are currently very active areas of research. Today, many impressive patient-specific case studies focusing on various different cardiovascular diseases, taking for example aortic dissection (see, e.g., [266, 385–388] among others) or abdominal aortic aneurysma (see, for example [389–392]) can be found in literature, while non-invasive parameter identification and propagation of uncertainties through the computationally expensive models is becoming more and more feasible given robust and fast solution algorithms for the individual subproblems and FSI [393–398].

FSI schemes and treatment of flows on moving domains as a central part of any FSI solver are frequently differentiated by the methodology used for resolving the moving and deforming interface or fluid domain boundary. Employing an Arbitrary Lagrangian–Eulerian (ALE) formulation is one of the many available alternatives often resorted to. Updating the current fluid domain and adapting the balance equations accordingly to account for the mapping allows considering a mapped fluid domain and therefore *tracking* the interface or boundary motion. A vast pool of studies found in literature utilise ALE methods (see, e.g., [187–194] among others), mostly for the main advantages which are (i) the possibility to design high-quality meshes including boundary layers capturing the expected high velocity gradients, (ii) ease of implementation only requiring an update of the computational grid’s position and minor changes required to account for the ALE formulation, and (iii) reuse of numerical methods and software developed for fluid flow solvers on stationary meshes. Due to these advantages, we herein introduce an ALE framework based on the PPE-based split-step scheme in ALE formulation as presented in Sec. 4.1.

In contrast to the ALE framework, the interface or boundary treatment in fluid flow simulations on moving domains might also be accomplished *capturing* the flow domain’s boundary applying immersed boundary methods [195–200], fictitious-domain methods [201–204] or XFEM [399–403] just to name the three most prominent available options. Interface-capturing techniques enforce the boundary conditions for the fluid subproblem without the need for resolving the domain’s boundary by the computational grid. Thus, a fixed background grid can be employed, enabling even topology changes of the fluid domain. On the other side, the fluid domain map is the Achilles’ heel of most ALE methods, which might break down or suffer from bad element aspect ratios present in the mapped grid under severe distortion, necessitating remeshing of the fluid domain and mapping data between grids. Despite us settling for an ALE approach, it shall still be mentioned that various interface-capturing techniques are applicable to the PPE-based split-step scheme to solve fluid flow problems on moving domains in a straight-forward manner as well, adapting the imposition of boundary conditions accordingly, which depends on the specific algorithm chosen. As long as the mesh quality can be kept at a reasonable level (not or only scarcely performing remeshing and mapping data), the ALE techniques have the advantage that standard methods can be applied to enforce boundary conditions, no bad cut scenarios must be handled, stabilisations to prevent severe ill-conditioning of the linear system are avoided and suitable boundary layers following the fluid domain motion can be employed.

On the fluid–structure interface, coupling conditions being the continuity of tractions, displacements and velocities are enforced following either monolithic or partitioned ap-

proaches. In monolithic schemes, all involved balance equations (of the fluid *and* structure) are assembled into one single system of equations, gathering contributions from all involved variables into the same matrix and right-hand side, typically within a Newton–Krylov approach. This enforces an inherent tight coupling of physical fields, but unfortunately comes with an increased implementation effort, possibly unusual data structures and more involved preconditioners. On the other hand, such an approach ensures the coupling conditions at each nonlinear iteration, since balance equations are solved simultaneously. Several variants to enforce the interface conditions in monolithic schemes include, e.g., Langrange multipliers [380, 399, 404–406], Nitsche’s method [407–409], mortar techniques [200, 401, 410], penalty approaches [411–413] or formulations enforcing interface conditions strongly through tailored function spaces and posing fluid and structure problems in a joint reference frame [9, 414–418].

Preconditioning the resulting monolithic linearised system for efficient solution independent of grid resolution and physical parameters is a challenging and highly active field of research, given that solving the fluid subproblem at high Reynolds numbers on its own is a complex matter. Based on the block structure of the assembled system, physics-based preconditioners often incorporating approximations of block-inverses via multigrid cycles can be designed taking the chosen setup into account (see, e.g., the works by [108, 187, 380, 404, 406, 417, 419–424]).

Partitioned coupling schemes, on the other hand, alleviate the development of efficient preconditioners by iteratively enforcing the interface conditions, repeatedly solving the involved subproblems and exchanging the updated interface data. They allow reusing already advanced state-of-the-art solution algorithms and preconditioners tailored to the specific subproblems and/or application and ease the inclusion of complex physics or discretisation methods in the individual fields [425–428]. Using separate solvers with clearly defined software interfaces and exchanging updates to the interface variables also allows simpler software design, where individual parts and solvers are exchanged with much less effort compared to a monolithic approach. This enables a black-box principle, where the individual subproblem solvers are treated as nonlinear functions mapping input to output interface data, ensuring convergence via an outer coupling loop without using any knowledge regarding the actual subproblem solver.

However, the partitioned paradigm shifts some of the intricacy to the outer coupling procedure—which can to some extent be interpreted as an advantage, since potent acceleration schemes are available. At the same time, this introduces an additional layer of numerical techniques and thereby further complicates matters. In the context of biomedical applications, the added-mass effect [365, 366, 429] is a well-known problem, notorious for severely impairing the coupling scheme’s performance or even hindering convergence completely for small time step sizes or a large fluid to solid density ratio ρ_f/ρ_s . Standard partitioned schemes such as the loosely-coupled, serial staggered method [367] and even implicit partitioned methods suffer from decreased convergence speed as demonstrated, e.g., in [88, 368, 369]. To counteract the hampered convergence behaviour in such cases, a number of remedies have been presented ranging from simple relaxation of the exchanged interface data with a fixed parameter [189] or via dynamic Aitken relaxation [88], over approaches targeting the nonlinear interface problem via interface (quasi-)Newton or Newton–Krylov solvers [430–436], to methods based

on Robin interface conditions [217, 218, 220] or introducing artificial/weak compressibility [221, 437–440].

Despite these tremendous research efforts and marvelous success over the past decade, performing several—typically 5 to 15—coupling iterations between fluid and solid subproblems per time step nonetheless remains costly despite the massive performance gains achieved with the currently available acceleration schemes. But not all involved fields have an equally strong effect on the coupling, and might thus be linearised, treated explicitly or pulled out of the outer coupling loop of the partitioned scheme, such that some subproblems are only solved once per time step. This fundamental idea greatly reduces the computational burden and was first introduced by Fernández et al. [441], leading to *semi-implicit* FSI algorithms. In such schemes, the fluid flow problem is solved via a projection or split-step method, decoupling the fluid’s velocity and pressure unknowns. Additionally, the fluid domain deformation on the interface is extrapolated from previous time steps, which then allows for coupling only the fluid pressure and solid deformation implicitly. This brilliant adaptation allows added-mass stable coupling of FSI problems in the biomedical regime ($\rho_f/\rho_s \approx 1$), since the fluid pressure is the dominant loading. Following this approach, most expensive problems being the possibly nonlinear, vector-valued fluid momentum balance and fluid mesh motion equations only need to be solved once per time step. Doing so does neither deteriorate accuracy nor stability and interestingly, almost identical results are obtained from implicit and semi-implicit partitioned strong coupling schemes.

Based on the rationale of semi-implicitly coupling some of the involved fields, rapid development was seen in the following years (see, e.g., [192, 225, 442–448]). In numerous challenging settings, these techniques were demonstrated to yield accurate and stable results while substantially increasing performance. Taking a further step in this direction, *fully explicit* treatment of interface conditions in FSI for problems with large added-mass effect was presented for shells (see, e.g., [219, 448–452]) and also for three-dimensional continua [222, 453–456]. Unfortunately, such fully explicit coupling schemes for bulk solids are to the best of the author’s knowledge as of now either limited to simple constitutive behaviour or only first-order accurate in time unless multiple correction steps are performed. Additionally, they most often rely on Robin transmission conditions with a user-defined parameter, which can strongly influence the physics of the FSI problem when set incorrectly and might necessitate even smaller time step sizes to appropriately capture transient solutions. Note here, however, that fully explicit FSI is a rapidly progressing research topic and these statements might need to be weakened or completely revoked soon. Additionally, the increase in performance and robustness comes at the price of partially giving up on the black-box design principle, or at least re-interpreting it in a more general way, allowing for several (“matching”) black-box solvers recovering the various fields’ unknowns separately.

Comparing monolithic and partitioned schemes in terms of performance and robustness is a delicate matter for several convoluted reasons. The two most striking ones being that first, constructing state-of-the-art FSI solvers is a complex task and potent software packages evolve over years if not decades, mostly following either the one or the other paradigm per design. And secondly, rapid developments in the involved fields cause the applied numerical methods being out of date very fast, such that using the same

subproblem solvers for the block inversion in physics-based preconditioned monolithic approaches and partitioned schemes might work against a fair comparison. Nonetheless, several studies aim at a comparison of the monolithic and partitioned design principles, see, e.g., [353, 405, 425, 435, 457–460], where monolithic schemes have been demonstrated to be more robust in scenarios where the added-mass-effect or the subproblems' interactions are strong. In scenarios where these effects are not dominant or when using one or multiple acceleration methods combined, partitioned schemes can be competitive. If the number of FSI coupling iterations per step can be kept at a reasonable level, meaning, around 5 to 15, the highly tailored subproblem solvers might in fact lead to a significant reduction in runtime for each of the problems, such that the potentially increased iteration count compared to the number of Newton steps and (block-)preconditioner applications in a monolithic scheme might be more than compensated. Since there is also a large body of literature dealing with comparisons of the various acceleration schemes (see, e.g. [5, 88, 217, 218, 351–353, 362, 425, 432, 436, 441, 459, 461, 462] and many others), comparing the multitude of monolithic and partitioned schemes with each other thus remains cumbersome and highly dependent on the targeted flow regime and physical problem.

However, the following statement still seems in place: *in the past*, monolithic approaches seemed to be the only reasonable way to tackle biomedical FSI applications. *Today*, with latest semi-implicit partitioned coupling approaches based on (accelerated) Robin coupling, the number of iterations and robustness could be significantly improved, while at the same time promising fully explicit schemes are being developed. Together with the reduced complexity, increased modularity and the possibility to employ tailored high-performance subproblem solvers in partitioned methods, we find that the decision of which coupling approach to employ is recently strongly shifting in favour of such partitioned schemes.

Although FSI applied to numerous engineering applications has been a major research topic in the past decade, approaches specifically targeting generalised Newtonian fluid flow are still scarcely found in literature [450, 463–466]. And while partitioned FSI approaches and some of the performant acceleration schemes do allow a straight-forward incorporation of complex constitutive laws by the black-box design, the promising concept of semi-implicit FSI builds upon methods decoupling fluid velocity and pressure, which then need to be capable of considering non-Newtonian models. Within such fluid solvers segregating velocity and pressure, necessary projection methods need to be free of nonphysical pressure boundary layers, rendering the PPE-based time splitting scheme presented in Ch. 3 particularly well fit.

In this context, we herein present a summary and extension of the recently published papers by the authors [5, 6], building the basis for this chapter and the corresponding numerical experiments in Ch. 8. One of the main contributions of this thesis is the novel split-step framework for partitioned, (semi-)implicit, Robin–Robin and Aitken or interface quasi-Newton-accelerated FSI scheme involving incompressible, generalised Newtonian fluid flows combining the PPE-based time splitting scheme with nearly incompressible hyperelastic and fibre-reinforced continua. Thus, we inherit the favourable properties of the split-step scheme and exploit that the fluid's velocity and pressure fields are completely decoupled using higher-order and possibly adaptive timestepping

and suitable extrapolation formulae. The flow solver further allows employing equal-order, standard C^0 -continuous finite elements, since the PPE-based fluid solver is not subject to the classical inf-sup condition as discussed in detail in Ch. 3. As shown in Sec. 4.1.2, fully consistent boundary and coupling conditions for the fluid subproblems are thus easily derived, preserving accuracy on the fluid–structure interface.

Consequently, semi-implicit and added-mass stable variants of the FSI scheme can be designed following Fernández et al. [441], implicitly, i.e., iteratively, coupling merely the solid displacement and fluid pressure. The remaining subproblems are treated in an explicit fashion, which enhances performance without degrading accuracy or stability. Moreover, divergence damping to improve mass conservation is applied in an identical fashion as shown for the fluid problem alone. The flexibility concerning the rheological law is also preserved, such that modifying the fluid material behaviour is yet again as simple as exchanging the right-hand side of the viscosity projection step.

Then, we consistently integrate well-established mathematical models and numerical techniques needed in biomedical applications into the PPE-based FSI scheme, preserving accuracy, efficiency and stability. Such overall frameworks are scarcely found in literature, with closely related frameworks by Astorino et al. [225] and Bertoglio et al. [267] being exceptions. The techniques to reduce instabilities arising in the practically relevant parameter range are (i) backflow stabilisation [236] (Sec. 4.5), and (ii) Galerkin Least-Squares stabilisation [33] (GLS) to counteract dominant convection (Sec. 4.3). To increase convergence speed as the added-mass effect becomes dominant, we consider *combined* (iii) Robin interface conditions [218, 219] (Secs. 4.2 and 6.3) and (iv) the Interface Quasi-Newton Inverse Least-Squares method [435] (IQN-ILS). Further models in the cardiovascular and general biomedical context considered are (i) three-element Windkessel models to account for the downstream vasculature’s resistance and capacitance (Sec. 4.6), (ii) remapping of prescribed inflow profiles to non-circular inlets [228] (Sec. 4.4), (iii) construction of material orientation vectors in tube-like geometries with subdivided cross-sections (Sec. 6.1), (iv) an algorithm to include prestress present in geometries constructed from medical images [318] (Sec. 6.4) and (v) viscoelastic support provided by the surrounding external tissue via Robin boundary conditions (Sec. 6.2).

Altogether, this leads to a robust framework involving all needed modelling aspects in the cardiovascular setting and can also be applied to other challenging problems in the general (bio-)medical context. The building blocks composing the FSI solver lead to standard problems in science and engineering, such that the resulting linear systems are easily tackled using off-the-shelf black-box preconditioning techniques available as open-source scientific software and thus bear high potential for further development, while this current work does not make use of massive parallelism.

7.1 Governing equations

The computational domain in the Lagrangian reference frame $\hat{\Omega}$ is composed of fluid and structure subdomains $\hat{\Omega}_f$ and $\hat{\Omega}_s$, respectively, such that $\hat{\Omega} = \hat{\Omega}_f \cup \hat{\Omega}_s$. In the spatial configuration at time t , we denote the two subdomains by Ω_f^t and Ω_s^t , where $\Omega = \Omega_f^t \cup \Omega_s^t$

holds accordingly. The joint interface of fluid and structure domains in the material and spatial configuration is then

$$\hat{\Sigma} := \partial\hat{\Omega}_f \cap \partial\hat{\Omega}_s \quad \text{and} \quad \Sigma^t := \partial\Omega_f^t \cap \partial\Omega_s^t,$$

Adopting an Arbitrary Lagrangian–Eulerian (ALE) method (cf. Sec. 4.1 or [188, 207]) to account for the motion of Ω_f^t and a total Lagrangian approach (cf. Ch. 5 or [276, 319]) to consider the structure’s balance equations in $\hat{\Omega}_s$, the deformations \mathbf{d}_f and \mathbf{d}_s from the Lagrangian reference frames $\hat{\Omega}_f$ and $\hat{\Omega}_s$ are introduced, such that

$$\phi_f : \hat{\mathbf{x}} \in \hat{\Omega}_f \rightarrow \mathbf{x} \in \Omega_f^t, \quad \phi_f(\hat{\mathbf{x}}, t) := \hat{\mathbf{x}} + \mathbf{d}_f(\hat{\mathbf{x}}, t), \quad (7.1)$$

$$\phi_s : \hat{\mathbf{x}} \in \hat{\Omega}_s \rightarrow \mathbf{x} \in \Omega_s^t, \quad \phi_s(\hat{\mathbf{x}}, t) := \hat{\mathbf{x}} + \mathbf{d}_s(\hat{\mathbf{x}}, t), \quad (7.2)$$

yield the corresponding spatial coordinates $\mathbf{x}(\hat{\mathbf{x}}, t)$ in $\hat{\Omega}$. In the structural subdomain $\hat{\Omega}_s$, the material displacement \mathbf{d}_s is a physical quantity of interest, while in $\hat{\Omega}_f$, the displacement of the domain \mathbf{d}_f is completely artificial. The field \mathbf{d}_f may be given on the entirety of $\partial\hat{\Omega}_f$ or on (at least) some part of $\partial\hat{\Omega}_f$ depending on the interface conditions enforced and the boundary conditions on the remaining $\partial\hat{\Omega}_f \setminus \hat{\Sigma}$. In both of these cases, the domain displacement needs to be extended into the interior of the fluid domain, considering for a prescribed motion of the boundary or taking the motion of the fluid–structure interface into account. For any ALE method, the construction of the ALE map (7.1) is a *crucial* ingredient, since non-invertible or low-quality meshes resulting from the extension lead to unreliable and erroneous results.

Similar to a classical total Lagrangian formulation, we can then define the deformation gradients \mathbf{F} and corresponding Jacobians J in each of the subdomains as

$$\begin{aligned} \mathbf{F}_f &:= \hat{\nabla}\phi_f = \mathbf{I} + \hat{\nabla}\mathbf{d}_f, & J_f &= \det \mathbf{F}_f, \\ \mathbf{F}_s &:= \hat{\nabla}\phi_s = \mathbf{I} + \hat{\nabla}\mathbf{d}_s, & J_s &= \det \mathbf{F}_s, \end{aligned}$$

where \mathbf{F}_s is again used to derive strain measures since ϕ_s follows the material motion. The map ϕ_f yields the spatial grid coordinates, and hence \mathbf{F}_f is merely descriptive to the mapping. Here, we again use the $\hat{\cdot}$ symbol to refer to the Lagrangian reference frame $\hat{\Omega}$, such that for a function $g(\mathbf{x}, t)$ with $\mathbf{x} \in \Omega^t$, its counterpart $\hat{g}(\hat{\mathbf{x}}, t)$ with $\hat{\mathbf{x}} \in \hat{\Omega}$ is given as $\hat{g}(\hat{\mathbf{x}}, t) = g(\mathbf{x}, t)$. However, if the meaning clearly derives from the definition, we omit the $\hat{\cdot}$ to not clutter the notation (e.g., \mathbf{d}_s , not $\hat{\mathbf{d}}_s$).

With the mappings defined, we can further proceed with formulating the interface coupling conditions, namely the kinematic coupling conditions, enforcing continuity of displacements and velocities,

$$\mathbf{d}_f = \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad \mathbf{u}_f = \mathbf{d}_t \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad (7.3)$$

and the kinetic coupling conditions equating tractions,

$$\boldsymbol{\sigma}_f \mathbf{n}_f = -\boldsymbol{\sigma}_s \mathbf{n}_s \quad \text{on } \Sigma^t, \quad \mathbf{P}_f \hat{\mathbf{n}}_f = -\mathbf{P}_s \hat{\mathbf{n}}_s \quad \text{on } \hat{\Sigma}. \quad (7.4)$$




Fig. 7.1: Nomenclature used for boundary segments in FSI, exemplarily shown for the aortic dissection case: overall geometry (a) and zoom-in on arch, showing cut structural domain $\hat{\Omega}_s$ with adventitia (red), media (peach) and flap (light blue) regions of the tissue (b) and discretised lumen $\hat{\Omega}_f$ (c).

The traction conditions may be recast with the help of Nanson's formula $\mathbf{P} = J\boldsymbol{\sigma}\mathbf{F}^{-\top}$ given in Eqn. (5.7), see Sec. 5.1, to yield relations on either $\hat{\Sigma}$ or Σ^t as

$$\boldsymbol{\sigma}_f \mathbf{n}_f = J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f \quad \text{on } \Sigma^t, \quad (7.5)$$

$$\mathbf{P}_s \hat{\mathbf{n}}_s = J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s \quad \text{on } \hat{\Sigma}, \quad (7.6)$$

depending on which data is easiest available. Robin interface conditions are derived combining the kinematic (7.3) and kinetic interface conditions (7.5)–(7.6) (see, e.g., [5, 6, 217–225] among others), which yields

$$\eta_f^R \mathbf{u}_f + \boldsymbol{\sigma}_f \mathbf{n}_f = \eta_s^R d_t \mathbf{d}_s + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f \quad \text{on } \Sigma^t, \quad (7.7)$$

$$\eta_s^R d_t \mathbf{d}_s + \mathbf{P}_s \hat{\mathbf{n}}_s = \eta_f^R \mathbf{u}_f + J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s \quad \text{on } \hat{\Sigma}. \quad (7.8)$$

Here, the scaling parameters η_s^R and η_f^R smoothly transition from the continuity of velocities ($\eta_s^R \rightarrow \infty$, $\eta_f^R \rightarrow \infty$) to traction balance ($\eta_s^R \rightarrow 0$, $\eta_f^R \rightarrow 0$) being enforced on the fluid–structure interface. Such Robin conditions are consistent, since for the exact solution of the FSI problem, both of the combined coupling conditions are fulfilled exactly. When linearising the problem, however, the enforced interface condition is merely approximated, converging to the correct solution as $h_e \rightarrow 0$ and $\Delta t \rightarrow 0$, as is thoroughly discussed in Sec. 7.2. Also, note here that the Robin interface conditions in FSI naturally lead to Robin conditions on the fluid–structure interface in the fluid and solid subproblems, such that the elaborations in Sec. 4.2 and 6.3 are directly applicable, especially in the case of partitioned coupling schemes being employed.

Regarding nomenclature, we will stick to the preceding sections, but omit the time interval in the domain specifications for brevity (e.g., “on Σ^t ” instead of “on $\Sigma^t \times (0, T]$ ”). Also, the problem formulations in this chapter target haemodynamic applications regarding the sets of boundary conditions, but are equally valid adopting the various other combinations of boundary conditions of the respective subproblems introduced in the respective Chs. 4 and 5 as long as they lead to a well-posed boundary value problem. A possible configuration in the context of aortic dissection is shown in Fig. 7.1, representative for the cardiovascular setting. Simpler problem definitions trivially derive from the presented scenario.

For the solid problem, we will prescribe displacements on the inlet and outlet faces, $\hat{\Gamma}_{\text{in},s}$ and $\hat{\Gamma}_{i,s}$, $i = 1, \dots, N_{\text{out}}$, summarised in $\hat{\Gamma}_{\text{out},s}$, such that

$$\hat{\Gamma}_{\text{out},s} := \bigcup_{i=1}^{N_{\text{out}}} \hat{\Gamma}_{i,s}, \quad \text{and} \quad \hat{\Gamma}_{D,s} := \hat{\Gamma}_{\text{in},s} \cup \hat{\Gamma}_{\text{out},s},$$

composes the Dirichlet boundary of the structure subproblem, where deformations are prescribed, i.e., $\mathbf{d}_s = \mathbf{g}_s$ holds. On the structure’s exterior, denoted by $\hat{\Gamma}_{\text{ext}}$, a Robin condition is enforced, namely,

$$\mathbf{P}_s \hat{\mathbf{n}}_s = -k_e \mathbf{d}_s - c_e \mathbf{d}_t \mathbf{d}_s - p_e \hat{\mathbf{n}}_s,$$

choosing $\hat{\Gamma}_{R,s} = \hat{\Gamma}_{\text{ext}}$ to consider for viscoelastic support and/or prescribing an exterior pressure (cf. Sec. 6.2). Alternatively, prescribing a first Piola–Kirchhoff traction via $\mathbf{P}_s \hat{\mathbf{n}}_s = \hat{\mathbf{t}}$, i.e., a Neumann condition might be considered setting $\hat{\Gamma}_{N,s} = \hat{\Gamma}_{\text{ext}}$, and is in fact slightly simpler since $\hat{\mathbf{t}}(\hat{\mathbf{x}}, t)$ is a given function.

As thoroughly elaborated in Sec. 5.1, the nonlinear initial boundary value problem involving the structure’s momentum balance equation in $\hat{\Omega}_s$ with the above mentioned boundary conditions can then be written as

$$\rho_s \mathbf{d}_{tt} \mathbf{d}_s - \hat{\nabla} \cdot \mathbf{P}_s = \hat{\mathbf{b}}_s \quad \text{in } \hat{\Omega}_s, \quad (7.9)$$

$$\mathbf{d}_s = \mathbf{g}_s \quad \text{on } \hat{\Gamma}_{D,s}, \quad (7.10)$$

$$\mathbf{P}_s \hat{\mathbf{n}}_s = -k_e \mathbf{d}_s - c_e \mathbf{d}_t \mathbf{d}_s - p_e \hat{\mathbf{n}}_s \quad \text{on } \hat{\Gamma}_{\text{ext}}, \quad (7.11)$$

$$\mathbf{P}_s \hat{\mathbf{n}}_s = \eta_s^R (\mathbf{u}_f - \mathbf{d}_t \mathbf{d}_s) + J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s \quad \text{on } \hat{\Sigma}, \quad (7.12)$$

$$\mathbf{d}_s(\hat{\mathbf{x}}, 0) = \mathbf{d}_{s,0} \quad \text{at } t = 0, \quad (7.13)$$

$$\mathbf{d}_t \mathbf{d}_s(\hat{\mathbf{x}}, 0) = \dot{\mathbf{d}}_{s,0} \quad \text{at } t = 0, \quad (7.14)$$

compare with Eqns. (5.9)–(5.13), where $\mathbf{d}_{s,0}$ and $\dot{\mathbf{d}}_{s,0}$ are given initial material displacement and velocity fields.

Similarly, the fluid subdomain’s boundary $\partial\Omega_f^t$ is decomposed into the inlet face $\Gamma_{\text{in},f}^t$, the fluid–structure interface Σ^t and N_{out} individual outlet faces, $\Gamma_{i,f}$, $i = 1, \dots, N_{\text{out}}$, for convenience gathered in $\Gamma_{\text{out},f}^t$, such that

$$\partial\Omega_f^t = \Gamma_{\text{in},f}^t \cup \Sigma^t \cup \Gamma_{\text{out},f}^t, \quad \text{with} \quad \Gamma_{\text{out},f}^t := \bigcup_{i=1}^{N_{\text{out}}} \Gamma_{i,f}.$$

The respective counterparts in the Lagrangian reference frame are indicated by a $\hat{\cdot}$ as shown in Fig. 7.1 and coincide with the spatial counterparts at initial time $t = 0$.

The ALE extension aims at obtaining a smooth field \mathbf{d}_f in $\hat{\Omega}_f$, given mesh displacements on the interface $\hat{\Sigma}$ and on the boundaries, such that an invertible, one-to-one mapping from $\hat{\Omega}_f$ to Ω_f^t can be constructed via $\mathbf{x}(\hat{\mathbf{x}}, t) = \hat{\mathbf{x}} + \mathbf{d}_f(\hat{\mathbf{x}}, t)$, see Eqn. (7.1). A standard choice is to consider for the arbitrary auxiliary mesh motion equation a simple harmonic extension with pure Dirichlet boundary conditions,

$$-\hat{\nabla} \cdot \hat{\nabla} \mathbf{d}_f = \mathbf{0} \quad \text{in } \hat{\Omega}_f, \quad (7.15)$$

$$\mathbf{d}_f = \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad (7.16)$$

$$\mathbf{d}_f = \mathbf{0} \quad \text{on } \Gamma_{\text{in},f}^t \cup \Gamma_{\text{out},f}^t, \quad (7.17)$$

being the most trivial choice. For the sake of presentation, Eqns. (7.15)–(7.17) suffice, but may fail in more complex scenarios.

In this regard, possible improvements can be achieved by enforcing merely the normal component of the displacement, i.e., $\mathbf{d}_f \cdot \mathbf{n}_f$ on the whole boundary $\partial\hat{\Omega}_f$, or replacing the mesh motion equation (7.15) by more elaborated versions or techniques specifically designed to preserve mesh quality [467–473]. Sliding subdomains complicate transferring and enforcing interface data in an ALE approach—especially when employing domain decomposition methods. Herein, we thus restrict ourselves to the standard displacement boundary condition on the full grid displacement vector \mathbf{d}_f . The second option to facilitate the ALE extension, however, received great attention in the past, leading to numerous alternatives to the prototypical harmonic extension being available (see, e.g., [467, 470–475]). For the sake of concision, any further remarks on this topic are postponed to Sec. 7.5, where some of the available options are compared.

Regardless of the mesh motion technique used to construct the fluid’s domain displacement \mathbf{d}_f , the mesh velocity is given by Eqn. (4.5),

$$\mathbf{u}_m(\tilde{\mathbf{x}}, t) := \partial_t \mathbf{x}|_{\tilde{\mathbf{x}}} = \partial_t \mathbf{d}_f(\hat{\mathbf{x}}(\tilde{\mathbf{x}}, t))|_{\tilde{\mathbf{x}}},$$

with mesh displacements available in grid points. Then, adopting the standard Navier–Stokes system in convective ALE form, the coupled problem reads

$$\rho_f \partial_t \mathbf{u}_f|_{\bar{\mathbf{x}}} + \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f - \mu_f [\Delta \mathbf{u}_f + \chi \nabla (\nabla \cdot \mathbf{u}_f)] - 2\nabla^s \mathbf{u}_f \nabla \mu + \nabla p_f = \mathbf{b}_f \quad \text{in } \Omega_f^t, \quad (7.18)$$

$$\nabla \cdot \mathbf{u}_f = 0 \quad \text{in } \Omega_f^t, \quad (7.19)$$

$$\mathbf{u}_f = \mathbf{g}_f \quad \text{on } \Gamma_{\text{in},f}^t, \quad (7.20)$$

$$(-p_f \mathbf{I} + \mu_f \nabla \mathbf{u}_f + \chi \nabla^\top \mathbf{u}_f) \mathbf{n}_f = \bar{\mathbf{t}}_f \quad \text{on } \Gamma_{\text{out},f}^t, \quad (7.21)$$

$$\eta_f^R (\mathbf{d}_t \mathbf{d}_s - \mathbf{u}_f) + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f = \boldsymbol{\sigma}_f \mathbf{n}_f \quad \text{on } \Sigma^t, \quad (7.22)$$

$$\mathbf{u}_f = \mathbf{u}_{f,0} \quad \text{at } t = 0, \quad (7.23)$$

with the fluid velocity and pressure, \mathbf{u}_f and p_f , the dynamic viscosity μ_f , body force $\mathbf{b}_f(\mathbf{x}, t)$, density ρ_f , convective velocity $\mathbf{c}_f := \mathbf{u}_f - \mathbf{u}_m$ and the scaling parameter χ to yield either the generalised Laplacian ($\chi = 0$) or stress-divergence form ($\chi = 1$), compare Eqns. (4.6) and (4.7)–(4.11). Dirichlet conditions are imposed on the inlet $\Gamma_{\text{in},f}^t$, while on the outlet faces, $\Gamma_{\text{out},f}^t$, the real or pseudo traction vector $\bar{\mathbf{t}}$ is prescribed. The latter condition can be easily replaced, incorporating models of the downstream vasculature and/or backflow stabilisation as discussed in Secs. 4.5 and 4.6, but is omitted here for the sake of brevity.

Note that the interface Robin condition involves the Cauchy traction $\boldsymbol{\sigma}_f \mathbf{n}_f$, which does not occur in a boundary term obtained via integration by parts when the generalised Laplacian form ($\chi = 0$) is employed. The Neumann condition on the outlet might also necessitate the use of the stress-divergence form, which largely depends on the desired problem at hand. Switching Eqn. (7.22) for a pure Dirichlet condition ($\eta_f^R \rightarrow \infty$), these complications can be resolved. An alternative also allowing for the Robin interface condition (7.22) consists of adding interface integrals, accounting for the difference between real and pseudo tractions, $\mu \nabla^\top \mathbf{u}_f$, as briefly discussed in Ch. 2.

In analogy to Secs. 4.1.2 and 4.2, system (7.18)–(7.23) can be reformulated based on a corresponding PPE with Robin data $\mathbf{h}_f = \eta_f^R \mathbf{d}_t \mathbf{d}_s + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f$ given on Σ^t and considering for the moving domain Ω_f^t , which reads

$$\rho_f \partial_t \mathbf{u}_f|_{\bar{\mathbf{x}}} + \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f - \mu_f [\Delta \mathbf{u}_f + \chi \nabla (\nabla \cdot \mathbf{u}_f)] - 2\nabla^s \mathbf{u}_f \nabla \mu + \nabla p_f = \mathbf{b}_f \quad \text{in } \Omega_f^t, \quad (7.24)$$

$$-\Delta p_f + \nabla \cdot [2\nabla^s \mathbf{u}_f \nabla \mu_f - \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f] - [\nabla \times (\nabla \times \mathbf{u}_f)] \cdot \nabla \mu_f = -\nabla \cdot \mathbf{b}_f \quad \text{in } \Omega_f^t, \quad (7.25)$$

$$\mathbf{u}_f = \mathbf{g}_f \quad \text{on } \Gamma_{\text{in},f}^t, \quad (7.26)$$

$$(-p_f \mathbf{I} + \mu_f \nabla \mathbf{u}_f + \chi \mu_f \nabla^\top \mathbf{u}_f) \mathbf{n}_f = \bar{\mathbf{t}}_f \quad \text{on } \Gamma_{\text{out},f}^t, \quad (7.27)$$

$$\eta_f^R (\mathbf{d}_t \mathbf{d}_s - \mathbf{u}_f) + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f = \boldsymbol{\sigma}_f \mathbf{n}_f \quad \text{on } \Sigma^t, \quad (7.28)$$

$$-\mu_f \nabla \cdot \mathbf{u}_f + \mathbf{n}_f \cdot [\mu_f (\nabla \mathbf{u}_f + \chi \nabla^\top \mathbf{u}_f) \mathbf{n}_f - \bar{\mathbf{t}}_f] = p_f \quad \text{on } \Gamma_{\text{out},f}^t, \quad (7.29)$$

$$\mathbf{n}_f \cdot [\mathbf{b}_f - \rho_f \mathbf{d}_t \mathbf{d}_s - \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f + 2\nabla^s \mathbf{u}_f \nabla \mu_f - \mu_f \nabla \times (\nabla \times \mathbf{u}_f)] = \mathbf{n}_f \cdot \nabla p_f \quad \text{on } \Sigma^t, \quad (7.30)$$

$$\mathbf{n}_f \cdot [\mathbf{b}_f - \rho_f \partial_t \mathbf{u}_f|_{\bar{\mathbf{x}}} - \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f + 2\nabla^s \mathbf{u}_f \nabla \mu_f - \mu_f \nabla \times (\nabla \times \mathbf{u}_f)] = \mathbf{n}_f \cdot \nabla p_f \quad \text{on } \Gamma_{\text{in},f}^t, \quad (7.31)$$

$$\mathbf{u}_f = \mathbf{u}_{f,0} \quad \text{at } t = 0, \quad (7.32)$$

similar to Eqns. (4.51)–(4.59) and with a divergence-free initial velocity field $\mathbf{u}_{f,0}$.

Notice, however, that the treatment of Robin interface conditions in the momentum balance and PPE steps is not simply assigning Robin interface conditions to the whole fluid subproblem, but rather enforcing Robin interface conditions in the fluid momentum

equation (7.28) and at the same time enforcing $\mathbf{u}_f = \mathbf{d}_t \mathbf{d}_s$ in the PPE (7.30), leading to a Neumann condition for the pressure on the interface. This combination lead to significant improvements in temporal stability of the scheme compared to a scheme enforcing Robin interface conditions within the PPE step and is within this work solely based on numerical observations. Also, the proposed strategy is comparable to the one adopted by Astorino et al. [225], who use a projection solver in semi-implicit FSI based on Robin coupling and Nitsche's method.

Omitting now the initial- and boundary conditions on the inlet and outlet sections in all subproblems for brevity, the FSI problem reads

$$\rho_s \mathbf{d}_{tt} \mathbf{d}_s - \hat{\nabla} \cdot \mathbf{P}_s = \hat{\mathbf{b}}_s \quad \text{in } \hat{\Omega}_s, \quad (7.33)$$

$$\rho_f \partial_t \mathbf{u}_f|_{\hat{x}} + \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f - \mu_f [\Delta \mathbf{u}_f + \chi \nabla (\nabla \cdot \mathbf{u}_f)] - 2 \nabla^s \mathbf{u}_f \nabla \mu + \nabla p_f = \mathbf{b}_f \quad \text{in } \Omega_f^t, \quad (7.34)$$

$$\nabla \cdot \mathbf{u}_f = 0 \quad \text{in } \Omega_f^t, \quad (7.35)$$

$$-\hat{\nabla} \cdot \hat{\nabla} \mathbf{d}_f = \mathbf{0} \quad \text{in } \hat{\Omega}_f, \quad (7.36)$$

$$\eta_s^R (\mathbf{u}_f - \mathbf{d}_t \mathbf{d}_s) + J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s = \mathbf{P}_s \hat{\mathbf{n}}_s \quad \text{on } \hat{\Sigma}, \quad (7.37)$$

$$\eta_f^R (\mathbf{d}_t \mathbf{d}_s - \mathbf{u}_f) + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f = \boldsymbol{\sigma}_f \mathbf{n}_f \quad \text{on } \Sigma^t, \quad (7.38)$$

$$\mathbf{d}_f = \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad (7.39)$$

using the Navier–Stokes equations for incompressible flows in their standard convective form resulting in a coupled velocity-pressure formulation of the fluid subproblem as discussed in Sec. 2.2. Alternatively, replacing the standard divergence-free constraint on the fluid velocity (7.35) by a suitable PPE (7.42) with a consistent interface Neumann condition on the pressure (7.44), one obtains the alternative system

$$\rho_s \mathbf{d}_{tt} \mathbf{d}_s - \hat{\nabla} \cdot \mathbf{P}_s = \hat{\mathbf{b}}_s \quad \text{in } \hat{\Omega}_s, \quad (7.40)$$

$$\rho_f \partial_t \mathbf{u}_f|_{\hat{x}} + \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f - \mu_f [\Delta \mathbf{u}_f + \chi \nabla (\nabla \cdot \mathbf{u}_f)] - 2 \nabla^s \mathbf{u}_f \nabla \mu + \nabla p_f = \mathbf{b}_f \quad \text{in } \Omega_f^t, \quad (7.41)$$

$$-\Delta p_f + \nabla \cdot [2 \nabla^s \mathbf{u}_f \nabla \mu_f - \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f] - [\nabla \times (\nabla \times \mathbf{u}_f)] \cdot \nabla \mu_f = -\nabla \cdot \mathbf{b}_f \quad \text{in } \Omega_f^t, \quad (7.42)$$

$$-\hat{\nabla} \cdot \hat{\nabla} \mathbf{d}_f = \mathbf{0} \quad \text{in } \hat{\Omega}_f, \quad (7.43)$$

$$\mathbf{n}_f \cdot [\mathbf{b}_f - \rho_f \mathbf{d}_{tt} \mathbf{d}_s - \rho_f (\nabla \mathbf{u}_f) \mathbf{c}_f + 2 \nabla^s \mathbf{u}_f \nabla \mu_f - \mu_f \nabla \times (\nabla \times \mathbf{u}_f)] = \mathbf{n}_f \cdot \nabla p_f \quad \text{on } \Sigma^t, \quad (7.44)$$

$$\eta_s^R (\mathbf{u}_f - \mathbf{d}_t \mathbf{d}_s) + J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s = \mathbf{P}_s \hat{\mathbf{n}}_s \quad \text{on } \hat{\Sigma}, \quad (7.45)$$

$$\eta_f^R (\mathbf{d}_t \mathbf{d}_s - \mathbf{u}_f) + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f = \boldsymbol{\sigma}_f \mathbf{n}_f \quad \text{on } \Sigma^t, \quad (7.46)$$

$$\mathbf{d}_f = \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad (7.47)$$

interpreting the fluid–structure interface as a Robin boundary for both momentum balance equations in the fluid and solid domains, respectively, and as a Dirichlet boundary for the fluid subproblem's PPE, leading to a Neumann condition on the pressure weakly enforcing $\partial_t \mathbf{u}_f|_{\hat{x}} = \mathbf{d}_{tt} \mathbf{d}_s$ on the fluid–structure interface.

At this stage, both systems (7.33)–(7.39) and (7.40)–(7.47) might be used in a partitioned approach, iterating through the subproblems and updating interface data and the fluid domain position accordingly. Based on the performance comparison in Sec. 3.5.4 and especially motivated by the possibility of semi-implicit coupling schemes enabled by the use of a split-step fluid solver, we herein focus entirely on the system involving the

split-step scheme, i.e., Eqns. (7.40)–(7.47). In addition to expected gains in efficiency, another reason for using the split-step scheme is that studies on similar schemes are rather scarcely found in literature. Investigating the interplay of Robin interface conditions, semi-implicit coupling, additional acceleration methods and modelling approaches frequently adopted in the biomedical context is thus a novel contribution of this work. Before targeting biomedical applications, though, we first introduce the basic variants of the partitioned coupling scheme in the following section.

7.2 Partitioned schemes based on a split-step method

Due to the partitioned design, the presented subproblem solvers can be employed easily, merely exchanging interface data. Therefore, it suffices to present the final formulations as derived in the preceding Secs. 4.2 and 6.3, adopting the interface Robin conditions present in the FSI problem. We discretise the time interval $I_t = (0, T]$ from initial time $t = 0$ to end time $t = T$ by decomposing it into N_t possibly non-uniform time steps denoted by $\Delta t^n = t^{n+1} - t^n$, $n = 1, \dots, N_t$. The structure's material acceleration, velocity and displacement at time $t = t^{n+1}$ are denoted by

$$\ddot{\mathbf{d}}_s^{n+1} := d_{tt} \mathbf{d}_s(\hat{\mathbf{x}}, t^{n+1}), \quad \dot{\mathbf{d}}_s^{n+1} := d_t \mathbf{d}_s(\hat{\mathbf{x}}, t^{n+1}), \quad \mathbf{d}_s^{n+1} := \mathbf{d}_s(\hat{\mathbf{x}}, t),$$

and approximated using Newmark formulae (5.44)–(5.45)

$$\begin{aligned} \ddot{\mathbf{d}}_s^{n+1} &\approx \frac{1}{\beta (\Delta t^n)^2} (\mathbf{d}_s^{n+1} - \mathbf{d}_s^n) - \frac{1}{\beta \Delta t^n} \dot{\mathbf{d}}_s^n + \left(1 - \frac{1}{2\beta}\right) \ddot{\mathbf{d}}_s^n, \\ \dot{\mathbf{d}}_s^{n+1} &\approx \frac{\gamma}{\beta \Delta t^n} (\mathbf{d}_s^{n+1} - \mathbf{d}_s^n) + \left(1 - \frac{\gamma}{\beta}\right) \dot{\mathbf{d}}_s^n + \Delta t^n \left(1 - \frac{\gamma}{2\beta}\right) \ddot{\mathbf{d}}_s^n, \end{aligned}$$

with coefficients given in Tab. 5.1. A generalised- α time integration scheme is employed, integrating the forcing terms via a trapezoidal rule in time. Following the steps in Secs. 5.3, 5.4 and 6.3, a semilinear form $\mathbf{R}(\mathbf{d})(\boldsymbol{\varphi})$ is obtained as

$$\begin{aligned} \mathbf{R}_s(\mathbf{d}_s)(\boldsymbol{\varphi}) &:= \langle \rho_s \boldsymbol{\varphi}, \alpha'_m \ddot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s) + \alpha_m \ddot{\mathbf{d}}_s^n \rangle_{\hat{\Omega}_s} \\ &- \langle \boldsymbol{\varphi}, \alpha'_f \hat{\mathbf{b}}_s^{n+1} + \alpha_f \hat{\mathbf{b}}_s^n \rangle_{\hat{\Omega}_s} + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \mathbf{P}_s(\mathbf{d}_s) + \alpha_f \mathbf{P}_s^n \rangle_{\hat{\Omega}_s} - \langle \alpha_f \boldsymbol{\varphi}, \mathbf{P}_s^n \hat{\mathbf{n}} \rangle_{\hat{\Gamma}_{\text{ext}} \cup \hat{\Sigma}} \\ &+ \langle \alpha'_f \boldsymbol{\varphi}, k_e \mathbf{d}_s + c_e \dot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s) + p_e \hat{\mathbf{n}}_s \rangle_{\hat{\Gamma}_{\text{ext}}} + \langle \alpha'_f \boldsymbol{\varphi}, \eta_s^R \dot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s) - \hat{\mathbf{h}}_s^{n+1} \rangle_{\hat{\Sigma}}, \end{aligned} \quad (7.48)$$

with a vector-valued test function $\boldsymbol{\varphi} \in [X(\hat{\Omega}_f)]^d \subset [H^1(\hat{\Omega}_f)]^d$, $\boldsymbol{\varphi}|_{\hat{\Gamma}_D} = \mathbf{0}$ and the Robin interface data dependent on the fluid's velocity, pressure and domain displacement as $\hat{\mathbf{h}}_s := \eta_s^R \mathbf{u}_f + J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s$, completely analogous to Eqn. (6.23).

Inserting the updated interface data into $\hat{\mathbf{h}}_s$, the Jacobian $J(\mathbf{d}_s)(\boldsymbol{\varphi})$ remains otherwise unchanged, such that each Newton step consists of finding $\delta \mathbf{d}_s^k \in [X(\hat{\Omega}_s)]^d$ with $\delta \mathbf{d}_s|_{\hat{\Gamma}_D} = \mathbf{0}$ given the last iterate $\mathbf{d}_s^{n+1,k}$ of the displacement \mathbf{d}_s^{n+1} at time $t = t^{n+1}$ and the last time steps' solutions \mathbf{d}_s^n , \mathbf{d}_s^n and \mathbf{d}_s^n , such that

$$\begin{aligned} & \langle \rho_s \boldsymbol{\varphi}, \frac{\alpha'_m}{\beta (\Delta t^n)^2} \delta \mathbf{d}_s \rangle_{\hat{\Omega}_s} + \langle \hat{\nabla} \boldsymbol{\varphi}, \alpha'_f \text{D}_d \mathbf{P}(\mathbf{d}_s^{n+1,k}) (\delta \mathbf{d}_s) \rangle_{\hat{\Omega}_s} \\ & + \langle \alpha'_f \boldsymbol{\varphi}, k_e \delta \mathbf{d}_s + \frac{c_e \gamma}{\Delta t^n \beta} \delta \mathbf{d}_s \rangle_{\hat{\Gamma}_{\text{ext}}} + \langle \alpha'_f \boldsymbol{\varphi}, \frac{\eta_s^R \gamma}{\Delta t^n \beta} \delta \mathbf{d}_s \rangle_{\hat{\Sigma}} = -\text{R}(\mathbf{d}_s^{n+1,k}) (\delta \mathbf{d}_s), \end{aligned} \quad (7.49)$$

for all $\boldsymbol{\varphi} \in [X(\hat{\Omega})]^d$, $\boldsymbol{\varphi}|_{\hat{\Gamma}_D} = \mathbf{0}$, in analogy to (6.19) and (6.24). Inhomogeneous boundary conditions are enforced on the initial guess, $\mathbf{d}_s^{n+1,0}$, and the iterate is updated following $\mathbf{d}_s^{n+1,k+1} = \mathbf{d}_s^{n+1,k} + \delta \mathbf{d}_s^k$ with homogeneous Dirichlet conditions.

Spatial discretisation is carried out in a standard manner, decomposing the structural domain in the Lagrangian reference frame into $N_{e,s}$ shape-regular, non-overlapping Lagrangian elements $\hat{\Omega}_e$,

$$\hat{\Omega}_s \approx \hat{\Omega}_{s,h} = \bigcup_{e=1}^{N_{e,s}} \hat{\Omega}_{e,s}.$$

Polynomial shape functions defined on simplicial or tensor-product reference elements are mapped via isoparametric mappings $\boldsymbol{\chi}_e(\boldsymbol{\xi})$, leading to standard Lagrangian finite elements and respective C^0 -continuous function spaces as

$$\begin{aligned} X_{h,s}(\hat{\Omega}_{s,h}) & := \left\{ q_h \in C^0(\hat{\Omega}_{s,h}) : q_h \circ \boldsymbol{\chi}_e(\boldsymbol{\xi})|_{\hat{\Omega}_{e,s}} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{P}^k(\Omega_{\text{ref}}) \quad \forall \hat{\Omega}_{e,s} \in \hat{\Omega}_{s,h} \right\}, \\ \text{or } X_{h,s}(\hat{\Omega}_{s,h}) & := \left\{ q_h \in C^0(\hat{\Omega}_{s,h}) : q_h \circ \boldsymbol{\chi}_e(\boldsymbol{\xi})|_{\hat{\Omega}_{e,s}} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{Q}^k(\Omega_{\text{ref}}) \quad \forall \hat{\Omega}_{e,s} \in \hat{\Omega}_{s,h} \right\}. \end{aligned}$$

Plugging in the finite element approximations $\boldsymbol{\varphi} \approx \boldsymbol{\varphi}_h$ and $\delta \mathbf{d} \approx \delta \mathbf{d}_h$, both being elements of $[X_{h,s}(\hat{\Omega}_s)]^d \subset [H^1(\hat{\Omega}_s)]^d$, Eqn. (7.49) directly yields the fully discrete problem to be solved in each iteration of Newton's method, reading

$$\begin{aligned} & \langle \rho_s \boldsymbol{\varphi}_h, \frac{\alpha'_m}{\beta (\Delta t^n)^2} \delta \mathbf{d}_{s,h} \rangle_{\hat{\Omega}_s} + \langle \hat{\nabla} \boldsymbol{\varphi}_h, \alpha'_f \text{D}_d \mathbf{P}(\mathbf{d}_{s,h}^{n+1,k}) (\delta \mathbf{d}_{s,h}) \rangle_{\hat{\Omega}_s} \\ & + \langle \alpha'_f \boldsymbol{\varphi}_h, k_e \delta \mathbf{d}_{s,h} + \frac{c_e \gamma}{\Delta t^n \beta} \delta \mathbf{d}_{s,h} \rangle_{\hat{\Gamma}_{\text{ext}}} + \langle \alpha'_f \boldsymbol{\varphi}_h, \frac{\eta_s^R \gamma}{\Delta t^n \beta} \delta \mathbf{d}_{s,h} \rangle_{\hat{\Sigma}} = -\text{R}(\mathbf{d}_{s,h}^{n+1,k}) (\delta \mathbf{d}_{s,h}), \end{aligned} \quad (7.50)$$

where we refrain from indicating the domain and boundary approximations by h for ease of notation, understanding that the integrals are taken over the finite element approximations of the domain and its boundaries.

In the fluid domain, a similar approach is chosen, decomposing the current configuration at time t into $N_{e,f}$ non-overlapping finite elements,

$$\Omega_f^t \approx \Omega_{f,h}^t = \bigcup_{e=1}^{N_{e,f}} \Omega_{e,f}^t,$$

which naturally gives rise to the function spaces

$$\begin{aligned} X_{h,f}^t(\Omega_{f,h}^t) &:= \left\{ q_h \in C^0(\overline{\Omega_{f,h}^t}) : q_h \circ \chi_e^t(\boldsymbol{\xi})|_{\Omega_{e,f}^t} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{P}^k(\Omega_{\text{ref}}) \quad \forall \Omega_{e,f}^t \in \Omega_{f,h}^t \right\}, \\ \text{or } X_{h,f}^t(\Omega_{f,h}^t) &:= \left\{ q_h \in C^0(\overline{\Omega_{f,h}^t}) : q_h \circ \chi_e^t(\boldsymbol{\xi})|_{\Omega_{e,f}^t} = q_{\text{ref}}(\boldsymbol{\xi})|_{\Omega_{\text{ref}}} \in \mathbb{Q}^k(\Omega_{\text{ref}}) \quad \forall \Omega_{e,f}^t \in \Omega_{f,h}^t \right\}, \end{aligned} \quad (7.51)$$

where $\chi_e^t(\boldsymbol{\xi})$ again denotes the isoparametric mapping from the reference element Ω_{ref} to a specific element $\Omega_{e,f}^t \subset \Omega_{f,h}^t$.

We adopt the ALE version of the split-step scheme including Robin interface conditions as introduced in Sec. 4.2. That is, we employ a BDF m scheme with coefficients listed in Tab. 3.1 and combine it with divergence damping, adding the gradient term from Leray projection to the past time steps' velocities only,

$$\partial_t \mathbf{u}_f|_{\bar{\mathbf{x}}} \approx \alpha_0^m \mathbf{u}_f^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m \left(\mathbf{u}_f^{n-j} - \nabla \psi^{n-j} \right)|_{\bar{\mathbf{x}}}.$$

The convective velocity \mathbf{c}_f summarises the fluid and mesh velocities as before, but the latter is conveniently expressed in terms of the ALE time derivative of the fluid domain displacement \mathbf{d}_f , i.e., $\mathbf{u}_m := \partial_t \mathbf{x}|_{\bar{\mathbf{x}}} \equiv \partial_t \mathbf{d}_f|_{\bar{\mathbf{x}}}$, yielding

$$\mathbf{c}_f^{n+1} = \mathbf{u}_f^{n+1} - \mathbf{u}_m^{n+1} \approx \mathbf{u}_f^{n+1} - \sum_{j=0}^m \alpha_j^m \mathbf{d}_f^{n+1-j}|_{\bar{\mathbf{x}}}. \quad (7.52)$$

Now, an essential step is to introduce extrapolations of all involved fields as

$$\begin{aligned} \mathbf{u}_f^{n+1} \approx \mathbf{u}_f^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{u}_f^{n-j}|_{\bar{\mathbf{x}}}, & p_f^{n+1} \approx p_f^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m p_f^{n-j}|_{\bar{\mathbf{x}}}, \\ \mu_f^{n+1} \approx \mu_f^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mu_f^{n-j}|_{\bar{\mathbf{x}}}, & \mathbf{d}_f^{n+1} \approx \mathbf{d}_f^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{d}_f^{n-j}|_{\bar{\mathbf{x}}}, \\ \mathbf{c}_f^{n+1} \approx \mathbf{c}_f^* &= \mathbf{u}_f^* - \alpha_0^m \mathbf{d}_f^* - \sum_{j=1}^{m-1} \alpha_{j+1}^m \mathbf{d}_f^{n-j}|_{\bar{\mathbf{x}}}, & \mathbf{d}_s^{n+1} \approx \mathbf{d}_s^* &= \sum_{j=0}^{m-1} \beta_{j+1}^m \mathbf{d}_s^{n-j} \end{aligned} \quad (7.53)$$

with coefficients given in Tab. 2.1. Suitably inserting these extrapolations leads to a number of benefits, namely, it (i) decouples the linear momentum balance and pressure Poisson equations and thus enables added-mass stable semi-implicit fluid–structure coupling schemes, (ii) linearises the convective term and (iii) accelerates convergence by constructing proper initial guesses.

Within this work, all numerical tests regarding FSI problems are considering the stress-divergence form of the fluid momentum balance equation, coupling all the individual components of \mathbf{u}_f . The reason being that on the fluid–structure interface, the real

traction corresponding to the Cauchy stress is required and stabilisations such as SUPG or GLS discussed in Sec. 4.3 are based on the strong form's time-discrete residual being

$$\begin{aligned} R_f(\mathbf{u}_f^{n+1}, p_f^{n+1}) &:= \rho_f \alpha_0^m \mathbf{u}_f^{n+1} + \rho_f \sum_{j=0}^{m-1} \alpha_{j+1}^m (\mathbf{u}_f^{n-j} - \nabla \psi^{n-j}) + \nabla p_f^{n+1} \\ &+ \rho_f (\nabla \mathbf{u}_f^{n+1}) \mathbf{c}_f^{n+1} - \mu_f^{n+1} \Delta \mathbf{u}_f^{n+1} - \chi \mu_f^{n+1} \nabla (\nabla \cdot \mathbf{u}_f) - 2 \nabla^s \mathbf{u}_f^{n+1} \nabla \mu_f^{n+1}, \end{aligned} \quad (7.54)$$

and the corresponding time-discrete operator $L_f(\varphi)$ acting on \mathbf{u}_f defined as

$$L_f(\varphi) := \rho_f \alpha_0^m \varphi + \rho_f (\nabla \varphi) \mathbf{c}_f^{n+1} - \mu_f^{n+1} \Delta \varphi - \chi \mu_f^{n+1} \nabla (\nabla \cdot \varphi) - 2 \nabla^s \varphi \nabla \mu_f^{n+1}, \quad (7.55)$$

in the case of GLS or, if we were to employ SUPG stabilisation,

$$L_f(\varphi) := \rho_f (\nabla \varphi) \mathbf{c}_f^{n+1}. \quad (7.56)$$

Since $R_f(\mathbf{u}_f^{n+1}, p_f^{n+1})$ contains $\nabla^s \mathbf{u}_f$ even for the generalised Laplacian form, components inevitably couple. The general rationale we follow is keeping the stabilisation as implicit as possible, while remaining linear in \mathbf{u}_f . One might alternatively use extrapolations for any component-coupling terms, e.g., $\nabla^s \mathbf{u}_f^{n+1} \approx \nabla^s \mathbf{u}_f^*$, or choose to stabilise the decoupled equations in a component-wise manner. Together with boundary integrals correcting the discrepancy between pseudo- and real tractions on the interface, it is thus indeed possible to consider for the individual components. Within the numerical tests presented in this work, however, only the stress-divergence form is considered, as gains in efficiency from the laid out strategy are expected rather low. The reason for this is the semi-implicit FSI scheme, where the velocity field is merely solved once per time step to begin with, and the algebraic multigrid methods were found very effective given the small time steps adapted to satisfy the CFL condition. Given that the various extrapolations and correction terms might still influence the scheme under certain conditions, the more conservative option of considering the stress-divergence form of the fluid's momentum balance equation is chosen. For matrix-based methods, storage savings considering only d instead of $d \times d$ individual velocity component blocks might not be achieved in any way, if a (often necessary) coupled and/or nonlinear mesh motion equation as introduced in Sec. 7.5 is considered.

Another motivation for choosing the (generalised) Laplacian form is that outflow profiles are better preserved when enforcing pseudo tractions (see, e.g., [22–25] and others regarding outflow conditions). For generalised Newtonian fluids, however, pseudo tractions do not enforce mean pressures at planar outlets exactly, but approximately [26]. In addition to that, the viscous contribution to the traction vector might be small compared to the isotropic contribution from the pressure prescribed—especially in haemodynamic applications (see Sec. 4.6)—and outlet faces might not be straight to begin with. Lastly, in case of re-entrant flow, i.e., when $\mathbf{u}_f \cdot \mathbf{n}_f < 0$, destabilising effects must be countered with backflow stabilisation as discussed in Sec. 4.5, also heavily affecting the outflow region.

Motivated by these arguments, correcting boundary integrals added on the outlets might be the preferred option as commented on in Sec. 4.2. However, correcting for pseudo

tractions is rather unlikely to improve the results in terms of predicting the real flow fields close to the outlets, if the outlet *as a whole* is held fixed, which impacts the solution drastically. Mending this last aspect requires introducing (visco-)elastic bedding similar to the Robin condition for external tissue support at the outlet faces as introduced by Moireau et al. [359], but is not considered herein.

With this, we directly formulate the fully discrete weak form of the momentum balance equation in the discretised fluid domain Ω_f^{n+1} (again omitting the h for concision) seeking for $\mathbf{u}_{f,h}^{n+1} \in [X_{h,f}^{n+1}(\Omega_f^{n+1})]^d$ with $\mathbf{u}_{f,h}^{n+1}|_{\Gamma_{in,f}^{n+1}} = \mathbf{g}_f^{n+1}$, such that

$$\begin{aligned}
0 = & \langle \rho_f \boldsymbol{\varphi}_h, \alpha_0^m \mathbf{u}_{f,h}^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m (\mathbf{u}_{f,h}^{n-j} - \nabla \psi_h^{n-j}) \rangle + \langle \nabla \mathbf{u}_{f,h}^{n+1} \mathbf{c}_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} - \langle \boldsymbol{\varphi}_h, \mathbf{b}_f^{n+1} \rangle_{\Omega_f^{n+1}} \\
& + (1 - \chi/2) \langle \nabla \boldsymbol{\varphi}_h + \chi \nabla^\top \boldsymbol{\varphi}_h, \mu_{f,h}^{n+1} \nabla \mathbf{u}_{f,h}^{n+1} + \chi \mu_{f,h}^{n+1} \nabla^\top \mathbf{u}_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} - (1 - \chi) \langle \boldsymbol{\varphi}_h, (\nabla^\top \mathbf{u}_{f,h}^{n+1}) \nabla \mu_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} \\
& - \langle \nabla \cdot \boldsymbol{\varphi}_h, p_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} + \langle \boldsymbol{\varphi}_h, \eta_f^R \mathbf{u}_{f,h}^{n+1} - \mathbf{h}_{f,h}^{n+1} + (1 - \chi) \mu_{f,h}^{n+1} (\nabla^\top \mathbf{u}_{f,h}^{n+1}) \mathbf{n}_f \rangle_{\Sigma^{n+1}} \\
& - \langle \boldsymbol{\varphi}_h, \bar{\mathbf{t}}_f^{n+1} - \alpha_b \frac{\rho_f}{2} \mathbf{u}_{f,h}^{n+1} (\mathbf{c}_{f,h}^{n+1} \cdot \mathbf{n}_f)_- \rangle_{\Gamma_{out,f}^{n+1}} + \sum_{e=1}^{N_{e,f}} \langle \tau_e \mathbf{L}_f(\boldsymbol{\varphi}_h), \mathbf{R}_f(\mathbf{u}_{f,h}^{n+1}, p_{f,h}^{n+1}) \rangle_{\Omega_{e,f}^{n+1}}, \quad (7.57)
\end{aligned}$$

for all $\boldsymbol{\varphi}_h \in [X_{h,f}^{n+1}(\Omega_f^{n+1})]^d$, where $\boldsymbol{\varphi}_h|_{\Gamma_{in,f}^{n+1}} = \mathbf{0}$ and the Robin data on the fluid–structure interface is given by $\mathbf{h}_f := \eta_f^R \mathbf{d}_t \mathbf{d}_s + J_s^{-1} \mathbf{P}_s \mathbf{F}_s^\top \mathbf{n}_f$. Eqn. (7.57) includes divergence damping via $\nabla \psi_h^{n-j}$ (see Sec. 3.3), backflow stabilisation with a parameter $\alpha_b \geq 1$ as introduced in Sec. 4.5, residual-based stabilisation with parameter τ_e as defined in Eqn. (4.62) for convection-dominant problems (see Sec. 4.3) and the interface integrals incorporating Robin transmission conditions and corrections completing the real tractions for the generalised Laplacian case (Sec. 4.2), all of which are covered in Ch. 4. Depending on the data available, either quantities from t^{n+1} or suitable extrapolations at t^{n+1} are used to recover a linearised problem in $\mathbf{u}_{f,h}^{n+1}$ only. Moreover, Windkessel models are easily included by simply prescribing $\bar{\mathbf{t}}_f^{n+1} = -p_{c,i}^{n+1} \mathbf{n}_f$ on $\Gamma_{i,f}^{n+1}$, i.e., prescribing appropriate (pseudo-)traction data on the outlets (see Sec. 4.6).

Note here that the fluid’s pressure and the variable ψ used for divergence suppression are separated unlike in the fluid-only scenario presented in Ch. 3 since the physical pressure is needed for the interface condition. Thus, the pressure $p_{f,h}^{n+1} \in X_{h,f}^{n+1}(\Omega_f^{n+1})$ is retrieved from the weak form of the PPE incorporating the structure’s acceleration on the fluid–structure interface, such that $p_{f,h}^{n+1}|_{\Gamma_{out,f}^{n+1}} = \zeta_h^{n+1}$ and

$$\begin{aligned}
\langle \nabla \boldsymbol{\varphi}_h, \nabla p_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} = & \langle \nabla \boldsymbol{\varphi}_h, 2 (\nabla^\top \mathbf{u}_{f,h}^{n+1}) \nabla \mu_{f,h}^{n+1} - \rho_f (\nabla \mathbf{u}_{f,h}^{n+1}) \mathbf{c}_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} \\
& - \rho_f \sum_{j=0}^m \langle \boldsymbol{\varphi}_h \mathbf{n}_f, \alpha_j^m \mathbf{g}_f^{n+1-j} \rangle_{\Gamma_{in,f}^{n+1}} - \langle \boldsymbol{\varphi}_h \mathbf{n}_f, \rho_f \ddot{\mathbf{d}}_{s,h}^{n+1} \rangle_{\Sigma^{n+1}} \\
& + \langle \mathbf{n}_f \times \nabla \boldsymbol{\varphi}_h, \mu_{f,h}^{n+1} \nabla \times \mathbf{u}_{f,h}^{n+1} \rangle_{\Sigma^{n+1} \cup \Gamma_{in,f}^{n+1}} \quad (7.58)
\end{aligned}$$

holds for all $\boldsymbol{\varphi}_h \in X_{h,f}^{n+1}(\Omega_f^{n+1})$, where $\boldsymbol{\varphi}_h|_{\Gamma_{out,f}^{n+1}} = 0$ and $\ddot{\mathbf{d}}_{s,h}^{n+1}$ is evaluated using $\mathbf{d}_{s,h}^{n+1}$. On the inlet face, $\Gamma_{in,f}^{n+1}$, a BDF is employed for the time derivative of the inlet data $\partial_t \mathbf{g}_f|_{\bar{\mathbf{x}}}$, while on the interface, Σ^{n+1} , a Newmark approximation of the structure’s ma-

terial acceleration $\text{d}_{tt}\mathbf{d}_s$ is used. That is, the fluid–structure interface is treated as a Dirichlet boundary for the PPE, while in the fluid’s momentum equation, Robin transmission conditions are enforced on Σ^{n+1} . As already mentioned, this is similar to the strategy proposed by Astorino et al. [225]. Adopting a Robin transmission condition on the fluid–structure interface for the PPE as well lead to divergence of the scheme within a few steps in our numerical tests and is thus not further investigated in this work.

Continuous Dirichlet boundary data to enforce on the pressure at the outlets $\Gamma_{\text{out},f}^{n+1}$ is constructed as shown in Sec. 4.5. For this purpose, we define $Y_{h,f}^{n+1}(\Gamma_{\text{out},f}^{n+1}) \subset H^{1/2}(\Gamma_{\text{out},f}^{n+1})$ composed by traces of functions belonging to $X_{h,f}^{n+1}(\Omega_f^{n+1})$ on $\Gamma_{\text{out},f}^{n+1}$. The fully discrete problem is then to find continuous pressure boundary data $\zeta_h^{n+1} \in Y_{h,f}^{n+1}(\Gamma_{\text{out},f}^{n+1})$, such that for all $s_h \in Y_{h,f}^{n+1}(\Gamma_{\text{out},f}^{n+1})$, there holds

$$\begin{aligned} \langle s_h, \zeta_h^{n+1} \rangle_{\Gamma_{\text{out},f}^{n+1}} &= -\langle s_h, \mu_{f,h}^{n+1} \nabla \cdot \mathbf{u}_{f,h}^{n+1} \rangle_{\Gamma_{\text{out},f}^{n+1}} \\ &+ \langle s_h \mathbf{n}_f, \mu_{f,h}^{n+1} (\nabla \mathbf{u}_{f,h}^{n+1} + \chi \nabla^\top \mathbf{u}_{f,h}^{n+1}) \mathbf{n}_f - \bar{\mathbf{t}}_f^{n+1} + \alpha_b \frac{\rho_f}{2} \mathbf{u}_{f,h}^{n+1} (\mathbf{c}_{f,h}^{n+1} \cdot \mathbf{n}_f)_- \rangle_{\Gamma_{\text{out},f}^{n+1}}. \end{aligned} \quad (7.59)$$

Here, backflow stabilisation contributes once again and Windkessel models can be included setting Neumann data $\bar{\mathbf{t}}_f^{n+1} = -p_{c,i}^{n+1} \mathbf{n}_f$ on the individual outlet faces $\Gamma_{i,f}^{n+1}$. The pressure Dirichlet data projection variable ζ_h^{n+1} is not defined on Σ^{n+1} , given that for the PPE, the fluid–structure interface is considered a pure Dirichlet boundary. Similar to the pure flow problem, we solve for the PPE’s boundary data in one single linear system, gathering the nodal DoFs to approximate disconnected sets of ζ_h^{n+1} on each outlet face.

The viscosity is projected onto a continuous space to recover $\mu_f^{k+1} \in X_{h,f}^{k+1}(\Omega_f^{k+1})$ via

$$\langle \varphi_h, \mu_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} = \langle \varphi_h, \eta (\hat{\gamma}(\nabla \mathbf{u}_{f,h}^{n+1})) \rangle_{\Omega_f^{n+1}}, \quad (7.60)$$

for all $\varphi_h \in X_{h,f}^{n+1}(\Omega_f^{n+1})$. Again, step (7.60) is only needed when considering generalised Newtonian fluids and modifying the rheological law is achieved with little effort.

The Leray projection step to modify the fluid velocity $\mathbf{u}_{f,h}^{n+1}$ for divergence suppression in the following step consists of finding $\psi_h^{n+1} \in X_{h,f}^{n+1}(\Omega_f^{n+1})$ using \mathbf{u}_f^{n+1} , such that $\psi_h^{n+1}|_{\Gamma_{\text{out},f}^{n+1}} = 0$, $\psi_h^{n+1}|_{\Sigma^\tau} = 0$ and

$$\langle \nabla \varphi_h, \nabla \psi_h^{n+1} \rangle_{\Omega_f^{n+1}} = -\langle \varphi_h, \nabla \cdot \mathbf{u}_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}}, \quad (7.61)$$

holds for all $\varphi_h \in X_{h,f}^{n+1}(\Omega_f^{n+1})$ with $\varphi_h|_{\Gamma_{\text{out},f}^{n+1}} = 0$ and $\varphi_h|_{\Sigma^\tau} = 0$. The variable ψ_h^{n+1} is used on the past time step’s velocities only as described in Sec. 3.3. This step is optional, but consists merely of a single Poisson problem in the fluid domain per time step. Due to the improvement in mass conservation at a relatively low cost, it is employed if not indicated otherwise.

The last missing component in the FSI scheme is the fully discrete weak form of the ALE mesh motion equation, where we introduced harmonic extension (7.15)–(7.17) for the sake of presentation. Hence, find the fluid domain's displacement $\mathbf{d}_{f,h}^{n+1} \in [X_{h,f}(\hat{\Omega}_f)]^d$, such that $\mathbf{d}_{f,h}^{n+1}|_{\hat{\Sigma}} = \mathbf{d}_{s,h}^{n+1}$, $\mathbf{d}_{f,h}|_{\partial\hat{\Omega}_f \setminus \hat{\Sigma}} = \mathbf{0}$ and there holds

$$\langle \hat{\nabla} \boldsymbol{\varphi}_h, \hat{\nabla} \mathbf{d}_{f,h}^{n+1} \rangle_{\hat{\Omega}_f} = 0, \quad (7.62)$$

for all $\boldsymbol{\varphi}_h \in [X_{h,f}(\hat{\Omega}_f)]^d$ with $\boldsymbol{\varphi}|_{\partial\hat{\Omega}_f} = \mathbf{0}$. More practical alternatives to Eqn. (7.62) are given in Sec. 7.5.

The discrete problems governing the Newton update to the structure's displacement $\mathbf{d}_{s,h}^{n+1}$ (7.50) and the fluid's velocity $\mathbf{d}_{f,h}^{n+1}$, pressure $p_{f,h}^{n+1}$, viscosity $\mu_{f,h}^{n+1}$ and auxiliary variables ζ_h^{n+1} and ψ_h^{n+1} in Eqns. (7.57)–(7.62) are repeatedly executed at each time step. Interface data is updated with the most recent iterates until the absolute and relative convergence criteria are fulfilled. Denoting, e.g., the discrete finite element vector of nodal unknowns containing the current iterate of the structural displacement at time t^{n+1} by $\underline{\mathbf{d}}_s^{n+1,k+1}$ and the last iterate by $\underline{\mathbf{d}}_s^{n+1,k}$, we define

$$\left\| \underline{\mathbf{d}}_s^{n+1,k+1} - \underline{\mathbf{d}}_s^{n+1,k} \right\| \leq \epsilon_{\text{abs}}^d \quad \text{and} \quad \left\| \underline{\mathbf{d}}_s^{n+1,k+1} - \underline{\mathbf{d}}_s^{n+1,k} \right\| \leq \epsilon_{\text{rel}}^d \left\| \underline{\mathbf{d}}_s^n \right\|, \quad (7.63)$$

as absolute and relative convergence criteria. Analogously, define finite element vectors of nodal unknowns of the fluid velocity and pressure to check convergence in the fluid velocity via

$$\left\| \underline{\mathbf{u}}_f^{n+1,k+1} - \underline{\mathbf{u}}_f^{n+1,k} \right\| \leq \epsilon_{\text{abs}}^u \quad \text{and} \quad \left\| \underline{\mathbf{u}}_f^{n+1,k+1} - \underline{\mathbf{u}}_f^{n+1,k} \right\| \leq \epsilon_{\text{rel}}^u \left\| \underline{\mathbf{u}}_f^n \right\|, \quad (7.64)$$

and in the fluid pressure via

$$\left\| \underline{\mathbf{p}}_f^{n+1,k+1} - \underline{\mathbf{p}}_f^{n+1,k} \right\| \leq \epsilon_{\text{abs}}^p \quad \text{and} \quad \left\| \underline{\mathbf{p}}_f^{n+1,k+1} - \underline{\mathbf{p}}_f^{n+1,k} \right\| \leq \epsilon_{\text{rel}}^p \left\| \underline{\mathbf{p}}_f^n \right\|. \quad (7.65)$$

A single time step from t^n to t^{n+1} of the fully implicit partitioned scheme using the PPE-based split-step scheme for the fluid subproblem can be written as

1. *Initial guess:* Extrapolate in time via (7.53) based on the past time steps' solutions \mathbf{d}_s^{n-j} , \mathbf{d}_f^{n-j} , \mathbf{u}_f^{n-j} , p_f^{n-j} and μ_f^{n-j} , where $j = 0, \dots, m-1$, and set those as initial guesses for all involved fields.
2. *Implicit coupling loop:*
WHILE not converged according to (7.63), (7.64) and (7.65) **DO**
 - a) *Mesh displacement:* Compute $\mathbf{d}_f^{n+1,k+1}$ extending the structural displacement $\mathbf{d}_s^{n+1,k}$ from $\hat{\Sigma}$ into $\hat{\Omega}_f$ solving (7.62).
 - b) *Mesh velocity:* Update the fluid domain Ω_f^{n+1} , the mesh velocity \mathbf{u}_m^{n+1} and the convective velocity \mathbf{c}_f^* via (7.52) using $\mathbf{u}_f^{n+1,k}$ and $\mathbf{d}_f^{n+1,k+1}$.
 - c) *Fluid viscosity:* Project the fluid's dynamic viscosity $\mu_f^{n+1,k+1}$ onto a continuous space via Eqn. (7.60) using $\mathbf{u}_f^{n+1,k}$.

- d) *Pressure boundary data*: Construct continuous boundary data $\zeta^{n+1,k+1}$ for the PPE via (7.59) given $\mu_f^{n+1,k+1}$, $\mathbf{u}_f^{n+1,k}$ and \mathbf{c}_f^* .
- e) *Fluid pressure*: Determine the pressure $p_f^{n+1,k+1}$ solving the PPE (7.58) using $\mu_f^{n+1,k+1}$, $\mathbf{u}_f^{n+1,k}$, \mathbf{c}_f^* and $\ddot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s^{n+1,k})$.
- f) *Structural deformation*: Execute a nonlinear solver, e.g., Newton's scheme repeatedly solving (7.50) using $\mathbf{u}_f^{n+1,k}$, $\mu_f^{n+1,k+1}$ and $p_f^{n+1,k+1}$ until converged according to (5.70).
- g) *Acceleration*: modify the updated $\mathbf{d}_s^{n+1,k+1}$ (optional, see Sec. 7.4).
- h) *Fluid velocity*: Compute the fluid's velocity $\mathbf{u}_f^{n+1,k+1}$ via Eqn. (7.57) using $\mu_f^{n+1,k+1}$, $p_f^{n+1,k+1}$, $\mathbf{d}_s^{n+1,k+1}$, \mathbf{c}_f^* for linearisation and backflow stabilisation and $\ddot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s^{n+1,k+1})$.

END WHILE

3. *Update converged iterates*: Set $\mathbf{d}_s^{n+1} = \mathbf{d}_s^{n+1,k+1}$, $\mathbf{d}_f^{n+1} = \mathbf{d}_f^{n+1,k+1}$, $\mathbf{u}_f^{n+1} = \mathbf{u}_f^{n+1,k+1}$, $p_f^{n+1} = p_f^{n+1,k+1}$ and $\mu_f^{n+1} = \mu_f^{n+1,k+1}$.
4. *Divergence suppression*: Update the auxiliary variable ψ^{n+1} given \mathbf{u}_f^{n+1} to be used in the next time step as ψ^n on \mathbf{u}_f^n .
5. *Advance in time*: Proceed to next time step.

The above scheme treats all involved fields including the convective velocity \mathbf{c} implicitly and couples each of the variables strongly, thus serves as a basis for the semi-implicit schemes presented in the following Sec. 7.3. The standard nonlinear term \mathbf{c}^{n+1} is reintroduced and linearised via simple Picard iteration. This linearisation strategy for the fluid momentum balance equation is equivalent to a single fixed-point iteration and suffices to resolve the nonlinearities over several FSI steps given small enough time steps. As in any multiphysics problem, acceleration schemes modifying the iterate(s) are applicable. In the context of FSI, often the interface displacement is relaxed, accelerating the coupling scheme's convergence as discussed in Sec. 7.4.

7.3 Semi-implicit variants of the coupling scheme

To design a semi-implicit variant of the implicit Robin–Robin (RR) FSI scheme in the spirit of Fernández et al. [441], a split-step or projection scheme such as the present PPE-based one has to be employed as a flow solver. This allows recovering the fluid's velocity and pressure independent of one another, such that updating the interface conditions in the fluid subproblem leads to an update in the boundary data enforced on the substep yielding the fluid pressure. For the fully discrete PPE (7.58), the updated solid material acceleration enters the boundary term $\langle \varphi \mathbf{n}_f, \rho_f \ddot{\mathbf{d}}_s^{n+1} \rangle_{\Sigma^{n+1}}$. Therefore, fluid pressure and structural displacement can be coupled directly, while treating the fluid velocity contribution explicit.

In the haemodynamic regime, the added-mass effect is large, since $\rho_f/\rho_s \approx 1$, time steps might be small and considered geometries are slender. Strongly coupling the fluid pressure and structural displacements, however, is a valid strategy to ensure a stable FSI scheme as numerically observed and proven for simplified model problems (see, e.g., [351, 441, 444, 476]). Combined with suitable extrapolation, temporal accuracy is preserved, while the explicit treatment reduces computational cost significantly. Hence, a first semi-implicit variant of the Robin–Robin scheme (SIRR) is constructed moving the ALE update, the viscosity projection and the fluid momentum step out of the coupling loop, similar to [192, 225, 441, 443–445, 447, 466]. A single time step then reads

-
1. *Initial guess*: Extrapolate via (7.53) based on \mathbf{d}_s^{n-j} , \mathbf{d}_f^{n-j} , \mathbf{u}_f^{n-j} , p_f^{n-j} and μ_f^{n-j} , where $j = 0, \dots, m-1$, and set those as initial guesses.
 2. *Mesh displacement*: Compute \mathbf{d}_f^{n+1} extending the structural displacement \mathbf{d}_s^* from $\hat{\Sigma}$ into $\hat{\Omega}_f$ solving the ALE mesh motion equation (7.62).
 3. *Mesh velocity*: Update the fluid domain Ω_f^{n+1} , the mesh velocity \mathbf{u}_{m}^{n+1} and the convective velocity \mathbf{c}_f^* via (7.52) using \mathbf{u}_f^* and \mathbf{d}_f^{n+1} .
 4. *Fluid viscosity*: Project the fluid’s dynamic viscosity μ_f^{n+1} onto a continuous space via Eqn. (7.60) using the extrapolated \mathbf{u}_f^* .
 5. *Pressure boundary data*: Construct continuous boundary data ζ^{n+1} for the PPE via (7.59) given μ_f^{n+1} , \mathbf{u}_f^* and \mathbf{c}_f^* .
 6. *Implicit coupling loop*:
WHILE not converged according to (7.63) *or* (7.65) **DO**
 - a) *Fluid pressure*: Determine $p_f^{n+1,k+1}$ solving (7.58) using μ_f^{n+1} , \mathbf{u}_f^* , \mathbf{c}_f^* , boundary data ζ^{n+1} and $\ddot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s^{n+1,k})$.
 - b) *Structural deformation*: Execute a nonlinear solver, e.g., Newton’s scheme repeatedly solving (7.50) using \mathbf{u}_f^* , μ_f^{n+1} and $p_f^{n+1,k+1}$ until converged according to (5.70).
 - c) *Acceleration*: modify the updated $\mathbf{d}_s^{n+1,k+1}$ (optional, see Sec. 7.4).**END WHILE**
 7. *Update converged iterates*: Set $\mathbf{d}_s^{n+1} = \mathbf{d}_s^{n+1,k+1}$ and $p_f^{n+1} = p_f^{n+1,k+1}$.
 8. *Fluid velocity*: Compute the fluid’s velocity \mathbf{u}_f^{n+1} via Eqn. (7.57) using the implicitly coupled iterates, \mathbf{c}_f^* for linearisation, backflow stabilisation and $\ddot{\mathbf{d}}_s^{n+1}(\mathbf{d}_s^{n+1})$.
 9. *Divergence suppression*: Update the auxiliary variable ψ^{n+1} given \mathbf{u}_f^{n+1} to be used in the next time step as ψ^n on \mathbf{u}_f^n .
 10. *Advance in time*: Proceed to next time step.
-

Compared to the pure flow solver discussed in Ch. 3, where pressure and velocity steps might be exchanged in their sequence, numerical experiments with the FSI scheme show

that the fluid momentum step considering the implicitly coupled structural displacement and fluid pressure is more robust than computing the fluid velocity before executing the coupling loop. This way, vast changes from the extrapolated, non-coupled pressure to the implicitly coupled pressure do not enter the fluid's momentum balance, which results in improved stability of the overall method—especially when resistance-based outflow models are considered.

In comparison to related methods, we can formulate the following distinct features of the proposed SIRR scheme based on a PPE for generalised Newtonian fluids:

1. Higher-order and adaptive timestepping schemes are directly available via standard timestepping schemes and suitable extrapolation formulae.
2. Interpolation with standard equal-order, C^0 -continuous Lagrangian finite elements is enabled, circumventing inf-sup conditions of velocity-pressure formulations without any stabilisations.
3. The semi-implicit design implicitly coupling only the fluid pressure and structural displacements in itself reduces computing times tremendously, while preserving stability properties and accuracy.
4. Standard acceleration strategies (see Sec. 7.4) acting on the structure's interface displacement are straight-forwardly applied in addition to the Robin–Robin coupling algorithm.
5. Rheological laws of generalised Newtonian fluids are effortlessly exchanged, since only the right-hand side in the viscosity projection step needs to be adapted.
6. Due to the partitioned design, adaptations in the structure's constitutive equation or adding models relevant for practical application (e.g., Windkessel outlets or viscoelastic support of the tissue) are carried out easily.
7. The PPE-based split-step scheme itself is iteration-free and features only a low number of steps.
8. Divergence damping hinders errors in mass balance from accumulating, but does neither spoil interface conditions nor necessitate an additional vector-valued projection step to recover a weakly divergence-free velocity field as required by standard projection schemes.
9. All involved problems yield well-known linear systems of equations and are hence effectively tackled by standard linear solvers and preconditioners. Naturally, off-the-shelf black-box preconditioning techniques available as open-source scientific software can be employed, which is the strategy followed herein.

The Robin–Robin coupling algorithm with parameters η_s^R and η_f^R on the structure and fluid side, respectively, includes the classical Dirichlet–Neumann (DN) strategy. In the limit case $\eta_f^R \rightarrow \infty$ and $\eta_s^R = 0$, the fluid subproblem involves $\mathbf{u}_f = \mathbf{d}_t \mathbf{d}_s$ on Σ^t , while the structure subproblem enforces $\mathbf{P}_s \hat{\mathbf{n}}_s = J_f \boldsymbol{\sigma}_f \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s$ on $\hat{\Sigma}$. In a fully discrete setting, however, η_f^R too large may lead to unnecessary ill-conditioning of the linear system of the fluid's momentum balance equation. On the other side, assigning a bounded value may not enforce the interface conditions strong enough at all times.

This is a problem easily mended by a slight adaption of the fully discrete fluid momentum balance equation, incorporating the Dirichlet boundary condition on the fluid–structure interface directly in to the function spaces. In analogy to (7.57), in this case we seek for $\mathbf{u}_{f,h}^{n+1} \in [X_{h,f}^{n+1}(\Omega_f^{n+1})]^d$, where now $\mathbf{u}_{f,h}^{n+1}|_{\Gamma_{\text{in},f}^{n+1}} = \mathbf{g}_f^{n+1}$ and $\mathbf{u}_{f,h}^{n+1}|_{\Sigma^{n+1}} = \mathbf{d}_{s,h}^{n+1}$ are embedded in the function space, such that

$$\begin{aligned}
0 = & \langle \rho_f \boldsymbol{\varphi}_h, \alpha_0^m \mathbf{u}_{f,h}^{n+1} + \sum_{j=0}^{m-1} \alpha_{j+1}^m (\mathbf{u}_{f,h}^{n-j} - \nabla \psi_h^{n-j}) \rangle + \langle \nabla \mathbf{u}_{f,h}^{n+1} \mathbf{c}_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} - \langle \boldsymbol{\varphi}_h, \mathbf{b}_f^{n+1} \rangle_{\Omega_f^{n+1}} \\
& + (1 - \chi/2) \langle \nabla \boldsymbol{\varphi}_h + \chi \nabla^\top \boldsymbol{\varphi}_h, \mu_{f,h}^{n+1} \nabla \mathbf{u}_{f,h}^{n+1} + \chi \mu_{f,h}^{n+1} \nabla^\top \mathbf{u}_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} - (1 - \chi) \langle \boldsymbol{\varphi}_h, (\nabla^\top \mathbf{u}_{f,h}^{n+1}) \nabla \mu_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} \\
& - \langle \nabla \cdot \boldsymbol{\varphi}, p_{f,h}^{n+1} \rangle_{\Omega_f^{n+1}} - \langle \boldsymbol{\varphi}_h, \bar{\mathbf{t}}_f^{n+1} - \alpha_b \frac{\rho_f}{2} \mathbf{u}_{f,h}^{n+1} (\mathbf{c}_{f,h}^{n+1} \cdot \mathbf{n}_f)_- \rangle_{\Gamma_{\text{out},f}^{n+1}} \\
& + \sum_{e=1}^{N_{e,f}} \langle \tau_e \mathbf{L}_f(\boldsymbol{\varphi}_h), \mathbf{R}_f(\mathbf{u}_{f,h}^{n+1}, p_{f,h}^{n+1}) \rangle_{\Omega_{e,f}^{n+1}}, \tag{7.66}
\end{aligned}$$

holds in the discretised fluid domain Ω_f^{n+1} (again omitting the h for concision) for all $\boldsymbol{\varphi}_h \in [X_{h,f}^{n+1}(\Omega_f^{n+1})]^d$, with $\boldsymbol{\varphi}_h|_{\Gamma_{\text{in},f}^{n+1}} = \mathbf{0}$ and $\boldsymbol{\varphi}|_{\Sigma^{n+1}} = \mathbf{0}$. Hence, the interface integral present in the Robin variant (7.57) vanishes. The remaining parts in this discrete problem remain completely unchanged, that is, Eqn. (7.66) includes divergence damping via $\nabla \psi_h^{n-j}$ (see Sec. 3.3), backflow stabilisation with a parameter $\alpha_b \geq 1$ (Sec. 4.5), residual-based stabilisation with parameter τ_e as defined in Eqn. (4.62) for convection-dominant problems (Sec. 4.3). No major changes apart from the function spaces are required, exchanging the fluid momentum step (7.57) by (7.66), yielding the Dirichlet–Robin variant of the split-step FSI scheme.

7.4 Accelerating FSI schemes

Acceleration/relaxation schemes adapting the exchanged interface data can significantly speed up convergence of the FSI coupling procedure within a (semi-)implicit partitioned approach. They remain largely independent of the specific choice of interface condition and can therefore be used in conjunction with Robin interface conditions. Already starting from a constant relaxation parameter as proposed by Le Tallec and Mouro [189], improvements are observed compared to the standard Dirichlet–Neumann scheme. In cases of high added-mass effect, Aitken’s acceleration [87–89], which dynamically adapts the relaxation parameter during the iteration process, further allows tackling challenging applications in the haemodynamic regime as shown by [88, 477].

Another important development emanated from interpreting the FSI coupling loop as a fixed point problem in the interface unknowns, such that interface Newton schemes and their (locally) quadratic convergence can be exploited. Going in this direction, interface quasi-Newton and Newton–Krylov solvers have been shown to both increase convergence speed of the iterative scheme, while at the same time rendering the coupling method more robust [430, 431, 433, 434]. A particular landmark contribution in this context is the Interface Quasi-Newton Inverse Least-Squares (IQN-ILS) method introduced

by Degroote et al. [435], approximating the inverse Jacobian directly based on input and output pairs of past coupling iterations. Reusing information from past iterates turns out to be extremely beneficial, but with this, two problems arise [425, 461, 478, 479]: first, the number of data pairs considered increases the computational effort spent in the acceleration scheme and data from past time steps might not be adequately representing the coupling iteration's inverse Jacobian of the *current* time step. Second, (almost) linearly dependent data vectors have the potential to completely annihilate any gains in efficiency, since the minimisation problem is not well conditioned and hence leads to unfavourable coefficients in the Newton approximation.

In this regard, improvements can be achieved by altering the selection of data pairs and/or filtering out linearly dependent data vectors from the (inverse) Jacobian approximation [432, 480, 481]. Doing so, data pairs from past time steps, which are *not* linearly dependent, but *contradict* the current Jacobian approximation still remain to be eliminated. In this regard, implicit reuse of the past time step's data lead to the so-called Interface Quasi-Newton Inverse Multi-Vector Jacobian method [461] and the Interface Quasi-Newton Implicit Multi-Vector Least-Squares method [436], which allow eliminating the problem dependent parameter regulating the number of time steps used for the Jacobian approximation or renders it less influential.

Within this work, we restrict ourselves to the standard IQN-ILS method [435] with filtering, closely following Degroote [425]. The comparisons presented in [425, 436, 482] indicate that the basic IQN-ILS with filtering can in fact yield a performance very close to more advanced schemes, but is implementationally much simpler. All tackled problems within this work were found rather insensitive to the number of time steps considered for the construction of the inverse Jacobian, especially when combining the IQN-ILS with filtering and Robin interface conditions. This combination of Robin interface conditions and interface quasi-Newton methods was only very recently considered by Spenke et al. [462] to the best of the author's knowledge. The novelty of this present work lies in the combination of IQN-ILS with semi-implicit schemes and Robin interface conditions, distinctively different from [462], which considers a fully implicit partitioned scheme involving a coupled velocity-pressure formulation of the Navier–Stokes equations.

Following the presentation in [425], we write the FSI coupling loop—be it the reduced pressure–structure coupling loop or the fully implicit coupling loop—as a nonlinear equation in the discrete structural displacement restricted to the interface,

$$\underline{\mathbf{d}}_{\Sigma}^k := \underline{\mathbf{d}}_s^k|_{\Sigma}, \quad \underline{\mathbf{d}}_{\Sigma}^k \in \mathbb{R}^{N_{\Sigma} \times 1},$$

with N_{Σ} being the number of displacement degrees of freedom on the interface. Since continuous function spaces are being used for the structural equations, the map from nodal vectors containing all degrees of freedom in the structural domain $\underline{\mathbf{d}}_s$ to the interface degrees of freedom $\underline{\mathbf{d}}_{\Sigma}$ is trivial.

To ease the definition of the interface equation and its residual, define the structural displacement and fluid interface operators as \mathbf{S} and \mathbf{F} , respectively. The latter operator \mathbf{F} summarises all parts of the fluid flow problem included in the FSI coupling loop.

Hence, it potentially includes the domain update, viscosity projection, PPE and fluid momentum balance steps, but might reduce to the PPE step only, if the semi-implicit variants of the FSI scheme are chosen. A single iteration of the complete FSI problem can then be written as

$$\tilde{\mathbf{d}}_{\Sigma}^k := \mathbf{S} \circ \mathbf{F} \left(\mathbf{d}_{\Sigma}^k \right).$$

The two most recent iterates of the interface displacement, the unmodified new iterate $\tilde{\mathbf{d}}_{\Sigma}^k$ and the last iterate \mathbf{d}_{Σ}^k , are used to formulate the interface residual \mathbf{r}_{Σ}^k in the structural displacement as the difference between them, i.e.,

$$\mathbf{r}_{\Sigma}^k := \mathbf{S} \circ \mathbf{F} \left(\mathbf{d}_{\Sigma}^k \right) - \mathbf{d}_{\Sigma}^k = \tilde{\mathbf{d}}_{\Sigma}^k - \mathbf{d}_{\Sigma}^k. \quad (7.67)$$

With this introduced, a relaxed update step reads

$$\mathbf{d}_{\Sigma}^{k+1} = \omega \tilde{\mathbf{d}}_{\Sigma}^k + (1 - \omega) \mathbf{d}_{\Sigma}^k, \quad (7.68)$$

following [189], where a fixed relaxation parameter ω in the range $0 < \omega < 1$ is selected to speed up convergence and/or prevent divergence of the coupling scheme. Naturally, this relaxation parameter depends on the problem at hand and an optimal choice for ω is likely to vary over time for a single problem considered. To improve upon this strategy, Mok and Wall [477] and later Küttler and Wall [88] proposed using Aitken's method to dynamically determine the relaxation parameter ω during the iterative process. The relaxation parameter ω_k is *recursively* defined dependent on the current and last iterate and the corresponding residuals as

$$\omega_k := -\omega_{k-1} \left\| \Delta \mathbf{r}_{\Sigma}^k \right\|^{-2} \mathbf{r}_{\Sigma}^{k-1} \cdot \Delta \mathbf{r}_{\Sigma}^k, \quad (7.69)$$

with the difference in the residual vectors $\Delta \mathbf{r}_{\Sigma}^k$ given by

$$\Delta \mathbf{r}_{\Sigma}^k := \mathbf{r}_{\Sigma}^k - \mathbf{r}_{\Sigma}^{k-1}. \quad (7.70)$$

Due to its effectiveness and simplicity, this method serves as a great starting point and rightfully so earned great popularity. For easier comparison, the Aitken-accelerated coupling scheme is given in Alg. 1.

In contrast to that, a Newton method can be applied to the interface residual (7.67), yielding a single Newton step of the form

$$\mathbf{d}_{\Sigma}^{k+1} = \tilde{\mathbf{d}}_{\Sigma}^k - \left(\frac{\mathbf{D}\mathbf{r}}{\mathbf{D}\mathbf{d}_{\Sigma}} \bigg|_{\mathbf{d}_{\Sigma}^k} \right)^{-1} \mathbf{r}_{\Sigma}^k, \quad (7.71)$$

directly written in terms of the inverted Jacobian of the interface residual evaluated with the last iterate. As laid out in [425], the Newton scheme (7.71) does neither require explicit computation of the interface Jacobian nor its inversion, but rather the application of its inverse to a vector.

Algorithm 1 Aitken-accelerated coupling scheme

- 1: Set $k = 0$ and initial relaxation factor ω_0
 - 2: Extrapolate all fields in time and set initial guess $\underline{\mathbf{d}}_{\Sigma}^k = \underline{\mathbf{d}}_s^*|_{\Sigma}$ via (7.53)
 - 3: **while** Not converged acc. to Eqns. (7.63)–(7.65) **do**
 - 4: Perform single coupling step: $\tilde{\underline{\mathbf{d}}}_{\Sigma}^k = \mathbf{S} \circ \mathbf{F}(\underline{\mathbf{d}}_{\Sigma}^k)$
 - 5: Update residual: $\mathbf{r}_{\Sigma}^k = \tilde{\underline{\mathbf{d}}}_{\Sigma}^k - \underline{\mathbf{d}}_{\Sigma}^k$
 - 6: **if** $k > 0$ **then**
 - 7: Update difference in residuals (7.70): $\Delta \mathbf{r}_{\Sigma}^k = \mathbf{r}_{\Sigma}^k - \mathbf{r}_{\Sigma}^{k-1}$
 - 8: Update relaxation factor (7.69): $\omega_k = -\omega_{k-1} \left\| \Delta \mathbf{r}_{\Sigma}^k \right\|^{-2} \mathbf{r}_{\Sigma}^{k-1} \cdot \Delta \mathbf{r}_{\Sigma}^k$
 - 9: **end if**
 - 10: Relaxation step (7.68): $\underline{\mathbf{d}}_{\Sigma}^{k+1} = \omega_k \tilde{\underline{\mathbf{d}}}_{\Sigma}^k + (1 - \omega_k) \underline{\mathbf{d}}_{\Sigma}^k$
 - 11: Update $\underline{\mathbf{d}}_s^{n+1, k+1}|_{\Sigma} = \underline{\mathbf{d}}_{\Sigma}^{k+1}$, $\underline{\mathbf{d}}_s^{n+1}$ and $\underline{\mathbf{d}}_s^{n+1}$ via (5.44)–(5.45)
 - 12: $k = k + 1$
 - 13: **end while**
-

The key aspect of the work by Degroote et al. [435] is to approximate this matrix-vector product using differences in past coupling step's unmodified iterates and residuals, which are collected as columns of the matrices $\mathbf{W}_k \in \mathbb{R}^{N_{\Sigma} \times k}$

$$\mathbf{W}_k := \left(\Delta \tilde{\underline{\mathbf{d}}}_{\Sigma}^k, \Delta \tilde{\underline{\mathbf{d}}}_{\Sigma}^{k-1}, \dots, \Delta \tilde{\underline{\mathbf{d}}}_{\Sigma}^1 \right), \quad \text{with} \quad \Delta \tilde{\underline{\mathbf{d}}}_{\Sigma}^k := \tilde{\underline{\mathbf{d}}}_{\Sigma}^k - \tilde{\underline{\mathbf{d}}}_{\Sigma}^{k-1} \quad (7.72)$$

and $\mathbf{V}_k \in \mathbb{R}^{N_{\Sigma} \times k}$

$$\mathbf{V}_k := \left(\Delta \mathbf{r}_{\Sigma}^k, \Delta \mathbf{r}_{\Sigma}^{k-1}, \dots, \Delta \mathbf{r}_{\Sigma}^1 \right), \quad (7.73)$$

which grow as k increases. The aim is of course reaching convergence, i.e., $\mathbf{r}_{\Sigma}^{k+1} \stackrel{!}{=} \underline{\mathbf{0}}$, where we have $\Delta \mathbf{r}_{\Sigma}^{k+1} \stackrel{!}{=} \underline{\mathbf{0}} - \mathbf{r}_{\Sigma}^k$. Based on the known collected differences in residuals, we approximate

$$\Delta \mathbf{r}_{\Sigma}^{k+1} \approx \mathbf{V}_k \underline{\mathbf{c}}^k, \quad (7.74)$$

linearly combining them via coefficients $\underline{\mathbf{c}}^k \in \mathbb{R}^{k \times 1}$, such that we can reformulate the Newton update (7.71) as

$$\underline{\mathbf{d}}_{\Sigma}^{k+1} = \tilde{\underline{\mathbf{d}}}_{\Sigma}^k + \mathbf{W}_k \underline{\mathbf{c}}^k. \quad (7.75)$$

Note here, however, that Eqn. (7.74) is an overdetermined system as \mathbf{V}_k is not square, such that an exact solution is unlikely to exist and only a minimizer can be found by means of solving a least-squares problem of the form

$$\min_{\underline{\mathbf{c}}^k} \left\| \mathbf{V}_k \underline{\mathbf{c}}^k + \mathbf{r}_{\Sigma}^k \right\|. \quad (7.76)$$

Using the normal equations, an explicit formulation of the approximated inverse Jacobian can be found as [478, 481, 483]

$$\left(\begin{array}{c} \text{D}\underline{\mathbf{r}}_{\Sigma} \\ \text{D}\underline{\mathbf{d}}_{\Sigma} \end{array} \bigg|_{\underline{\mathbf{d}}_{\Sigma}^k} \right)^{-1} \approx \mathbf{W}_k \left(\mathbf{V}_k^{\top} \mathbf{V}_k \right)^{-1} \mathbf{V}_k^{\top}.$$

As the size of \mathbf{V}^k and/or linearly dependencies increase, the preferred strategy, however, is to consider a QR-decomposition of \mathbf{V}_k ,

$$\mathbf{QR} = \mathbf{V}_k, \quad (7.77)$$

with an orthogonal matrix $\mathbf{Q} \in \mathbb{R}^{N_{\Sigma} \times k}$ and upper triangular $\mathbf{R} \in \mathbb{R}^{k \times k}$. To compute the QR-decomposition, Householder transformations [425, 484], (modified) Gram–Schmidt orthogonalisation [432] or Givens rotations [484] are viable alternatives. Within this work, the last option is considered due to the simple reason of a matrix-free implementation—meaning, the potentially big \mathbf{Q} is not stored explicitly—being available in `deal.II` [96–99].

An important aspect during the construction of the QR-decomposition is the detection and removal of linearly dependent vectors on the fly [425, 432]. The overall strategy closely follows [425], rebuilding the decomposition at every time step, starting with the most recent data vectors and adding older vectors only, if the to-be-added column of the upper triangular matrix \mathbf{R} yields a diagonal entry $[\mathbf{R}]_{kk} \geq \epsilon_{\text{QR}}$. This allows effective reuse of data from the q past time steps and reduces the memory footprint by removing old and/or linearly dependent data from \mathbf{V}_k and \mathbf{W}_k . Rebuilding the QR-decomposition at every time step performs less operations on the more recent data. Thus, these newer contributions are less affected by numerical round-off whilst forming the decomposition. It is *crucial* to note here also that the difference vectors do *not* span multiple time steps and that linearly independent, but old data pairs can impair the scheme’s performance depending on q being the number of time steps data is considered from.

With the QR-decomposition based on the filtered \mathbf{V}_k considering data pairs from the last q time steps, we can rewrite (7.74) as

$$\begin{aligned} \Delta_{\underline{\mathbf{r}}_{\Sigma}}^{k+1} &\stackrel{!}{=} -\underline{\mathbf{r}}_{\Sigma}^k \approx \mathbf{V}_k \underline{\mathbf{c}}^k = \mathbf{QR} \underline{\mathbf{c}}^k \Rightarrow \mathbf{Q}^{\top} \mathbf{QR} \underline{\mathbf{c}}^k = -\mathbf{Q}^{\top} \underline{\mathbf{r}}_{\Sigma}^k \\ &\Rightarrow \mathbf{R} \underline{\mathbf{c}}^k = -\mathbf{Q}^{\top} \underline{\mathbf{r}}_{\Sigma}^k \end{aligned} \quad (7.78)$$

due to the orthogonality of \mathbf{Q} , that is $\mathbf{Q}^{-1} = \mathbf{Q}^{\top}$. To compute the action of \mathbf{Q}^{\top} onto $\underline{\mathbf{r}}^k$, we make use of

$$\mathbf{V}_k = \mathbf{QR} \Rightarrow \mathbf{V}_k^{\top} = \mathbf{R}^{\top} \mathbf{Q}^{\top} \Rightarrow \mathbf{R}^{-\top} \mathbf{V}_k^{\top} = \mathbf{Q}^{\top}, \quad (7.79)$$

avoiding explicit use of \mathbf{Q} . Meaning, we first compute a vector $\underline{\mathbf{a}} := \mathbf{V}_k^{\top} \underline{\mathbf{r}}_{\Sigma}^k$ and perform an additional back-insertion pass solving $\mathbf{R}^{\top} \underline{\mathbf{b}} = -\underline{\mathbf{a}}$. This yields $\underline{\mathbf{b}} := -\mathbf{Q}^{-\top} \underline{\mathbf{r}}_{\Sigma}^k$ and the final step is then to solve $\mathbf{R} \underline{\mathbf{c}}^k = \underline{\mathbf{b}}$ to recover the coefficient vector $\underline{\mathbf{c}}^k$. With this, the full IQN-ILS-accelerated coupling scheme can then be summarised as shown in Alg. 2.

Algorithm 2 IQN-ILS-accelerated coupling scheme

-
- 1: Set $k = 0$ and initial relaxation factor ω_0
 - 2: Extrapolate all fields in time and set initial guess $\underline{\mathbf{d}}_\Sigma^k = \underline{\mathbf{d}}_s^*|_\Sigma$ via (7.53)
 - 3: **while** Not converged acc. to Eqns.(7.63)–(7.65) **do**
 - 4: Perform single coupling step: $\tilde{\underline{\mathbf{d}}}_\Sigma^k = \mathbf{S} \circ \mathbf{F}(\underline{\mathbf{d}}_\Sigma^k)$
 - 5: Remove columns from \mathbf{V}_k and \mathbf{W}_k corresponding to time step $n - q - 1$
 - 6: Update residual: $\underline{\mathbf{r}}_\Sigma^k = \tilde{\underline{\mathbf{d}}}_\Sigma^k - \underline{\mathbf{d}}_\Sigma^k$
 - 7: Update difference vectors: $\Delta \tilde{\underline{\mathbf{d}}}_\Sigma^k = \tilde{\underline{\mathbf{d}}}_\Sigma^k - \tilde{\underline{\mathbf{d}}}_\Sigma^{k-1}$ and $\Delta \underline{\mathbf{r}}_\Sigma^k = \underline{\mathbf{r}}_\Sigma^k - \underline{\mathbf{r}}_\Sigma^{k-1}$
 - 8: Insert new data pairs into \mathbf{V}_k and \mathbf{W}_k as defined in (7.72) and (7.73)
 - 9: **if** $k = 0$ and $(q = 0$ or $n = 0)$ **then**
 - 10: Initial relaxation step: $\underline{\mathbf{d}}_\Sigma^{k+1} = \omega_0 \tilde{\underline{\mathbf{d}}}_\Sigma^k + (1 - \omega_0) \underline{\mathbf{d}}_\Sigma^k$
 - 11: **else**
 - 12: Perform QR-decomposition (7.77), adding column k if $[\mathbf{R}]_{kk} \geq \epsilon_{\text{QR}}$
 - 13: Remove columns from \mathbf{V}_k and \mathbf{W}_k that were rejected
 - 14: Compute $-\mathbf{Q}^\top \underline{\mathbf{r}}_\Sigma^k$ in (7.78) by solving $\mathbf{R}^\top \underline{\mathbf{d}} = -\underline{\mathbf{a}}$ with $\underline{\mathbf{a}} := \mathbf{V}^\top \underline{\mathbf{r}}_\Sigma^k$
 - 15: Solve $\mathbf{R} \underline{\mathbf{c}}^k = \underline{\mathbf{b}}^k$ (7.78)
 - 16: IQN step (7.75): $\underline{\mathbf{d}}_\Sigma^{k+1} = \tilde{\underline{\mathbf{d}}}_\Sigma^k + \mathbf{W}_k \underline{\mathbf{c}}^k$
 - 17: **end if**
 - 18: Update $\underline{\mathbf{d}}_s^{n+1, k+1}|_\Sigma = \underline{\mathbf{d}}_\Sigma^{k+1}$, $\underline{\mathbf{d}}_s^{n+1}$ and $\underline{\mathbf{d}}_s^{n+1}$ via (5.44)–(5.45)
 - 19: $k = k + 1$
 - 20: **end while**
-

Let us close this section with a few remarks on the algorithm and its use within this work: first, the simple filtering approach with parameter ϵ_{QR} rejecting columns which otherwise lead to small diagonal values in \mathbf{R} is found rather robust. In our numerical tests, we exclusively choose $\epsilon_{\text{QR}} = 10^{-16}$, which seems surprisingly low, but was found robust in combination with considering data from $q = \mathcal{O}(10)$ time steps.

Also, the additional back-insertion pass to avoid explicit use of \mathbf{Q}^\top does not influence the overall computing time, as the time spent in the whole IQN-ILS acceleration is insignificant compared to the application of a single coupling step. IQN-ILS tremendously speeds-up the iterative coupling process leading to significantly reduced overall simulation times and also renders the coupling algorithm more robust. That being said, we also want to emphasise that the comparably low computational effort connected to IQN-ILS acceleration does heavily depend on the ratio of the interface degrees of freedom to the bulk degrees of freedom in the fluid and solid subproblems.

Along the lines of Spenke et al. [462], we observe that combining Robin interface conditions and interface quasi-Newton methods eases the problem-dependent choice of q , which leads to the parameter choice being rather uncritical in all our numerical tests. If the share of relative computing time spent with IQN-ILS were bigger, the Interface Quasi-Newton Implicit Multi-Vector Least-Squares method [436] might be an attractive alternative. This method approximates the inverse Jacobian without critical tuning parameters and scales linearly with the problem size, but is on the other side slightly more complex in its implementation and hence not considered herein.

7.5 Mesh motion in ALE methods

Updating the moving computational domain is a centerpiece of any interface tracking approach within an FSI framework, as the extension of the interface displacement into the fluid domain can quickly become a delicate matter. Mesh motion techniques employed in this regard have to preserve mesh quality and bijectivity, since loss of invertibility of the ALE extension renders any results meaningless and often leads to divergence prior to the mapping's breakdown. Therefore, numerous methods have been devised to update the computational grid, ranging from interpolation methods employing radial basis functions [485–487] or based on transfinite mappings [187, 488, 489], spring analogies [490–495] or solving artificial boundary value problems with appropriate boundary data. The latter approach is from the perspective of code reuse the most attractive choice, and has hence seen vast progress in the recent past. Among the most popular choices are harmonic extension [496, 497], employing the Laplace operator, parabolic extension [467] in analogy to the heat equation, pseudo linear elasticity problems [474, 475, 498, 499], nonlinear (hyper-)elasticity [468–472] and biharmonic extension [473, 500], solving a fourth-order operator. Since the operator and potential pseudo material parameters considered are arbitrary, obvious tuning possibilities arise. Frequently used strategies include Jacobian-based stiffening [471, 474, 475, 501], stiffening small or distorted elements by dropping or taking powers of the (finite element) map's Jacobian to the computational domain, or considering an optimized zero-stress state [468, 470, 472] to counteract element distortion via more involved hyperelastic constitutive laws, pseudo growth and remodelling. Comparative studies have been carried out by [25, 467, 471, 473, 502–504] among others, but do not include the full spectrum of available methods and can thus only give an incomplete picture of the schemes currently available.

Within this work, we focus on the third mentioned family of approaches, namely, solving an artificial boundary value problem. Despite the semi-implicit setup only requiring an extension of the interface displacement once per time step, and hyperelastic, fibre-reinforced tissue models being available as potential candidates for a highly robust mesh motion technique, the most simple material models and local stiffening techniques are adopted. This is due to (i) reduced computational complexity in terms of element integration and linear system solve, and (ii) the introduced (comparably simple) extension techniques are already able to capture the fluid domain motion while maintaining a reasonable mesh quality for all considered numerical examples.

For the sake of presentation, Eqns. (7.15)–(7.17) introduce merely the basic harmonic extension involving a simple Laplace operator,

$$\begin{aligned}
 -\hat{\nabla} \cdot \hat{\nabla} \mathbf{d}_f &= \mathbf{0} && \text{in } \hat{\Omega}_f \\
 \mathbf{d}_f &= \mathbf{d}_s && \text{on } \hat{\Sigma} \\
 \mathbf{d}_f &= \mathbf{0} && \text{on } \Gamma_{\text{in},f}^t \cup \Gamma_{\text{out},f}^t
 \end{aligned}$$

where pure Dirichlet conditions are enforced on the entirety of the fluid reference domain's boundary in a first simplified formulation.

The corresponding weak problem is thus to find the fluid domain's displacement $\mathbf{d}_{f,h}^{n+1} \in [X_{h,f}(\hat{\Omega}_f)]^d$, such that $\mathbf{d}_{f,h}^{n+1}|_{\hat{\Sigma}} = \mathbf{d}_{s,h}^{n+1}$, $\mathbf{d}_{f,h}|_{\partial\hat{\Omega}_f \setminus \hat{\Sigma}} = \mathbf{0}$ and

$$\langle \hat{\nabla} \boldsymbol{\varphi}_h, \hat{\nabla} \mathbf{d}_{f,h}^{n+1} \rangle_{\hat{\Omega}_f} = 0,$$

for all $\boldsymbol{\varphi}_h \in [X_{h,f}(\hat{\Omega}_f)]^d$ with $\boldsymbol{\varphi}|_{\partial\hat{\Omega}_f} = \mathbf{0}$ and $X_{h,f}(\hat{\Omega}_f)$ again adopting discrete spaces of continuous nodal functions on the reference fluid domain $\hat{\Omega}_f$, cf. Eqns.(7.51) and (7.62). Now, tremendous improvements can be achieved introducing a variable pseudo diffusivity, following [85, 417, 467, 473, 475]. Assuming smaller elements present at the boundary, in corners or on the fluid–structure interface, where boundary layers are expected and potentially larger displacements are seen, we can shield smaller elements from excessive distortion increasing their local diffusivity depending on their volume $h_e := |\Omega_e|$, setting

$$c_h := c_{\min} + (c_{\max} - c_{\min}) \frac{c_0 - c_1}{c_1 (c_0 - 1)}, \quad (7.80)$$

$$\text{with } c_1 := 1 + (c_0 - 1) \sin\left(\frac{\pi}{2} \frac{h_e - h_{\min}}{h_{\max} - h_{\min}}\right), \quad (7.81)$$

to enforce upper and lower bounds on the element-local stiffness parameters, giving c_{\min} for h_{\max} and c_{\max} for h_{\min} with an additional scaling parameter c_0 to modify the transition between minimal and maximal values as shown in Fig. 7.2. Similarly, elements already undergoing large distortion are assigned increased diffusivity/stiffness based on the ALE mapping's Jacobian $J_f := \det \mathbf{F}_f$, defining

$$c_J := J_f^{-1}, \quad c_J := J_f + J_f^{-1} \quad \text{or} \quad c_J := J_f^{-2}, \quad (7.82)$$

following, or very much similar to [85, 417, 473] as demonstrated in Fig. 7.2 for varying Jacobians. As can be seen, the scaling through c_h (7.80) allows precise tuning via minimal/maximal target values with an additional transition parameter c_0 . The Jacobian stiffening based entirely on the ALE map leads to an increased parameter, as $J_f \rightarrow 0$ as desired. The diffusion or stiffening parameter is then given as

$$c_{m,e}(\mathbf{d}_f, h_e) := c_J(\mathbf{d}_f) c_h(h_e), \quad (7.83)$$

which directly leads to a corresponding strong form

$$-\hat{\nabla} \cdot (c_m \hat{\nabla} \mathbf{d}_f) = \mathbf{0} \quad \text{in } \hat{\Omega}_f, \quad (7.84)$$

$$\mathbf{d}_f = \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad (7.85)$$

$$\mathbf{d}_f \cdot \mathbf{n}_f = \mathbf{0} \quad \text{on } \Gamma_{\text{in},f}^t \cup \Gamma_{\text{out},f}^t, \quad (7.86)$$

with a continuous, but nonlinear diffusion coefficient $c_m(\mathbf{d}_f, h_e)$ and including a zero normal displacement boundary condition on the inlets and outlets (7.86), i.e., $\partial\hat{\Omega}_f \setminus \hat{\Sigma}$. Such a boundary condition allows the mesh to slide along the domain's boundaries, which may ease construction of a valid fluid domain and/or allows considering for deformable inlets and outlets. On the interface, the no-slip condition leads to $\mathbf{u}_f = \mathbf{d}_t \mathbf{d}_s$ and might

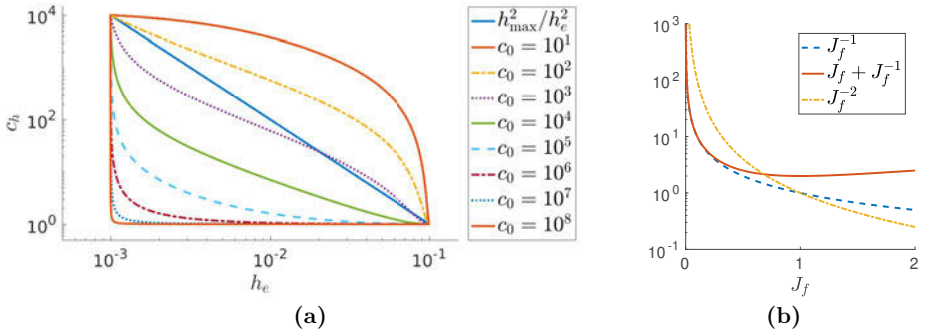


Fig. 7.2: Scaling parameter c_h (a) for $h_{\max}/h_{\min} = 100$ compared to scaling h_{\max}^2/h_e^2 introduced by Jendoubi et al. [467] and Jacobian stiffening variants (b).

be further altered allowing for slip, enforcing $\mathbf{u}_f \cdot \mathbf{n}_f = d_t \mathbf{d}_s \cdot \mathbf{n}_f$, but in any case, the mesh motion might be considered independently, enforcing continuity of displacement in normal direction only, replacing (7.85). To circumvent the need for data mapping algorithms, however, we herein stick to $\mathbf{d}_f = \mathbf{d}_s$ on the interface and choose no-slip interface conditions $\mathbf{u}_f = d_t \mathbf{d}_s$, the latter being motivated by physics and the former facilitating implementation.

The weak form of the harmonic extension with Jacobian stiffening and weakly enforced zero normal displacement on inlets and outlets then leads to the problem of finding the fluid domain displacement $\mathbf{d}_{f,h}^{n+1} \in [X_{h,f}(\hat{\Omega}_f)]^d$, such that $\mathbf{d}_{f,h}^{n+1}|_{\hat{\Sigma}} = \mathbf{d}_{s,h}^{n+1}$ and

$$\langle c_{m,e}(\mathbf{d}_{f,h}^n, h_e) \hat{\nabla} \varphi_h, \hat{\nabla} \mathbf{d}_{f,h}^{n+1} \rangle_{\hat{\Omega}_f} + c_n \langle \varphi_h \cdot \hat{\mathbf{n}}_f, \mathbf{d}_{f,h}^{n+1} \cdot \hat{\mathbf{n}}_f \rangle_{\Gamma_{in,f}^t \cup \Gamma_{out,f}^t} = 0, \quad (7.87)$$

for all $\varphi_h \in [X_{h,f}(\hat{\Omega}_f)]^d$ with $\varphi|_{\hat{\Sigma}} = \mathbf{0}$. Here, $c_n > 0$ is a large enough penalty parameter enforcing $\mathbf{d}_f \cdot \hat{\mathbf{n}}_f = 0$, distinctly different from the rudimentary approach in Eqn. (7.62) strongly enforcing $\mathbf{d}_f = \mathbf{0}$ on inlets and outlets. Due to the diffusion coefficient $c_{m,e}$ depending on the displacement \mathbf{d}_f itself via the Jacobian J_f , Eqn. (7.87) would be nonlinear, if discretised fully implicitly. Assuming that the mesh motion is continuous in time, we *always* linearise $c_{m,e}$ using the last time step's domain displacement $\mathbf{d}_{f,h}^n$.

In a similar manner, a (non-)linear elasticity problem can be applied to extend the interface motion into the fluid domain, in strong form reading

$$-\hat{\nabla} \cdot \mathbf{P}_m(\mathbf{d}_f) = \mathbf{0} \quad \text{in } \hat{\Omega}_f, \quad (7.88)$$

$$\mathbf{d}_f = \mathbf{d}_s \quad \text{on } \hat{\Sigma}, \quad (7.89)$$

$$\mathbf{d}_f \cdot \mathbf{n}_f = \mathbf{0} \quad \text{on } \Gamma_{in,f}^t \cup \Gamma_{out,f}^t. \quad (7.90)$$

Here, \mathbf{P}_m denotes the pseudo stress tensor as defined in Eqn. (5.19) as

$$\mathbf{P}_m(\mathbf{d}) := \lambda_m \mathbf{I} \hat{\nabla} \cdot \mathbf{d} + \mu_m (\hat{\nabla} \mathbf{d} + \hat{\nabla}^\top \mathbf{d}), \quad (7.91)$$

with λ_m and μ_m being pseudo Lamé parameters. In analogy to Eqn. (5.16) for elastic homogeneous isotropic materials, we define

$$\lambda_m := \frac{E_m \nu_m}{(1 - 2\nu_m)(1 + \nu_m)} \quad , \quad \mu_m := \frac{E_m}{2(1 + \nu_m)}. \quad (7.92)$$

Introducing a variable stiffness dependent on the ALE mapping's Jacobian J_f and the element size h_e , we consider a pseudo Young's modulus $E_{m,e}(\mathbf{d}_f, h_e) := c_{m,e}(\mathbf{d}_f, h_e)$ per element given via Eqns. (7.80)–(7.83) and Poisson's ratio $\nu_{m,e}(h_e)$

$$\nu_{m,e} := \nu_0 - \nu_1 \frac{c_{\min}}{c_h(h_e)}, \quad (7.93)$$

such that $\nu_{m,e}(h_{\min}) \approx \nu_0$ (for $c_{\max} \gg c_{\min}$) and $\nu_{m,e}(h_{\max}) = \nu_0 - \nu_1$, adopted from [467]. This immediately leads to the mesh motion equation based on a pseudo elasticity problem: find $\mathbf{d}_{f,h}^{n+1} \in [X_{h,f}(\hat{\Omega}_f)]^d$, such that $\mathbf{d}_{f,h}^{n+1}|_{\hat{\Sigma}} = \mathbf{d}_{s,h}^{n+1}$ and

$$\langle \hat{\nabla} \boldsymbol{\varphi}_h, \mathbf{P}_m(\mathbf{d}_{f,h}^{n+1}) \rangle_{\hat{\Omega}_f} + c_n \langle \boldsymbol{\varphi}_h \cdot \hat{\mathbf{n}}_f, \mathbf{d}_{f,h}^{n+1} \cdot \hat{\mathbf{n}}_f \rangle_{\Gamma_{\text{in},f}^t \cup \Gamma_{\text{out},f}^t} = 0, \quad (7.94)$$

for all $\boldsymbol{\varphi}_h \in [X_{h,f}(\hat{\Omega}_f)]^d$ with $\boldsymbol{\varphi}|_{\hat{\Sigma}} = \mathbf{0}$, where $c_n > 0$ is again a large enough penalty parameter to weakly enforce zero normal displacements, i.e., $\mathbf{d}_f \cdot \hat{\mathbf{n}}_f = 0$. The pseudo material parameters being the element-local pseudo Young's modulus $E_{m,e}(\mathbf{d}_{f,h}^n, h_e)$ and Poisson's ratio $\nu_{m,e}(h_e)$ are again linearised.

Naturally, this idea may be extended towards (nearly-)incompressible hyperelastic continua [470–472] by choosing a suitable representation of \mathbf{P}_m . All of the constitutive relations discussed in Sec. 5.1 may be employed, where we focus on a nearly incompressible neo-Hookean solid for the sake of brevity and since we do not employ (necessarily oriented) fibers in the ALE extension, which then leads to

$$\mathbf{P}_m(\mathbf{d}) := \kappa_m/2(J^2 - 1)\mathbf{F}^{-\top} + \mu_m J^{-2/3} \left(\mathbf{F} - 1/3 I_1 \mathbf{F}^{-\top} \right), \quad (7.95)$$

compare Eqns. (5.31) and (5.34). The pseudo bulk modulus κ_m derives from an analogous relation to Eqn. (5.17) for homogeneous isotropic materials, yielding

$$\kappa_m = \frac{E_m}{3(1 - 2\nu_m)}, \quad (7.96)$$

depending on the pseudo Young's modulus E_m and Poisson's ratio ν_m , yielding an identical representation of the corresponding weak formulation as shown in Eqn. (7.94). Notice, however, that \mathbf{P}_m is nonlinear when considering pseudo hyperelasticity, such that a nonlinear solution method, e.g., a Newton scheme has to be employed, repeatedly solving the Jacobian system until convergence criteria similar to Eqn. (5.70),

$$\left\| \underline{\mathbf{d}}_f^{n+1,k+1} - \underline{\mathbf{d}}_f^{n+1,k} \right\| \leq \epsilon_{\text{abs}}^m \quad \text{and} \quad \left\| \underline{\mathbf{d}}_f^{n+1,k+1} - \underline{\mathbf{d}}_f^{n+1,k} \right\| \leq \epsilon_{\text{rel}}^m \left\| \underline{\mathbf{d}}_f^n \right\|, \quad (7.97)$$

are met. The interested reader is referred to Sec. 5.4 for remarks on the linearisation of the Piola–Kirchhoff stress tensor, $\text{D}_d \mathbf{P}_m$, which are directly applicable, since the pseudo material parameters are linearised in time using $\mathbf{d}_{f,h}^n$.

In the applied Newton scheme, we enforce the interface condition on the last iterate of the fluid domain displacement, i.e., set $\mathbf{d}_{f,h}^{n+1,k} = \mathbf{d}_{s,h}^{n+1}$ on $\hat{\Sigma}$, which gives the discrete weak form: find the Newton update $\delta \mathbf{d}_{f,h} \in [X_{h,s}(\hat{\Omega}_f)]^d$ with $\delta \mathbf{d}_{f,h}|_{\hat{\Sigma}} = \mathbf{0}$, such that for all $\varphi_h \in [X_{h,s}(\hat{\Omega}_f)]^d$ with $\varphi_h|_{\hat{\Sigma}} = \mathbf{0}$, there holds

$$\begin{aligned} & \langle \hat{\nabla} \varphi_h, \text{D}_d \mathbf{P}_m(\mathbf{d}_{f,h}^{n+1,k})(\delta \mathbf{d}_{f,h}) \rangle_{\hat{\Omega}_f} + c_n \langle \varphi_h \cdot \hat{\mathbf{n}}_f, \delta \mathbf{d}_{f,h} \cdot \hat{\mathbf{n}}_f \rangle_{\Gamma_{\text{in},f}^t \cup \Gamma_{\text{out},f}^t} \\ & = - \langle \hat{\nabla} \varphi_h, \mathbf{P}_m(\mathbf{d}_{f,h}^{n+1,k}) \rangle_{\hat{\Omega}_f} - c_n \langle \varphi_h \cdot \hat{\mathbf{n}}_f, \mathbf{d}_{f,h}^{n+1,k} \cdot \hat{\mathbf{n}}_f \rangle_{\Gamma_{\text{in},f}^t \cup \Gamma_{\text{out},f}^t}, \end{aligned} \quad (7.98)$$

yielding $\delta \mathbf{d}_{f,h}^k$ to update the Newton iterate in the standard way, that is

$$\mathbf{d}_{f,h}^{n+1,k+1} = \mathbf{d}_{f,h}^{n+1,k} + \delta \mathbf{d}_{f,h}^k.$$

Given the freedom in selecting a specific mesh motion equation and parameters, it is also clear that the nonlinear problem of pseudo hyperelasticity is not necessarily solved up until low tolerances ϵ_{abs}^m and ϵ_{rel}^m are reached, but rather once per timestep (as proposed by Shamanskiy and Simeon [471]), keeping the associated numerical effort low.

Formulating the mesh motion equation in the reference domain $\hat{\Omega}_f$ might lead to hard-to-solve problems in both the linear and nonlinear case due to Jacobian stiffening and general nonlinear terms stemming from the constitutive equation employed. Taking the most recent solution $\mathbf{d}_{f,h}^{n+1,k}$ or $\mathbf{d}_{f,h}^n$ as initial guess for the linear solvers and Newton’s method, these effects are mitigated. However, having the current fluid domain configuration Ω_f^t and not the reference configuration $\hat{\Omega}_f$ available without mapping the computational domain back and forth or storing the fluid grid twice might be advantageous and is thus an alternative worthwhile considering. Additionally, when formulating the mesh motion equation in the current configuration Ω_f^t , element sizes and hence the diffusion/stiffening coefficients do *also* account for geometric changes on top of the Jacobian-based stiffening. As pointed out in [471], following this approach might lead to distortions accumulating in the grid, depending on the specific problem at hand—as is the case when the (non-)linear solver tolerances are not adequately chosen.

Further options to design appropriate extension types incorporate time derivatives in the balance equations, leading to initial boundary value problems to be solved at each call of the mesh update. The initial boundary value problem and all added parameters being arbitrary again, a vast pool of possible combinations arise with prominent examples being parabolic extension introduced by Jendoubi et al. [467] or various types of pseudo time-dependent hyperelasticity, as presented, e.g., by Takizawa et al. [470]. Herein, we combine all the presented spatial operators with first and second material time derivatives, that is, we add $\text{d}_t \mathbf{d}_f$ or $\text{d}_{tt} \mathbf{d}_f$ to the strong form of the mesh motion equations.

Since we are not necessarily interested in time accurate methods, we employ a BDF1, i.e., backward Euler time discretisation for the first and Chung–Hulbert- α time integration for the second time derivative. In the latter case, we set $\rho_\infty = 0$ to achieve maximum numerical damping for high-frequency modes, which yields $\alpha_m = -1$ and $\alpha_f = 0$, such that $\beta = 1$ and $\gamma = 3/2$. Then, the instationary momentum balance equation for pseudo elasticity with zero body force reads in weak form

$$0 = \langle \varphi, \frac{2}{\Delta t^2} (\mathbf{d}^{n+1} - \mathbf{d}^n) - \frac{2}{\Delta t} \dot{\mathbf{d}}^n \rangle + \langle \hat{\nabla} \varphi, \mathbf{P}^{n+1} \rangle,$$

where $\ddot{\mathbf{d}}_f^n$ cancels, the stress term is fully implicit and Neumann boundary terms are either set to zero or do not arise in the first place. In our numerical tests presented in Sec. 8.4, we further tested setting $\dot{\mathbf{d}}_s^n = 0$ as initial condition in *each* time step of the overall timestepping scheme, which did not significantly alter obtained results. Doing so, mesh velocities and accelerations are not carried over time steps in the FSI algorithm and one ends up with an element integration routine very much similar to a BDF1/backward Euler time stepping scheme in the first two pseudo time steps, but with adapted time step size. To remain competitive with other extension types and to not reach the quasi-stationary state (equivalent to the stationary counterparts), only a single or very few pseudo time steps are to be executed. For consistency, the initial condition for \mathbf{d}_f^n is chosen as \mathbf{d}_f^n , when mapping from $\hat{\Omega}_f$, and $\mathbf{0}$ otherwise for any extension technique considered.

To compare the mesh quality and possible accumulation of distortions resulting from the various extension types, several quality criteria are available. These are based on, e.g., min./max. node distances yielding a measure of aspect ratio, angles between faces or edges to determine skewness or ratios of Jacobians within an element to judge the strength of the mapping from reference to real element [505–507]. To detect non-invertibility of the (incremental) ALE mapping, its determinant J_f alone is sufficient, but it only includes the map from the ALE map’s reference configuration to Ω_f^t . Since the computational grid in its reference configuration $\hat{\Omega}_f$ might already include elements with bad aspect ratio or high skewness, we judge element quality by

$$Q_J := \min_{e=1, \dots, N_{e,f}} Q_{J,e},$$

$$\text{with } Q_{J,e} := \frac{\min_{q=1, \dots, N_q} (J_{\text{ref}}|_{\chi_e^{-1}(\hat{\mathbf{x}}_q)} J_f|_{\hat{\mathbf{x}}_q})}{\max_{q=1, \dots, N_q} (J_{\text{ref}}|_{\chi_e^{-1}(\hat{\mathbf{x}}_q)} J_f|_{\hat{\mathbf{x}}_q})}, \quad (7.99)$$

where N_q is the number of Gauss integration points, $\hat{\mathbf{x}}_q$ denotes their position in $\hat{\Omega}_e$ and $\chi_e^{-1}(\hat{\mathbf{x}}_q)$ the corresponding position in the reference finite element Ω_{ref} . In an isoparametric setting, the finite element map from the reference element Ω_{ref} to the current element $\hat{\Omega}_e$ is given as

$$\chi_e(\boldsymbol{\xi}) := \sum_{i=1}^{N_{\text{ref}}} N_i(\boldsymbol{\xi}) \hat{\mathbf{x}}_i,$$

with N_{ref} nodal shape functions $N_i(\boldsymbol{\xi})$. Hence, the Jacobian of the finite element map is given in the standard way, $J_{\text{ref}} := \det(\nabla_{\boldsymbol{\xi}} \boldsymbol{\chi}_e)$, and is straight-forwardly included in the quality criterion. Considering an extension from Ω_f^n to Ω_f^{n+1} , a similar trail of thoughts can be followed leading to an identical implementation merely changing the domain of integration and corresponding displacement increment considered. In this case, the displacement increment is defined on Ω_f^n and J_f in Eqn. (7.99) is replaced by the determinant of the displacement increment, which is distinctly different from the total ALE mapping's Jacobian.

Since the ALE extension given the interface displacement can be investigated separately from any FSI or flow problem on the moving domain, a numerical example with prescribed solid motion adopted from Jendoubi et al. [467] is presented in Sec. 8.4, postponing any related discussions for the moment.

8 The coupled problem: computational results

The aim of this section is to thoroughly investigate the proposed scheme and its variants for their accuracy, robustness and influence of various parameters included in the formulation. Similar to the numerical tests carried out in preceding Secs. 2.7 and 3.5, the implementation is based on the finite element toolbox `deal.II` [96–99] and makes heavy use of Trilinos’ ML package [111] for preconditioning the linear system solves by algebraic multigrid methods.

First, the fundamental behaviour of the scheme is investigated, whereas more complex applications in biomechanics are tackled in Secs. 8.5–8.8. The problems considered first are taken from [455, 456] to demonstrate the convergence behaviour of the method, then, variations of a classical benchmark in FSI being pulsatile flow through a straight vessel (see, e.g., [256, 405, 463]) are studied. Afterwards, the proposed ALE extension types with various options for stiffening are tested in a two-dimensional setup with prescribed interface displacements taken from [467], before moving on to more challenging setups in haemodynamics. These employ physiological parameters in realistic scenarios of pulsatile flow through an abdominal aortic aneurysm (Sec. 8.5), an iliac bifurcation (Sec. 8.6) and patient-specific aortic dissection (Sec. 8.7). Lastly, Sec. 8.8 summarises investigations on the PPE-based coupling scheme’s performance in the context of human phonation, showcasing applicability to aeroelastic problems.

8.1 Analytical solution: rectangular piston

The analytical solution presented in [455, 456] can be interpreted as an elastic piston undergoing periodic motion in vertical direction, driving an incompressible Newtonian fluid, which exits and enters freely over the boundary of the computational domain. The domain is $\hat{\Omega} = \hat{\Omega}_f \cup \hat{\Omega}_s = [0, L] \times [-H, H]$, where $L = 1$ and $H = 0.5$, and the fluid occupies $\hat{y} \geq 0$ at $t = 0$, while the structure’s Lagrangian reference frame is in $\hat{y} \leq 0$ as shown in Fig. 8.1.

An exact solution to the FSI problem can be found assuming vertical motion only, i.e., setting horizontal velocities and displacements to zero, and defining the interface motion $d_{\Sigma,y}$ according to

$$d_{\Sigma,y} := a \sin(\omega t) \quad \text{with} \quad a := 2\alpha \sin \frac{\omega H}{c_p}, \quad c_p := \sqrt{\frac{\lambda_s + 2\mu_s}{\rho_s}}, \quad (8.1)$$




Fig. 8.1: Rectangular piston problem: finite element mesh at refinement level 4.

leading to an oscillatory motion in vertical direction with frequency ω and maximum deformation a . The vertical displacements of the structure are then

$$d_{s,y}(\hat{y}, t) = f\left(t - \frac{\hat{y} + H}{c_p}\right) - f\left(t + \frac{\hat{y} + H}{c_p}\right) \quad \text{with} \quad f(\tau) = \alpha \cos(\omega\tau),$$

whereas the vertical velocity component of the fluid is dependent on time only, given by $u_{f,y}(t) = \partial_t d_{\Sigma,y}(t)$ and results directly from the incompressibility constraint. Then, the fluid pressure is defined as

$$p_f(y, t) = \frac{(H - y)p_{\Sigma} + (y - d_{\Sigma,y})p_H}{H - d_{\Sigma,y}} [(H - y)p_{\Sigma} + (y - d_{\Sigma,y})p_H]$$

$$\text{with} \quad p_{\Sigma} = -(\lambda_s + 2\mu_s)\partial_{\hat{y}}d_{s,y}(0, t)$$

$$\text{and} \quad p_H = -\rho_f [H - d_{s,y}(0, t)] \partial_{tt}d_{s,y}(0, t) + (\lambda_s + 2\mu_s)\partial_{\hat{y}}d_{s,y}(0, t).$$

For details on the derivation of this exact solution, the interested reader is referred to the original publications [455, 456]. Then, we can directly proceed with defining the material parameters chosen. The fluid has a density of $\rho_f = 1 \text{ kg/m}^3$ and a dynamic viscosity of $\mu_f \equiv \eta_{\infty} = 0.1 \text{ Pa s}$. The linear elastic solid is given a density of $\rho_s = 100 \text{ kg/m}^3$, a Young's modulus of $E_s = 5 \text{ kPa}$ and a Poisson's ratio of $\nu_s = 0.3$. Regarding the remaining parameters to define the piston motion, we choose an amplitude of $a = 0.005$ and set a frequency of $\omega = \pi$.

On the exterior boundaries of the domain, we prescribe Dirichlet conditions, that is, set the structure's material displacements as given by the exact solution and linearly interpolate the prescribed interface displacement and zero displacement enforced on the upper boundary, that is, at $y = H/2$, to enforce appropriate fluid domain displacements. Nonzero Neumann boundary conditions are computed from the exact solution, prescribing tractions at the left and right boundaries of the fluid domain at $x = 0$ and $x = L$, while the fluid velocity is prescribed at $y = H/2$.

In accordance to previous experiments, we measure convergence rates in the maximum L^2 -error over all time steps $n = 1, \dots, N_t$ in the fluid velocity \mathbf{u}_f , fluid pressure p_f and structural displacement \mathbf{d}_s ,

$$e_{\mathbf{u}_f} := \max_{n=1, \dots, N_t} \left\| \mathbf{u}_f^{n+1} - \mathbf{u}_{f,h}^{n+1} \right\|_{L^2(\Omega_f^{n+1})}, \quad e_{p_f} := \max_{n=1, \dots, N_t} \left\| p_f^{n+1} - p_{f,h}^{n+1} \right\|_{L^2(\Omega_f^{n+1})},$$

$$e_{\mathbf{d}_s} := \max_{n=1, \dots, N_t} \left\| \mathbf{d}_s^{n+1} - \mathbf{d}_{s,h}^{n+1} \right\|_{L^2(\Omega_s)},$$

in the considered time interval from $t = 0$ to $T = 0.5$ s. For ease of comparison, these errors are compared to the ideal orders of convergence indicated by triangles in the respective plots.

The spatial discretisation is carried out employing equal-order interpolation via standard continuous Q_1Q_1 finite element pairs for all involved unknowns, i.e., the fluid velocity \mathbf{u}_f and pressure p_f , the auxiliary variable ψ used for divergence suppression, the fluid’s viscosity μ_f , the structural displacement \mathbf{d}_s and an analogous discretisation of the pressure trace used for the projection of boundary data. The time interval is decomposed into uniform steps, taking a BDF2 timestepping scheme for the fluid and generalised- α time integration schemes as laid out in Sec. 7.2. The fluid and structure are in this initial setup coupled implicitly using the standard implicit Dirichlet–Neumann coupling scheme (IDN) with Aitken’s relaxation, where we start off with a comparison of generalised- α time integrators.

In a first series of computations, the Newmark- β (N- β) scheme is used when refining the time step. Expected second-order convergence rates in the structural displacement, the fluid velocity and the fluid pressure are observed as can be seen in Fig. 8.2




Fig. 8.2: N- β time integration with implicit DN coupling yields expected second-order temporal convergence in solid displacements, fluid velocities and pressure.

The generalised- α time integration scheme allows switching between (i) the Newmark- β (N- β) method by Newmark [344], (ii) the HHT- α method introduced by Hilber et al. [345], (iii) the WBZ- α method by Wood et al. [346], and (iv) the CH- α method by Chung and Hulbert [342] as discussed in Sec. 5.3. The latter three schemes introduce high-frequency dissipation, controlled by the user-defined spectral radius in the high frequency limit ρ_∞ . Setting $\rho_\infty = 1$ yields the N- β scheme with no dissipation, whereas $\rho_\infty = 0$ results in maximum numerical damping introduced, the so-called asymptotic annihilation case. Choosing a practically relevant (user-specified) high-frequency dissipation by setting the spectral radius in the high-frequency limit, second-order convergence rates in the velocities, displacements and the pressure are maintained, but an increase in the saturation error is observed, which further depends on the specific scheme employed. In Fig. 8.3, convergence studies are shown, which showcase this effect using $\rho_\infty = 0.98$.




Fig. 8.3: Generalised- α time integration with implicit DN coupling: choosing a spectral radius in the high frequency limit of $\rho_\infty = 0.98$ results in an increased saturation error compared to $\rho_\infty = 1.0$ (equivalent to N- β time integration).

For $\rho_\infty = 1$, the HHT- α and WBZ- α timestepping schemes yield identical results as the N- β integrator, but employing maximum high-frequency dissipation $\rho_\infty = 0$, linear temporal convergence in the fluid pressure p_f is observed using the WBZ- α and CH- α methods as depicted in Fig. 8.4. For the HHT- α scheme, only values $\rho_\infty \in [1/2, 1]$ are permitted. Velocities and displacements converge quadratically, with the CH- α scheme reaching the saturation error in e_{d_s} and e_{u_f} the fastest.

Due to the fact that the structure's material acceleration $d_{tt}\mathbf{d}_s$ enters the PPE via the interface integral term $\langle \varphi \mathbf{n}_f, \rho_f d_{tt}\mathbf{d}_s^{m+1} \rangle_{\Sigma^t}$, introducing numerical dissipation affects the temporal orders of convergence in the pressure, since the accuracy in the structural acceleration reduces. For the considered generalised- α time integration schemes, only first-order accuracy in time is achieved when $\alpha_f \neq \alpha_m$ [343], limiting the temporal accuracy in the fluid pressure when using $\rho_\infty \neq 1$.




Fig. 8.4: Generalised- α time integration with IDN coupling and various ρ_∞ : $\rho_\infty \rightarrow 0$ leads to earlier error saturation and reduces the pressure’s temporal accuracy from 2 to 1.

Investigating temporal convergence rates of the implicit Robin–Robin coupling scheme (IRR), analogous relations hold and are thus not repeated here. The temporal convergence behavior of the semi-implicit variants of the Robin–Robin (SIRR) and Dirichlet–Neumann (SIDN) coupling schemes, however, needs to be commented. A key aspect when considering semi-implicit FSI schemes is preserving added-mass stability *and* accuracy. Treating the fluid mesh motion, the viscosity projection, divergence suppression, the fluid momentum balance equation and therefore also the pressure boundary projection explicit, temporal convergence rates are still preserved for the most part as can be seen from Fig. 8.5.

Comparing the semi-implicit variants to the implicit Dirichlet–Neumann scheme, we observe an increase in the saturation errors by an order of magnitude with the employed solver settings in this specific example. When using the SIRR coupling scheme, saturation errors are influenced by the choice of Robin parameters η_s^R and η_f^R , which we set to $\eta_f^R = \rho_s/\Delta t^n$ and $\eta_s^R = \rho_f/\Delta t^n$ (again, similar to [218]). The potential influence of Robin parameters dependent on the time step size is investigated in Sec. 8.3, whereas here, the influence is little.

These tests render semi-implicit schemes very attractive due to the substantial decrease in implicitly coupled subproblems while preserving accuracy to a great extent. This is especially relevant for practical applications as indicated by various authors (see, e.g., [192, 225, 441, 443–445, 447] and others). For now, a further discussion regarding possible performance gains is postponed to Secs. 8.3 and 8.5–8.8, which are devoted to comparing the coupling variants and practical application of the FSI scheme in the biomedical context.

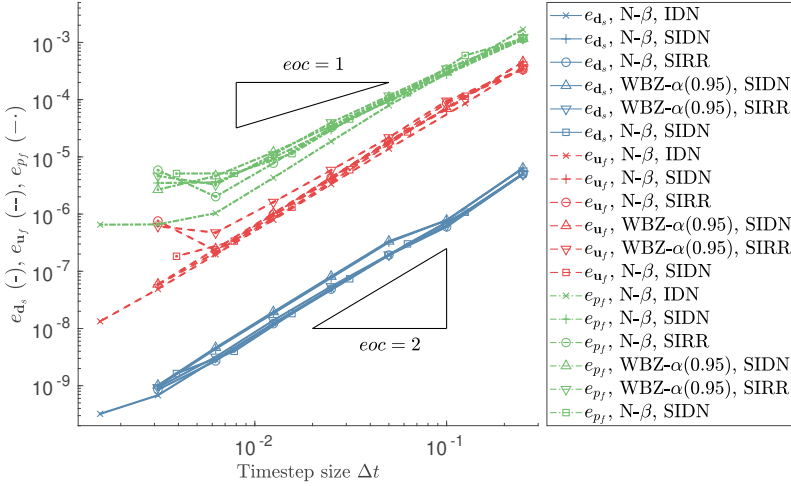


Fig. 8.5: $N\text{-}\beta$ and $WBZ\text{-}\alpha$ time integration with $\rho_\infty = 0.95$: SIDN and SIRR schemes result in almost identical errors compared to the IDN scheme.

8.2 Analytical solution: circular piston

To investigate spatial convergence rates, we consider an analytical solution presented by Serino et al. [455, 456], referring to the original work for a more detailed discussion. This numerical example describes the pulsating motion of a circular piston, with the structure in $\hat{r} = 0$ to the interface at $\hat{r}_\Sigma = r_\Sigma(t = 0)$, which drives the fluid surrounding the piston from $r_\Sigma(t)$ to $r = R$ as depicted in Fig. 8.6. Due to the pulsating piston, the region occupied by the fluid is changing, such that the incompressible fluid is allowed to enter and exit the domain freely.

For this simplified setup of an incompressible Newtonian fluid interacting with a linear elastic structure, an analytic solution can be formulated in terms of the structure's radial displacement

$$\mathbf{d}_{s,r}(\hat{r}, t) := \beta \mathcal{J}_1 \left(\frac{\omega \hat{r}}{c_p} \right) \sin(\omega t), \quad \text{such that} \quad r_\Sigma(t) = r_\Sigma^0 + \mathbf{d}_{s,r}(r_\Sigma^0, t),$$

with β scaling the amplitude, frequency ω , r_Σ^0 being the initial piston radius, c_p as defined in Eqn. (8.1) and \mathcal{J}_1 denoting the first-order Bessel function of the first kind, leading to a radial fluid velocity and corresponding pressure given as

$$\mathbf{u}_{f,r}(r, t) = R/r V(t), \quad \text{where} \quad V(t) = \frac{\omega \beta}{R} r_\Sigma(t) \mathcal{J}_1 \left(\frac{\omega r_\Sigma^0}{c_p} \right) \cos(\omega t),$$

$$p_f(r, t) = P(t) + \rho_f/2(1 - R^2/r^2) V^2(t) + \rho_f R \log(R/r) \partial_t V(t),$$

with $P(t) = -\rho_f/2(1 - R^2/r_\Sigma^2(t)) V^2(t) - \rho_f R \log(R/r_\Sigma(t)) \partial_t V(t)$

$$- \beta \sin(\omega t) \left[(\lambda_s + 2\mu_s)\omega/c_p \mathcal{J}'_1\left(\frac{\omega r_\Sigma^0}{c_p}\right) + \lambda_s/r_\Sigma^0 \mathcal{J}_1\left(\frac{\omega r_\Sigma^0}{c_p}\right) \right],$$

with \mathcal{J}'_1 denoting the first derivative of \mathcal{J}_1 .




Fig. 8.6: Circular piston problem: finite element mesh at refinement level 4.

The piston size and motion is set specifying $r_\Sigma^0 = 0.5$, $\omega = \pi$ and $\beta = 0.1$ and material parameters are selected as follows. The fluid density is chosen as $\rho_f = 1 \text{ g/m}^3$, the dynamic viscosity is $\mu_f = \eta_\infty = 0.5 \text{ mPa s}$, while the structure's density is $\rho_s = 1 \text{ kg/m}^3$ and elastic parameters are selected as $E_s = 100 \text{ kPa}$ and $\nu_s = 0.3$. To keep time integration errors at a minimum, we compute the solution from $t = 0$ to $t = 0.05 \text{ s}$, subdivided into 80 uniform time steps $\Delta t^n = 0.625 \text{ ms}$ and employ BDF2 and generalised- α time integration schemes combined with suitable, second-order accurate extrapolation.

The convergence studies are depicted in Figs. 8.7 and 8.8, including Q_1Q_1 , Q_2Q_1 and P_1P_1 finite elements paired with various (semi-)implicit Dirichlet–Neumann and Robin–Robin coupling schemes. Robin parameters are again selected as $\eta_f^R = \rho_s/\Delta t^n$ for the fluid and $\eta_s^R = \rho_f/\Delta t^n$ for the solid similar to [218].

With these parameters, the exact solution of \mathbf{d}_s up to time integration errors is directly recovered. However, these parameters are deliberately chosen to measure convergence rates of the fluid variables easily. As can be seen from Figs. 8.7 and 8.8, \mathbf{u}_f converges with rates of $roc = 2$ for any element pair or coupling scheme considered. The fluid pressure converges with rates of 1 for the equal order interpolations, i.e., Q_1Q_1 and P_1P_1 finite element pairs, which is exactly as expected when using (bi-)linear, equal-order finite elements.

With the Q_2Q_1 pairing, a convergence rate of 2 in the L^2 -error of the velocity might be interpreted as suboptimal by mistake, as one might expect 3 in an isoparametric setting, but in fact, both the L^2 -norm and the H^1 -seminorm converge with orders of 2 (cf. [3, 4, 20, 21]). The pressure convergence rate is slightly suboptimal, yielding

$eoc \approx 1.5$, where one might expect 2 (compare Fig. 3.1). This, however, is caused by the fact that linear finite elements are employed in the mapping even if the displacement approximation is higher-order accurate. The discrepancy is accepted at this point, since the focus herein lies on lower and equal-order interpolations. It is expected that the PPE-based split-step scheme used as the flow solver giving rates of 2 in this norm yields similar results in the FSI setting as well, if higher-order finite element mappings are employed.

Regarding the convergence rates of the solid displacements, setting different problem parameters yields the expected optimal rates in \mathbf{d}_s , which are omitted here for brevity. Balancing the solver's tolerances and physical parameters in order to showcase convergence rates of all involved fields at the same time is found rather cumbersome due to the problem setup and relative sizes of physical quantities and tolerance choices.




Fig. 8.7: Circular piston problem: spatial convergence rates obtained using Q_1Q_1 IDN, Q_1Q_1 SIRR and Q_1Q_1 SIDN element pairs and schemes.

8.3 Flow through a straight vessel

The numerical examples presented here consider blood flow through a straight vessel, which is a prominent benchmark example for FSI in the cardiovascular context (see, e.g., [256, 430, 463] among many others). Since this simplified setup serves as a testing bed for the coupling algorithms and various material models, the experiments are grouped into two scenarios: first, we report results for a pressure pulse in analogy to the classical benchmark configuration, only considering a short time interval. Second, pulsatile flow reproducing a flow regime in the physiological range is tackled. In both scenarios, the physical solution obtained with various material models is investigated and we compare the (semi-)implicit Dirichlet–Robin, Dirichlet–Neumann and Robin–Robin variants of the coupling scheme in terms of performance combining them with either IQN-ILS or Aitken acceleration.




Fig. 8.8: Circular piston problem: spatial convergence rates obtained using Q_1Q_1 IDN, P_1P_1 SIDN and Q_2Q_1 SIDN element pairs and schemes.

Two spatial discretisations as shown in Figs. 8.9 and 8.10 containing only hexahedra or tetrahedra are constructed, employing Q_1Q_1 or P_1P_1 finite element pairs. The computational domain has a length of $l = 5$ cm, with the lumen in its reference configuration having a radius of $r_i = 0.5$ cm and two tissue layers $\hat{\Omega}_{s,1}$ and $\hat{\Omega}_{s,2}$ of thickness $h_s = 0.05$ cm, leading to an outer radius of $r_o = 0.6$ cm.

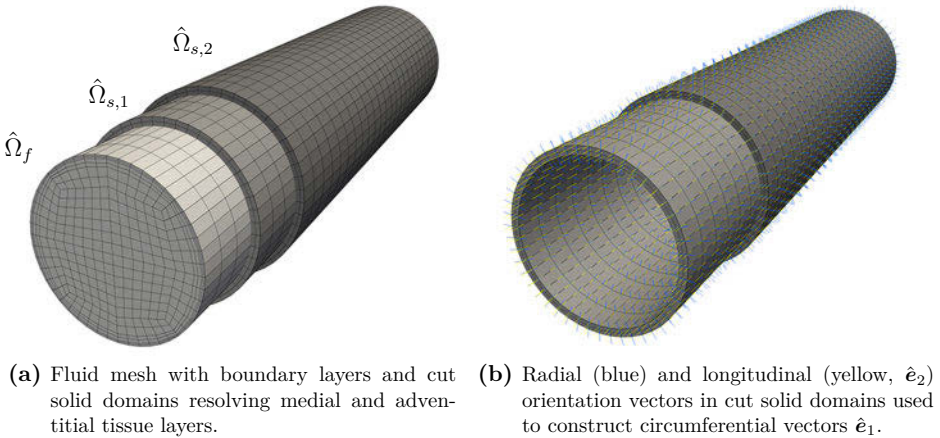



Fig. 8.9: Hexahedral finite element mesh considered for the straight vessel tests.

Concerning the material parameters, the fluid’s and tissue layers’ densities are selected as $\rho_f = 1060$ kg/m³ and $\rho_s = 1200$ kg/m³. Depending on the constitutive models applied, further parameters are specified. The Newtonian fluid with viscosity $\mu_f = 3.5$ mPa s (see, e.g., [18, 163, 463] regarding this choice) is compared to a fluid obeying the Carreau law with upper and lower viscosity limits $\eta_0 = 56$ mPa s and $\eta_\infty = 3.45$ mPa s and parameters $\lambda_f = 3.313$ s and $n = 0.3568$ taken from [508].



(a) Fluid mesh with boundary layers and cut solid domains resolving medial and adventitial tissue layers. (b) Longitudinal (blue, \hat{e}_2) and circumferential (yellow, \hat{e}_1) orientation vectors in cut solid domains.

Fig. 8.10: Tetrahedral finite element mesh considered for the straight vessel tests.

For the solid phase, linear elastic (E) material behaviour or a St. Venant–Kirchhoff model (SVK), both with Young’s modulus $E_s = 300$ kPa (if not specified otherwise) and Poisson ratio $\nu_s = 0.3$ are compared to neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) models. In the latter two cases, a Poisson’s ratio of $\nu_s = 0.499$ is selected to approximate incompressible material behaviour. The tissue layers are assigned individual shear rates of $\mu_{s,1} = 62.1$ kPa and $\mu_{s,2} = 21.6$ kPa, respectively. Fibre reinforcement is specified for the inner and outer tissue layer with individual parameters $k_1 = 1.4$ kPa, $k_2 = 22.1$ and $\kappa_{c,1} = 0.12$, $\alpha_{c,1} = 27.47^\circ$ or $\kappa_{c,2} = 0.25$, $\alpha_{c,2} = 52.88^\circ$ following [11, 338, 509].

The basis vectors for constructing the fibre orientation are depicted in Figs. 8.9(b) and 8.10(b), where the dual-Laplace approach is used for the mesh consisting of hexahedral elements only. In the tetrahedral mesh, the approach incorporating an averaged extrapolated interface normal is applied to obtain the longitudinal and circumferential vectors using standard parameters $\alpha_{\text{tol}} = 120^\circ$, $N_{\text{avg}} = 5$ and $h_{l,i} = h_n = 100$. For this geometry, both algorithms give very good results due to the simple topology, which could be further improved via spatial refinement.

8.3.1 Pressure pulse

In this first series of numerical tests, the fluid flow is driven by a smoothly ramped pressure drop enforced via $\mathbf{t}_f = (-p_{\text{in}}, 0, 0)^\top$, where

$$p_{\text{in}}(t) = \begin{cases} \bar{p}_{\text{in}} \left[\frac{1}{2} - \cos(500\pi t) \right] & \text{for } t \leq 0.002, \\ \bar{p}_{\text{in}} & \text{for } 0.002 < t < 0.007, \\ \bar{p}_{\text{in}} \left[\frac{1}{2} + \cos(200\pi t) \right] & \text{for } 0.007 \leq t \leq 0.009, \\ 0 & \text{otherwise,} \end{cases} \quad (8.2)$$

with $\bar{p}_{\text{in}} = 10 \text{ mmHg} \approx 1333.22 \text{ Pa}$ at the inlet and zero traction at $\hat{x} = l$. The structure is fixed at inlets and outlets, while zero Neumann conditions are enforced on the exterior surface of the tissue, i.e., at $\hat{y}^2 + \hat{z}^2 = r_o^2$.

Second-order time integration is employed combining 30 uniform steps $\Delta t^n = 0.5 \text{ ms}$ of BDF2/CH- α schemes with linear extrapolation of all fields, where $\rho_\infty = 0$ is set to counteract pressure oscillations in time caused by the sudden drop in pressure. At this point, neither GLS nor backflow stabilisation are considered in the fluid solver, and fluid and structure are coupled until relative convergence criteria of $\epsilon_{\text{rel}}^d = \epsilon_{\text{rel}}^u = \epsilon_{\text{rel}}^p = 10^{-4}$ or absolute criteria $\epsilon_{\text{abs}}^d = \epsilon_{\text{abs}}^u = \epsilon_{\text{abs}}^p = 10^{-7}$ are fulfilled. The acceleration schemes are initialised with relaxation parameter $\omega_0 = 0.01$, and Interface Quasi-Newton Inverse Least-Squares (IQN-ILS) acceleration is employed with the filtering parameter set to $\epsilon_{QR} = 10^{-16}$ and incorporating data from at most $q = 5$ time steps. In case of nonlinear tissue models, absolute and relative Newton tolerances in the structure's nonlinear solver of $\epsilon_{\text{rel}}^{nl} = 10^{-3}$ and $\epsilon_{\text{abs}}^{nl} = 10^{-7}$ are found sufficient.

First, snapshots of the solution are depicted in Fig. 8.11, showing the pressure pulse in a Carreau fluid traveling the vessel modelled by Holzapfel–Gasser–Ogden (HGO) material using Q_1Q_1 interpolation. Different from some variants of the benchmark, parameters in the haemodynamic regime yield larger displacements given the lower material stiffness, as shown in Fig. 8.12(a), taking the displacement norm in $\hat{\mathbf{x}} = (l/2, 0, r_i)$ as reference. The maximum displacement of $\approx 30 \%$ of the vessel wall's thickness leads to rather small strains, such that linear elasticity (E) and the St. Venant–Kirchhoff (SVK) model give almost identical values for the pressure and displacement in the reference point at the interface (see Figs. 8.12(a) and 8.12(b)). For the same reason, the stiffening effect of the fibre reinforcement is visible, but does not alter the system's response drastically. The fluid's viscosity, however, varies by a factor of 10 when employing the Carreau model as can be seen from Fig. 8.11, but in the present flow regime, the structure's displacement is dominated by the fluid's pressure. Also, the lack of backflow stabilisation leads to instabilities at the inlet due to the continuity of mass as the inlet pressure is rapidly reduced, but did not lead to divergence of the solver.

These results do not allow interpretations regarding the advanced constitutive models being necessary in the haemodynamic setting or not—they are only meant to showcase versatility of the partitioned coupling framework by switching between constitutive models with ease. In fact, even when considering a more realistic flow regime in this simplistic setup, one can not draw conclusions or judge upon complex material modelling being necessary.

Turning our attention now to the coupling scheme's iteration counts under variation of the constitutive models used, we observe that the semi-implicit schemes reduce overall steps executed. Since the PPE step does not feature a Robin interface condition owing to unfavourable stability properties of this variant, negligible differences between the semi-implicit Dirichlet–Neumann (SIDN) and semi-implicit Robin–Robin (SIRR) scheme are seen in Fig. 8.13. The Robin parameters are for these computations conservatively set to $\eta_f^R = 10^2 \rho_s / \Delta t^n$ and $\eta_s^R = 10^{-4} \rho_f / \Delta t^n$. Note here, that the Robin condition in the fluid momentum balance step does not improve convergence behavior, but rather increases coupling iteration counts. For $\eta_f^R \rightarrow \infty$, the Dirichlet–Robin scheme is recovered, and

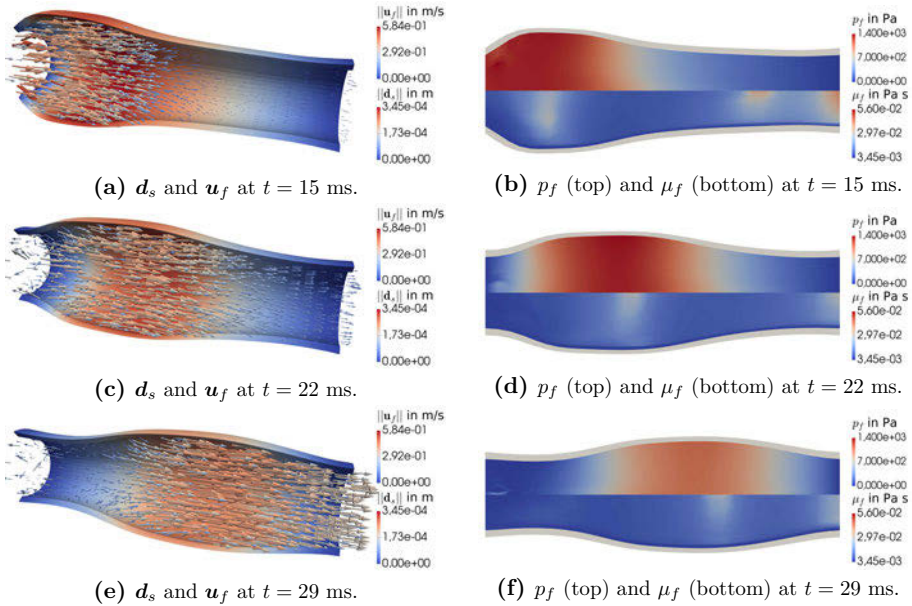


Fig. 8.11: Pressure pulse benchmark with Carreau fluid and HGO solid: snapshots of the solution at $t = 15, 22, 29$ ms (deformation scaled by 10): structural displacements d_s and fluid velocity u_f in cut tube (left), pressure p_f and viscosity μ_f in slice at $y = 0$ (right).

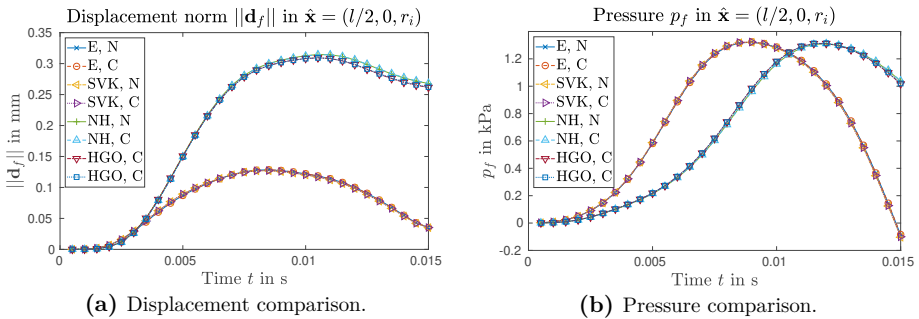


Fig. 8.12: Pressure-pulse benchmark with various constitutive models: linear elasticity (E), St. Venant–Kirchhoff (SVK), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) material models for the solid phase and Newtonian (N) or Carreau (C) fluids. Material nonlinearities have a limited effect.

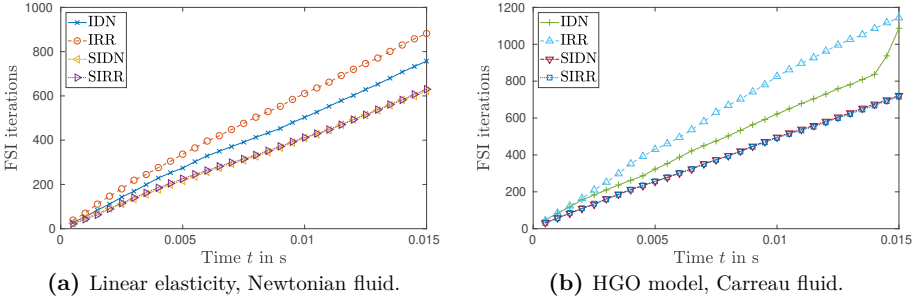


Fig. 8.13: Pressure pulse benchmark: accumulated FSI iterations using implicit Dirichlet–Neumann (IDN) or Robin–Robin (IRR) variants and semi-implicit counterparts (SIDN, SIRR), all with Aitken’s relaxation and Robin variants with parameters $\eta_f^R = 10^2 \rho_s / \Delta t^n$ and $\eta_s^R = 10^{-4} \rho_f / \Delta t^n$.




Fig. 8.14: Pressure pulse benchmark: accumulated FSI iterations using a fully implicit Robin–Neumann coupling scheme with $\eta_f^R = \alpha_R \rho_s / \Delta t$, linear elasticity and a Newtonian fluid.

if $\eta_s^R \rightarrow 0$ additionally, the classical Dirichlet–Neumann scheme results, yielding the respective coupling iteration counts. This rather surprising fact does not go in line with the literature (see, e.g., [217, 218, 220, 462]), but Robin parameters might be set inappropriate for this specific setup and moreover, studies on Robin-accelerated pressure-projection or PPE-based schemes are scarcely found in literature.

To investigate this irritating aspect further, we choose a fully implicit Robin–Neumann scheme, scaling the Robin parameter used in the fluid momentum balance step with α_R to give $\eta_f^R = \alpha_R \rho_s / \Delta t^n$. The results presented in Fig. 8.14 indicate, that the Robin condition in the fluid subproblem does not accelerate convergence *at all*, but rather leads to higher iteration counts or possible divergence if η_f^R is chosen too low. Interestingly, the works by Badia et al. [218] and Spenke et al. [462] considered coupled velocity-pressure formulations, which we suspect to influence how Robin conditions in the fluid subproblem affect the coupling scheme. In our PPE-based framework, again, Robin conditions in the PPE did not result in a stable scheme, and the pressure mostly dominating the interface tractions reduces influence of the fluid momentum step’s Robin condition—which is even more so the case for semi-implicit variants.

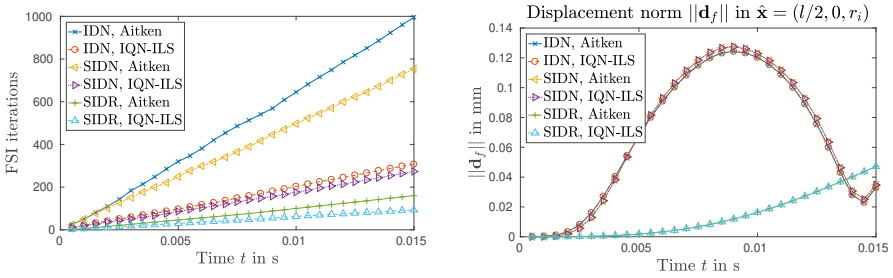


Fig. 8.15: Pressure pulse benchmark: accumulated FSI coupling steps (left) and displacement in reference point (right). Semi-implicit Dirichlet–Neumann (SIDN) and Robin–Neumann (SIRD) variants and IQN-ILS acceleration outperform the classical implicit Dirichlet–Neumann (IDN) scheme and Aitken relaxation in terms of coupling iterations. Robin interface conditions accelerate convergence further, but the solution might change drastically with large η_s^R .

Therefore, options for accelerating convergence considered within this work are (i) IQN-ILS or Aitken’s acceleration and (ii) the Robin interface condition in the structure’s momentum balance equation. For the sake of brevity, we limit discussions to a Newtonian fluid and linear elastic solid, which is admissible in this simplified setup, since nonlinearities do not affect the coupling scheme’s performance. To showcase applicability of P_1P_1 interpolation, these tests are carried out using the tetrahedral mesh. Aitken’s acceleration and the IQN-ILS method and the (semi-)implicit variants are compared in Fig. 8.15. As expected, IQN-ILS acceleration outperforms classical Aitken’s method, and the semi-implicit variants do not only reduce the number of required coupling iterations, but are naturally much cheaper in execution per coupling iteration. In this specific example, this combination resulted in a reduction of computing times by a factor of 20, as highlighted in Tab. 8.1.

However, care has to be taken when going for the semi-implicit Dirichlet–Robin (SIDR) variant, since the obtained solution might be altered taking a Robin parameter too large. This is caused by shifting from enforcing both continuity of tractions (interface traction in the structure’s momentum balance) and velocities (interface integral term in the PPE) to enforcing the continuity of velocities only. This change in weighting of semi-implicit dynamic and explicit kinematic coupling conditions dependent on the Robin parameter η_s^R , essentially reducing the implicit pressure’s significance to the solid’s momentum balance, bears great potential for acceleration, but may also alter results quite drastically.

Under these viewpoints, Fig. 8.15 clearly indicates that (i) the acceleration schemes applied (Aitken or IQN-ILS) approximately yield the same solution given identical stopping criteria, (ii) the effect of semi-implicit coupling is almost insignificant in terms of the physical solution obtained, but (iii) the Robin interface condition in the structure’s momentum balance equation might influence the system’s response. It is thus obvious that selecting η_s^R determines the computational results in a variety of ways, i.e., influences the coupling convergence speed *and* physical solution if chosen inappropriately.

Tab. 8.1: Total FSI iterations and computing time for pressure pulse in a straight cylindrical vessel: the semi-implicit Dirichlet–Neumann (SIDN) and Dirichlet–Robin (SIDR) schemes and IQN-ILS acceleration outperform the implicit Dirichlet–Neumann (IDN) scheme and Aitken relaxation.

	abs. FSI steps	rel. FSI steps	abs. time	rel. time
IDN, Aitken	971	10.67	1.03×10^3 s	20.44
IDN, IQN-ILS	296	3.25	3.06×10^2 s	6.07
SIDN, Aitken	730	8.02	3.19×10^2 s	6.33
SIDN, IQN-ILS	262	2.89	1.27×10^2 s	2.52
SIDR, Aitken	158	1.74	7.46×10^1 s	1.48
SIDR, IQN-ILS	91	1.00	5.04×10^1 s	1.00

The reason behind the decreasing coupling iterations when increasing the Robin parameter η_s^R is rooted in the integral term in the structure’s momentum balance equation,

$$\alpha'_f \langle \varphi, \eta_s^R (\mathbf{u}_f^* - \dot{\mathbf{d}}_s^{n+1}) + J_f \boldsymbol{\sigma}_f (\mathbf{u}_f^*, p_f^{n+1, k+1}, \mu_f^{n+1}) \mathbf{F}_f^{-\top} \hat{\mathbf{n}}_s \rangle_{\hat{\Sigma}},$$

transitioning to an explicit kinematic constraint as $\eta_s^R \rightarrow \infty$ enforced via a penalty method. Simultaneously, decreasing the time step size Δt^n reduces temporal errors involved in this term, given that \mathbf{u}_f^* approximates \mathbf{u}_f^{n+1} better. To investigate this relation in more detail, we again introduce the Robin parameter as $\eta_s^R = \alpha_R \rho_f / \Delta t^n$ with α_R allowing a transition from the classical Neumann interface condition as $\eta_s^R \rightarrow 0$ to the explicit penalty enforcing continuity of velocities.

Fig. 8.16 shows the expected influence on the FSI coupling iteration count, whereas the obtained solution (here exemplarily shown in terms of the displacement in the reference point) might not show a comparably uniform trend when decreasing the Robin parameter using the fully second-order accurate scheme. Here, we see a nonlinear relation between the Robin parameter and the resulting displacement in the reference point when decreasing the Robin parameter from 10 to 0. When choosing a lower-order extrapolation of the fluid velocity, which we will denote as m_u , this effect is mitigated as can be seen from corresponding experiments summarised in Fig. 8.17. The solutions corresponding to the Robin parameters transition much smoother and more uniform to the Dirichlet–Neumann case.

These observations are naturally limited to this test case, but this numerical evidence suggests that a higher-order extrapolation in the Robin interface term increases temporal accuracy, but at the same time introduces stricter limitations on the time step size. An influence of the inlet and outlet pressure can be ruled out, since the pressure drop is weakly enforced via standard boundary integrals, which becomes relevant in the pulsating flow with Windkessel outlets considered in Sec. 8.3.2.

But before moving on to a more realistic setting, let us demonstrate the important relations that (i) keeping Δt^n fixed and decreasing η_s^R transitions to Dirichlet–Neumann

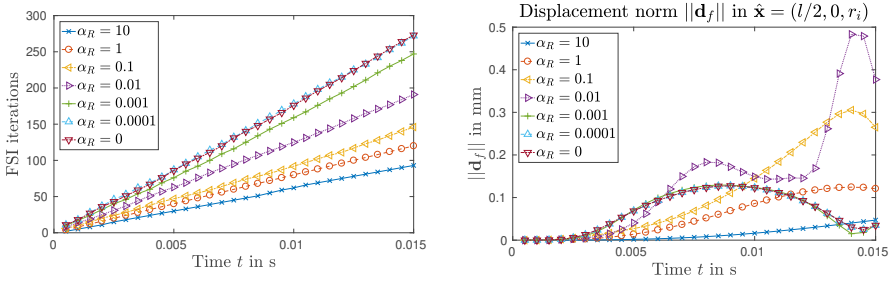


Fig. 8.16: Pressure pulse benchmark: accumulated FSI coupling steps (left) and displacement in reference point (right) obtained via the SIDR scheme with $\eta_s^R = \alpha_R \rho_f / \Delta t$ and second-order fluid velocity extrapolation ($m_u = 2$).

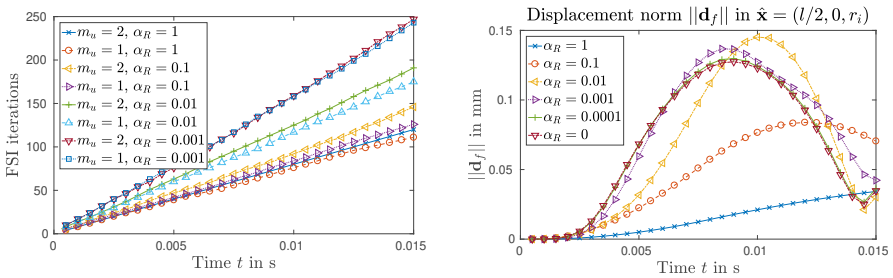


Fig. 8.17: Pressure pulse benchmark: accumulated FSI coupling steps with fluid velocity extrapolation $m_u = 1$ or 2 (left) and displacement in reference point with $m_u = 1$ (right), obtained via the SIDR using $\eta_s^R = \alpha_R \rho_f / \Delta t$.

coupling, where $m_u = 1$ improves stability and (ii) choosing smaller Δt^n and holding η_s^R constant reduces differences between Dirichlet–Neumann and Dirichlet–Robin schemes. To showcase these relations, we compare SIDR and SIDN schemes taking the latter as reference and computing the mean difference in the reference point’s displacement on the interface at $\hat{\mathbf{x}} = (l/2, 0, r_i)$ over all time steps using identical algorithmic parameters only differing in the Robin parameter used. Results are given in Tab. 8.2, showing that second-order extrapolation of \mathbf{u}_f^* ($m_u = 2$) leads to divergent simulations with Robin parameters too high due to explicit kinematic coupling but reduces the difference to the SIDN reference solution. For a lower-order extrapolation in time, $m_u = 1$, no simulation diverges and differences decrease with $\Delta t^n \rightarrow 0$ as expected.

Tab. 8.2: Average difference in displacements in $\hat{\mathbf{x}} = (l/2, 0, r_i)$ over all time steps between the SIDN reference and SIDR schemes with Robin parameter $\eta_s^R \neq \eta_s^R(\Delta t^n)$. The fully second-order scheme ($m = m_u = 2$) is less robust than choosing $m_u = 1$, i.e., $\mathbf{u}_f^* = \mathbf{u}_f^n$. With $\Delta t^n \rightarrow 0$, differences between SIDN and SIDR vanish.

m_u	η_s^R	$\Delta t^n = 0.125$ ms	$\Delta t^n = 0.25$ ms	$\Delta t^n = 5$ ms	$\Delta t^n = 1$ ms
2	2000	diverged	diverged	5.77×10^{-2}	6.04×10^{-2}
2	200	diverged	diverged	1.05×10^{-1}	diverged
2	20	5.42×10^{-2}	diverged	9.80×10^{-2}	diverged
2	2	7.36×10^{-4}	2.40×10^{-3}	7.58×10^{-3}	1.76×10^{-2}
2	0.2	7.08×10^{-5}	2.13×10^{-4}	4.98×10^{-4}	1.76×10^{-2}
1	2000	6.09×10^{-2}	6.25×10^{-2}	6.24×10^{-2}	6.15×10^{-2}
1	200	6.40×10^{-2}	5.75×10^{-2}	5.25×10^{-2}	5.19×10^{-2}
1	20	3.24×10^{-2}	3.98×10^{-2}	4.64×10^{-2}	4.89×10^{-2}
1	2	4.18×10^{-3}	7.28×10^{-3}	1.37×10^{-2}	1.91×10^{-2}
1	0.2	4.78×10^{-4}	7.64×10^{-4}	1.22×10^{-3}	$2.11 \cdot 10^{-3}$

These results are promising in the sense that using Robin interface conditions, accuracy can be preserved carefully setting η_s^R . That being said, let us briefly add some remarks regarding this parameter choice. The starting point $\eta_s^R = \alpha_R \rho_f / \Delta t^n$ derived by Badia et al. [218] is based on simplified model problems, where α_R is a suitably scaled maximum eigenvalue of the added-mass operator. From this viewpoint, the α_R chosen in the tests is rather high. On the other hand, our scheme shares the penalty-like term on the explicit kinematic coupling in the structure’s momentum balance step with the method proposed by Astorino et al. [225], which uses a projection scheme to advance the fluid in time. In their formulation, however, coupling conditions are enforced via Nitsche’s method with a parameter $10^{\mu_f/h_e}$. Now, following this interpretation and using the current parameters, $\eta_s^R \approx O(1000)$ lies in the expected range and would require $\alpha_R \approx O(0.01)$ in $\eta_s^R = \alpha_R \rho_f / \Delta t^n$.

In addition to that, we also want to emphasise that neither of the two solutions are converged in space or time, such that computing the mean difference in a single point

can merely indicate the actual errors in the solutions presented. Connecting back to the discussion before, a Robin-accelerated coupling scheme with high Robin parameter and extrapolation \mathbf{u}_f^* of desired order m_u is still considered being of the same temporal order of accuracy as a comparable Dirichlet–Neumann scheme. The specific temporal error might differ, but as the time step and element sizes go to zero, both the kinematic and dynamic interface conditions are fulfilled and the Dirichlet–Robin and Dirichlet–Neumann schemes both converge to the same solution.

8.3.2 Pulsatile flow

Now, having thoroughly investigated the pressure pulse benchmark, let us move on to a more realistic problem setup using the same basic geometry. Based on the insights just gained in the foregoing studies, we limit our experiments here to IQN-ILS acceleration and Robin conditions enforced in the structure’s momentum balance equation. It is of central interest at this point to investigate under which conditions the semi-implicit Dirichlet–Robin scheme can be applied without heavily altering the obtained results and how it affects the coupling scheme’s convergence in a realistic scenario.

To create a setting resembling blood flow through a vessel more accurately, we change boundary conditions for the fluid and solid subproblems. Here, a parabolic profile is prescribed at the inlet with $\mathbf{u}_f = (u_1, 0, 0)^\top$ and

$$u_1(r, t) = \bar{u}_{\text{in}}(t) (1 - r^2/r_i^2), \quad \bar{u}_{\text{in}}(t) = \eta_t(t) [4 + 6 \sin^2(2\pi t + 0.1\pi)],$$

which smoothly ramps the velocity from an initially quiescent state via

$$\eta_t(t) = \begin{cases} \sin^2\left(\frac{\pi t}{2T_r}\right) & \text{for } t \leq T_r, \\ 1 & \text{otherwise,} \end{cases} \quad (8.3)$$

from $t = 0$ to $T_r = 0.2$ s. Then, the inlet velocity reaches a maximum of $u_{\text{max}} \approx 0.5$ m/s as can be seen from Fig. 8.18, showing the time evolution of \bar{u}_{in} .

At the outlet, i.e., at $\hat{x} = l$, the pressure is determined via a three-element Windkessel model with resistances of $R_p = 1.5 \times 10^8$ Pa s/m³, $R_d = 8 \times 10^8$ Pa s/m³ and a capacitance of $C = 1.5 \times 10^{-10}$ m³/Pa together with $p_p|_{t=0} = p_d = 0$ Pa to reproduce a pressure fluctuation in the physiologically relevant range from ≈ 13 to 9 kPa, that is, $\approx 97, 5$ to 67.5 mmHg. To withstand this increased pressure load compared to the previous setting, viscoelastic support with $k_e = 10^7$ N/m³ and $c_e = 10^5$ Ns/m³ is taken into account. In this simplified setup, both the external pressure in the Robin condition on the solid’s exterior and the prestress in the vessel are neglected. Further, the structure’s Young’s modulus is increased to $E_s = 400$ kPa, while the remaining tissue and fluid parameters are not altered.




Fig. 8.18: Flow through a straight vessel: maximum inlet velocity \bar{u}_{in} scaling the parabolic inflow profile during three cycles from $t = 0$ to $t = 1.5$ s.

Regarding changes in the numerical setup, GLS and backflow stabilisation are combined with adaptive time step control. The time step size is chosen to reach a desired CFL_{\max}

$$\Delta t^{n+1} = \max \left\{ \Delta t_{\max}, 0.98 \Delta t^n \min \left\{ 1.02, \frac{CFL_{\max}}{\max_e \{CFL_e\}} \right\} \right\},$$

satisfying the CFL condition CFL_e

$$CFL_e = \max_i \left\{ \frac{|u_i - u_{m,i}| \Delta t^n}{h_i} \right\}, \quad i = 1, \dots, d,$$

with the directional element length h_i , see Eqns. (3.37) and (3.38), without repeating time steps violating this condition.

The adaptive timestepping is started from $\Delta t^0 = 5 \times 10^{-4}$ and the time interval from $t = 0$ to $t = 1.5$ s is chosen, i.e., a total of three cardiac cycles. The Robin parameter includes a scaling coefficient α_R such that $\eta_s^R = \alpha_R \frac{\rho_f}{\Delta t^n}$ and Newton tolerances in the structure's nonlinear solver set to $\epsilon_{\text{rel}}^{nl} = 10^{-4}$ and $\epsilon_{\text{abs}}^{nl} = 10^{-7}$ are found sufficient.

Starting again with a comparison of constitutive models, we employ linear elastic (E), St. Venant–Kirchhoff (SVK), neo-Hookean (NH) and Holzapfel–Gasser–Ogden (HGO) solid models, combining them with a Newtonian (N) or a Carreau (C) fluid with parameters as introduced in the beginning of this numerical example. Based on the previous findings, we combine the BDF2 and CH- α ($\rho_\infty = 0$) schemes with linear extrapolation for all involved fields except for the fluid velocity ($m_u = 1$) to improve robustness. In the reference point $\hat{\mathbf{x}} = (l/2, 0, r_i)$ on the fluid–structure interface, we compare the structural displacement \mathbf{d}_s and the fluid pressure p_f obtained.

The chosen settings lead to a displacement in the reference point of ≈ 1 mm, which equals 20% of the lumen's radius and 100% of the combined layers' thickness. In this more realistic setup leading to larger displacements and strains, differences are observed between displacement responses depending on the structural constitutive equation employed, as can be seen from Fig. 8.19. The measured quantities are not influenced by the fluid model, but the material parameters used for the different solids lead to significant changes. The fibre contributions in the HGO model render it stiffer, especially under

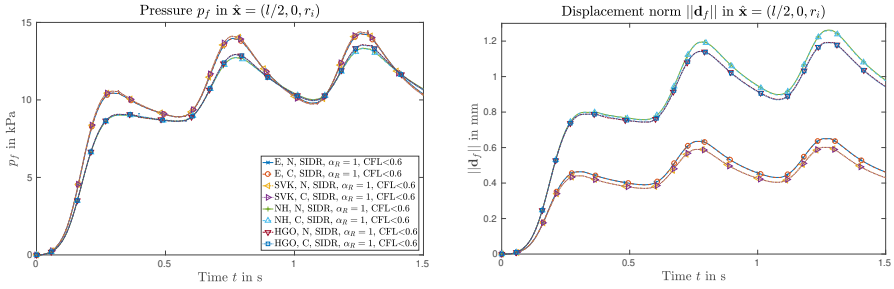


Fig. 8.19: Flow through a straight vessel: pressure (left) and displacement (right) in the reference point using the SIDR scheme with $\eta_s^R = \alpha_R \rho_f / \Delta t^n$ and adaptive timestepping ($\text{CFL}_{\max} = 0.6$). Comparison of linear elasticity (E), St. Venant–Kirchhoff (SVK), neo-Hookean (NH) and Holzapfel–Gasser–Ogden (HGO) solid models, combining them with Newtonian (N) or a Carreau (C) fluids.

Tab. 8.3: Flow through a straight vessel: comparison of various constitutive models in terms of time step number N_t and FSI coupling iterations N_{FSI} , mean FSI coupling steps per time step using the SIDR coupling scheme with adaptive timestepping targeting $\text{CFL}_{\max} = 0.6$.

	E,N	E,C	SVK,N	SVK,C	NH,N	NH,C	HGO,N	HGO,C
N_t	1125	1122	1128	1124	1084	1085	1084	1085
N_{FSI}	2551	2517	2587	2510	2408	2372	2367	2482
N_{FSI}/N_t	2.27	2.24	2.29	2.23	2.22	2.19	2.18	2.29

large fibre strains and the linear elastic solid with its linear strain measure overestimates the displacement. The differences in expansion lead to an altered volumetric flow over the outlet boundary, such that the transient pressure wave is changed as well.

Comparing the overall number of executed coupling steps as depicted in Fig. 8.20, we observe a negligible difference in the time steps and coupling iterations needed. As highlighted in Tab. 8.3, the mean number of FSI iterations per time step for all considered combinations is ≈ 2.2 , which is on average enough to reduce the pressure and displacement interface residuals by 10^4 owing to the combination of the SIDR scheme, IQN-ILS acceleration and a second-order accurate extrapolation of \mathbf{d}_s and p_f used as initial guess.

With this, let us turn our attention to a detailed comparison of SIDN and SIDR schemes using only St. Venant–Kirchhoff and Newtonian models to reduce the numerical effort. To diminish the influence of the exterior viscoelastic support, set $c_e = k_e = 0$ and consider the first-order scheme (BDF1 and trivial “extrapolation”) for increased temporal stability in this case with zero added damping. We observe once again that the RN variant converges to the DN one when, e.g., $\eta_s^R = 20000$ is held fixed and the time step size is reduced as Fig. 8.21 highlights.




Fig. 8.20: Flow through a straight vessel: accumulated FSI iterations using the SIDR scheme with $\eta_s^R = \alpha_R \rho_f / \Delta t^n$ and adaptive timestepping ($CFL_{max} = 0.6$). Comparison of linear elasticity (E), St. Venant–Kirchhoff (SVK), neo-Hookean (NH) and Holzapfel–Gasser–Ogden (HGO) solid models, combining them with Newtonian (N) or Carreau (C) fluids.

Therein, a zoom on the last cardiac cycle’s peak displacement and pressure ($t = 1.26$ s) shows a decrease in the phase shift as $\Delta t \rightarrow 0$ for the SIDR scheme. Being more sensible to added-mass instabilities, the pressure computed by the SIDN scheme shows temporal oscillations around the assumed exact solution given the small time step size. The Dirichlet–Neumann scheme is increasingly unstable as $\Delta t \rightarrow 0$, but the clearly visible pressure kinks in both the SIDR and SIDN variants can be linked to backflow stabilisation suddenly being activated in case $\mathbf{u}_f \cdot \mathbf{n}_f < 0$ at the outlet.

With adaptive timestepping according to the CFL condition, a good match in phase and amplitude are achieved even for $\alpha_R = 10$, which gave almost unacceptable results in the pressure pulse benchmark. In this more realistic scenario with much lower frequencies dominant in the system’s response, all obtained reference values match well on the scale




Fig. 8.21: Flow through a straight vessel: zoom-in on pressure and displacement in reference point comparing SIDR variants with $\eta_s^R = 20000$ to the SIDN scheme. Solutions converge as $\Delta t \rightarrow 0$ and Robin variants suffer less from added-mass instability, while the abruptly activated backflow stabilisation yields a pressure kink.

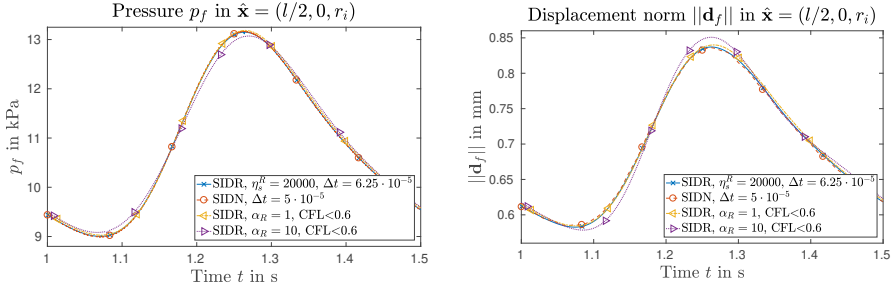


Fig. 8.22: Flow through a straight vessel: displacement and pressure in the reference point in the third cardiac cycle ($t = 1$ to 1.5 s). Adaptive timestepping with $\text{CFL}_{\max} = 0.6$ and SIDN or SIDR coupling with $\eta_s^R = \alpha_R \frac{\rho_f}{\Delta t^n}$ yields good agreement in phase and amplitude.

of one cardiac cycle's fluctuations as shown in Fig. 8.22. The fluctuations in the pressure computed with the SIDN scheme are also no longer visible on this scale and are in fact reduced due to larger time steps used.

The number of executed time steps of course drastically differs between the schemes with or without adaptive timestepping. A central aspect is here, that the Robin interface condition increases stability, easing adaptive time step selection. Given the excellent performance of the IQN-ILS acceleration with filtering in this setup, both the Dirichlet–Neumann and the Dirichlet–Robin variants converge in 2 to 4 coupling iterations per time step. The computational cost per time step is almost identical as well, such that the number of executed time steps to reach the target interval's end dominates the computational cost as can be seen from Tab. 8.4.

Therefore, the SIDR variant targeting $\text{CFL}_{\max} = 0.6$ and consequently needing 1880 time steps being the fastest option is not surprising. Compared to the SIDN reference (with small fixed time step size), it is roughly 26 times faster. Taking an SIDR scheme with fixed time step size of $\Delta t = 2.5 \times 10^{-4}$ yielding a comparable solution accuracy as reference, a speed-up of 5 is achieved. Such numbers are naturally limited to this specific example, since the achievable performance increase is problem-dependent. Up to some degree, however, an improvement can be expected if a periodic inflow with a large enough difference between minimum and maximum CFL numbers is considered, since a fixed time step size needs to account for the maximum CFL over all time steps in the considered time interval.

Let us close these investigations and summarise our observations:

1. In the cardiovascular setting, the semi-implicit coupling schemes are in good agreement with the fully implicit reference solutions with suitable time step size small enough, obeying a standard CFL-condition. For practical applications, the SIDR scheme can then yield substantial speed-ups compared to the SIDN counterpart while being more robust with respect to the added-mass effect.
2. In the PPE-based FSI scheme, $\eta_s^R > 1$ has large potential to reduce iteration counts as it shifts to a fully explicit kinematic coupling. Hence, the obtained solution is

altered when $\eta_s^R \gg 1$. But even then, the Dirichlet–Robin variant converges to the Dirichlet–Neumann reference as $\Delta t^n \rightarrow 0$.

3. Considering Windkessel outlets and/or adaptive timestepping, expected differences in the first- and second-order timestepping and extrapolation scheme’s stability properties become visible. Such effects can be reduced by considering time-averaged Robin parameters, averaging of adapted time step sizes (see, e.g., [151]) or considering adaptive two-element Windkessel models [272].

Tab. 8.4: Flow through a straight vessel: comparison of SIDN and SIDR coupling schemes with uniform or adaptive timestepping ($\text{CFL}_{\max} = 0.6$) in terms of absolute (relative) number of time steps, FSI coupling iterations and computing time. Adaptive timestepping allows reaching comparable accuracy in much fewer time steps, yielding a tremendous speed-up.

	time steps	FSI steps	time in 10^3 s
SIDN, $\Delta t = 1 \times 10^{-4}$ s	15000 (12.7)	45101 (12.8)	89.5 (12.8)
SIDN, $\Delta t = 5 \times 10^{-5}$ s	30000 (25.4)	91431 (26.0)	180.3 (25.8)
SIDR ($\eta_s^R = 20000$), $\Delta t = 1 \times 10^{-3}$ s	1500 (1.3)	4474 (1.3)	8.9 (1.3)
SIDR ($\eta_s^R = 20000$), $\Delta t = 5 \times 10^{-4}$ s	3000 (2.5)	8950 (2.5)	17.8 (2.5)
SIDR ($\eta_s^R = 20000$), $\Delta t = 2.5 \times 10^{-4}$ s	6000 (5.1)	17943 (5.0)	35.7 (5.1)
SIDR ($\eta_s^R = 20000$), $\Delta t = 1.25 \times 10^{-4}$ s	12000 (10.2)	35956 (10.2)	71.5 (10.2)
SIDR ($\eta_s^R = 20000$), $\Delta t = 6.25 \times 10^{-5}$ s	24000 (20.3)	72150 (20.5)	143.2 (20.5)
SIDR ($\alpha_R = 1$), $\text{CFL}_{\max} = 0.6$	1180 (1.0)	3516 (1.0)	7.0 (1.0)
SIDR ($\alpha_R = 10$), $\text{CFL}_{\max} = 0.6$	1184 (1.0)	3517 (1.0)	7.0 (1.0)

8.4 ALE extension: beam in a channel

As already pointed out, ALE extensions and their comparison can be examined considering prescribed interface data only, allowing more efficient investigations on the options presented in Sec. 7.5. Here, we focus on the various extension types being harmonic extension and pseudo (hyper-/linear) elasticity, all of which are formulated on $\hat{\Omega}_f$ and alternatively on Ω_f^t . Moreover, we shed some light on the influence of additional options such as Jacobian stiffening or the parameters in element-size-based assignment of diffusion or pseudo material coefficients. Note that we purposely present a multitude of possible combinations to give detailed insights into the individual choice’s properties as a robust ALE extension is *vital* in FSI. However, the ALE mapping is within this work investigated independent of the remaining FSI algorithm, which is again considered from Sec. 8.5 onwards.

The deformation of a beam in a channel under cross-flow is considered as proposed by Jendoubi et al. [467], which is inspired by [510]. The computational domain has a



Fig. 8.23: Beam in cross-flow: fluid (light grey) and structural (dark grey) domains in initial (top) and final (bottom) configurations at $t = 1$.

length of $L = 6$ and a height of $H = 1$ with a beam positioned at $L/2$, length $l = 0.45$ and height $h = 0.01$, such that the structural domain $\hat{\Omega}_s$ occupies $[3, 3.01] \times [0, 0.45]$ as depicted in Fig. 8.23. The structure undergoes a prescribed deformation \mathbf{d}_s which depends on time $t \in (0, 1]$ as

$$\mathbf{d}_s = \left(0.55 t \frac{\hat{y}^2}{l^2}, 0 \right)^\top,$$

where \hat{y} denotes the vertical coordinate which corresponds to the beam's longitudinal coordinate. Thus, the beam's deflection is a quadratic function of its length as shown in Fig. 8.23 being zero at the bottom of the channel and maximal at the beam's tip at all times.

A finite element mesh containing only quadrilaterals is constructed by the help of Gmsh [511], with a boundary layer to resolve steep velocity gradients as would be expected in the flow field typical for such a configuration. A close-up of the computational grid and corresponding quality criteria as defined in Eqn. (7.99) with $J_f = 1$ is shown in Fig. 8.24. Compared to the setup in Jendoubi et al. [467], we decompose the time interval into 1000 steps, motivated by the fact that in practical FSI problems, time steps are likely to be small compared to the system's dominating response frequencies. Then, the time history of the domain motion can be used for linearisation, stiffness updates and to generate good initial guesses for any involved (non-)linear solvers.

We compare individual extension types, Jacobian stiffening and stiffness scaling dependent on the element size for formulations on the initial reference configuration $\hat{\Omega}_f$ and the current domain Ω_f^t based on Q_f , measuring the quality in the final Ω_f^{n+1} , which would be considered in the following time step's PPE and momentum balance equations within the full FSI scheme. As pointed out by Shamanskiy and Simeon [471], accumulated distortion effects do not influence formulations on $\hat{\Omega}_f$. On the other side, formulations based on Ω_f^n include stiffening based on the updated element sizes and hence respect changes in the elements' volumes, when searching for an increment to the displacement field \mathbf{d}_f from Ω_f^n to Ω_f^{n+1} .




Fig. 8.24: Beam in cross-flow: finite element mesh with boundary layer zoomed-in at the beam’s tip with fluid (light grey) and structural (dark grey) domains resolved by the grid (left) and initial mesh quality criterion $Q_{J,e}$ with $\mathbf{d}_f = \mathbf{0}$ in (7.99) (right, high values in red).

The quality criterion $Q_{J,e}$ is sampled in each integration point of the elements, such that final configurations at $t = 1$ might be reached despite the fact that individual points in $\hat{\Omega}_f$ may exist, where the ALE map is not invertible, i.e., $J_f < 0$. This discrepancy is ignored, since we are interested in a relative score for comparison of individual mapping types only and thus allow configurations such as the one shown in Fig. 8.25, which still gives a quality criterion of $Q_J := \min_{e=1,\dots,N_{e,f}} Q_{J,e} \approx 0.1$.




Fig. 8.25: Beam in cross-flow: mesh quality criterion $Q_{J,e}$ (7.99) in Ω_f^n , beam region (left) and zoom-in on the beam’s tip (right).

In a first series of computations, the parameters entering the element-size-based diffusion or stiffness coefficients in harmonic extension and pseudo linear elasticity as given by Eqn. (7.80) are varied. More specifically, the minimal coefficient is held fixed at $c_{\min} = 1$, while the maximum c_{\max} , used in the smallest element with element size $\min_{e=1,\dots,N_{e,f}} h_e$, and the additional scaling parameter c_0 , modifying the transition from c_{\min} to c_{\max} , are varied from 10 to 10^9 . The only parameter remaining is the Poisson’s ratio in pseudo linear elasticity, where $\nu_m = 0.3$ is chosen as a starting point. The simulation is stopped, whenever $Q_J > 0$ is violated or $t = 1$ is reached, and the time corresponding to the last successfully completed step is recorded. In this way, each combination of c_{\max} and c_0 yields a time t_{end} and summing over all t_{end} for each of the combinations considered quantifies the robustness with respect to parameter choices, yielding a simplified overall

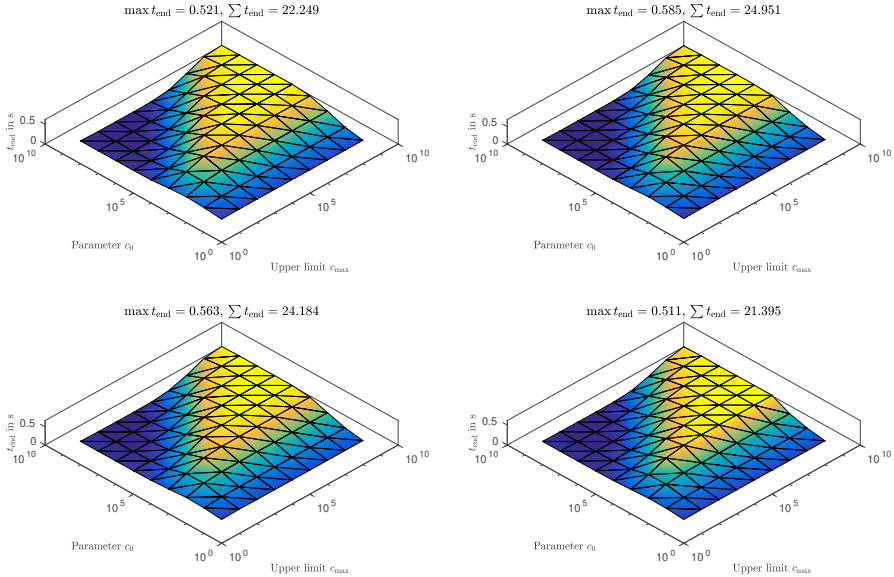


Fig. 8.26: Time $t \leq 1$ reached with $Q_J > 0$ using (from top left to bottom right) no Jacobian stiffening, J_f^{-1} , $J_f^{-1} + J_f$ or J_f^{-2} : harmonic extension in $\hat{\Omega}_f$.

score denoted as $\sum t_{\text{end}}$ in the following. Naturally, some parameter combinations do not result in $t_{\text{end}} = 1$, such that also the maximum time reached with any method in the parameter sweep, $\max t_{\text{end}}$, is also relevant.

Constructing the ALE extension from $\hat{\Omega}_f$, both harmonic extension and pseudo linear elasticity clearly profit from $c_0 \geq 10$ and $c_{\text{max}} \geq 10$, see Figs. 8.26 and 8.27. Here, high c_0 combined with high c_{max} , that is, a small number of elements with coefficients close to c_{max} and the remaining elements assigned c_{min} , lead to bad conditioning of the linear systems due to high gradients in the coefficients. In addition to that, scaling parameters on the upper end of the spectrum also lead to increasingly bad conditioning of the system matrices, since no further scaling is performed in the present case. Mapping from $\hat{\Omega}_f$, pseudo linear elasticity without Jacobian stiffening performs best, closely followed by the two variants involving first (negative) powers of the Jacobian. Pseudo linear elasticity with J_f^{-2} considered for stiffening yields results comparable to harmonic extension, which are slightly worse. Overall, parameters c_0 and c_{max} have a rather predictable effect on the simulation end time reached, while Jacobian stiffening with J_f^{-2} results in local coefficients with sharp gradients.

In comparison to that, constructing the extension from Ω_f^n profits tremendously from incorporating updated stiffness parameters given element sizes h_e in Ω_f^n as shown in Figs. 8.28 and 8.29. Jacobian stiffening again renders the extension more robust, where J_f and $J_f^{-1} + J_f$ again perform best. The interval from $t = 0$ to $t = 1$ can be completed employing both pseudo linear elastic or harmonic extension, where the latter performs better despite its simplicity. Jacobian stiffening with J_f^{-2} is again not delivering results

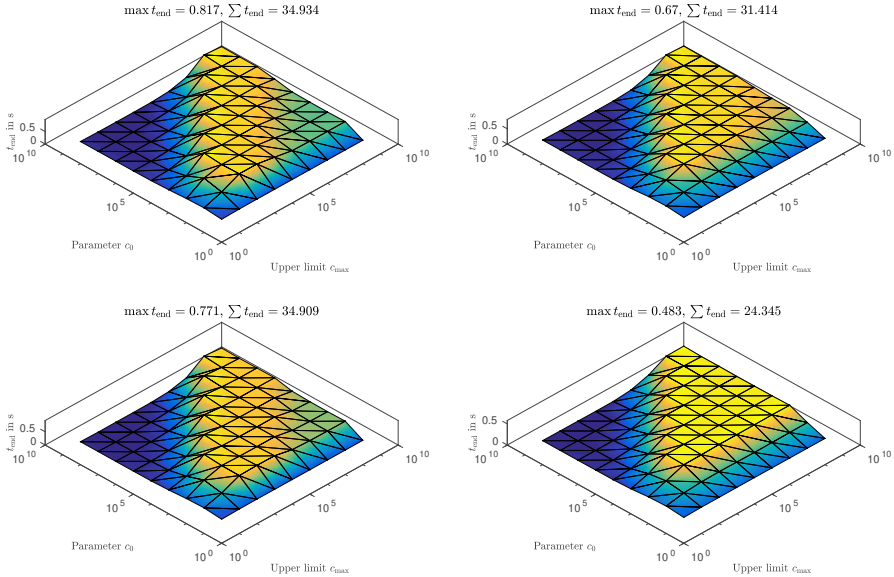


Fig. 8.27: Time $t \leq 1$ reached with $Q_J > 0$ using (from top left to bottom right) no Jacobian stiffening, J_f^{-1} , $J_f^{-1} + J_f$ or J_f^{-2} : pseudo linear elasticity in $\hat{\Omega}_f$.

as good as the variants involving first powers of J_f , which motivates neglecting this variant from now on to ease comparison. Again, combinations of c_0 and c_{max} too high lead to ill-conditioned linear systems, which require increased effort solving.

It is however, rather surprising that an appropriate scaling function c_h (7.80) combined with Jacobian stiffening c_J (7.82) is capable of extending the time interval from almost immediate abortion ($t_{\text{end}} \approx 0.06$ due to the thin boundary layer) obtained via harmonic extension in $\hat{\Omega}_f$ with a uniform diffusion coefficient and no Jacobian stiffening all the way to $t_{\text{end}} = 1$ when parameters are suitably chosen. This test is rather challenging, since we consider a boundary layer mesh, i.e., larger and abrupt changes in element sizes, quadrilateral elements and a rather unforgiving element quality metric suitable for higher-order finite elements.

To highlight parameter combinations c_0 and c_{max} yielding a better quality, Figs. 8.30 and 8.31 depict simulation times reached until a stricter quality threshold of $Q_J > 0.1$ is violated. From these graphs, the trend observed in the previous studies continues, i.e., $c_0 \geq 10$ and $c_{\text{max}} \geq 10$ can improve the obtained results, if suitably chosen. Here, “suitably chosen” is to be understood as the parameters being problem-dependent, but in the present numerical tests, $10^2 \leq c_0 < c_{\text{max}} \leq 10^4$ and combinations in the vicinity of this region performed reasonably well for all possible combinations tested.

The next series of numerical tests in this two-dimensional setup is aimed at the proper selection of $\nu_m \in (-1, 0.5)$ in pseudo elasticity, as different parameters are employed in literature, also including the auxetic range $-1 < \nu_m < 0$ (see, e.g., [85, 473]). So, again employing the stricter quality criterion $Q_J > 0.1$ and varying Poisson’s ratio ν_m ,

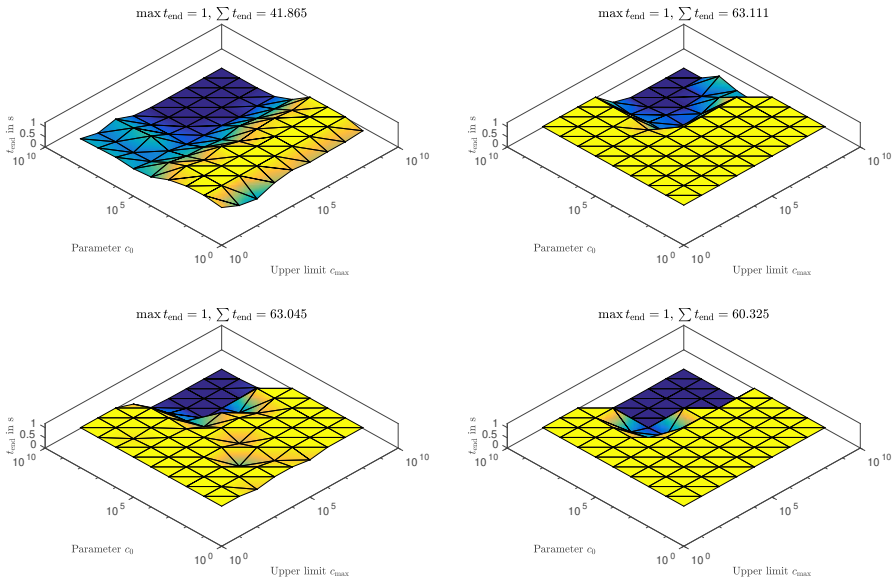


Fig. 8.28: Time $t \leq 1$ reached with $Q_J > 0$ using (from top left to bottom right) no Jacobian stiffening, J_f^{-1} , $J_f^{-1} + J_f$ or J_f^{-2} : harmonic extension in Ω_f^n .

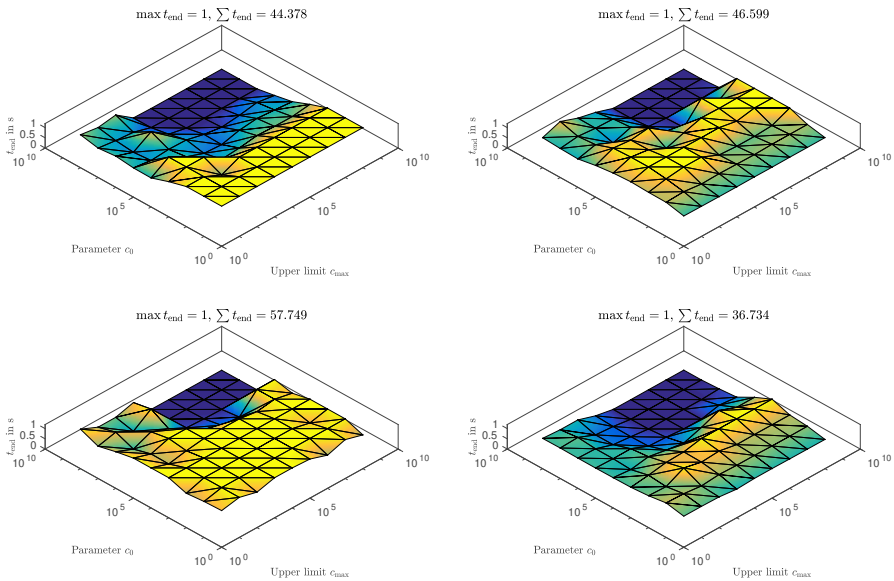


Fig. 8.29: Time $t \leq 1$ reached with $Q_J > 0$ using (from top left to bottom right) no Jacobian stiffening, J_f^{-1} , $J_f^{-1} + J_f$ or J_f^{-2} : pseudo linear elasticity in Ω_f^n .

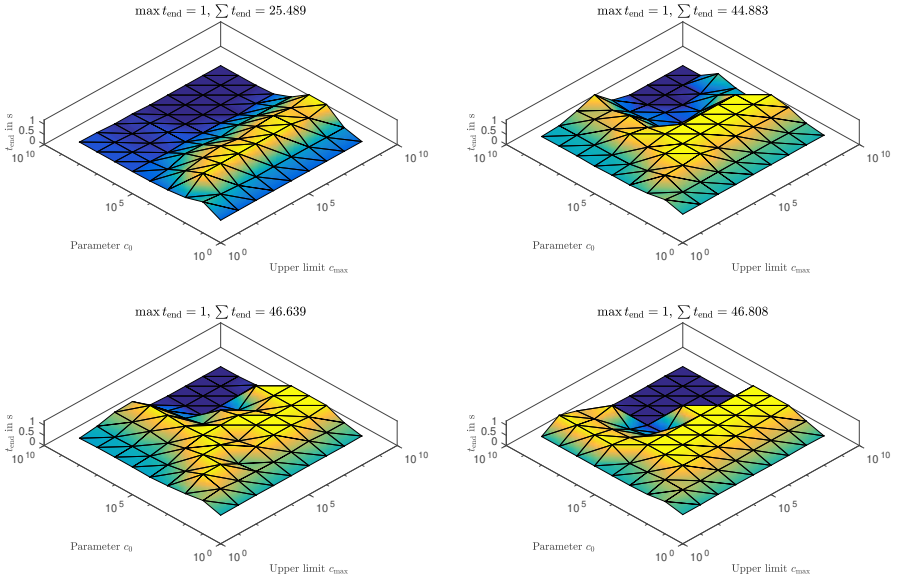


Fig. 8.30: Time $t \leq 1$ reached with $Q_J > 0.1$ using (from top left to bottom right) no Jacobian stiffening, J_f^{-1} , $J_f^{-1} + J_f$ or J_f^{-2} : harmonic extension in Ω_f^n .

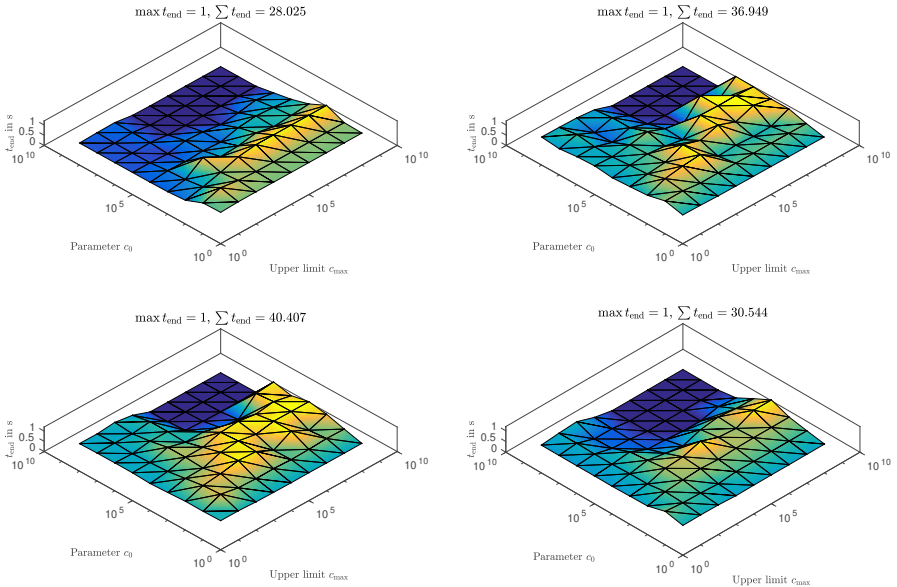


Fig. 8.31: Time $t \leq 1$ reached with $Q_J > 0.1$ using (from top left to bottom right) no Jacobian stiffening, J_f^{-1} , $J_f^{-1} + J_f$ or J_f^{-2} : pseudo linear elasticity in Ω_f^n .

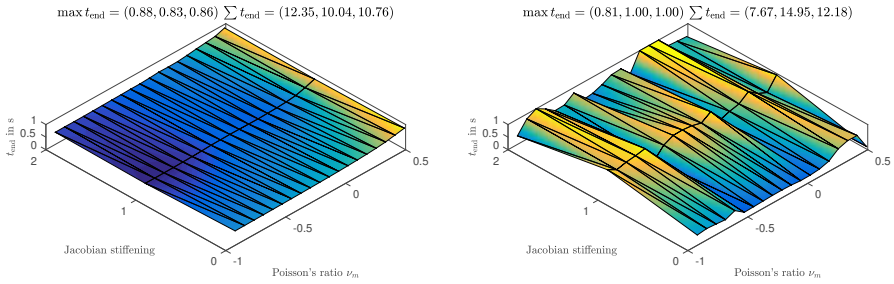


Fig. 8.32: Time $t \leq 1$ reached with $Q_J > 0.1$ using pseudo linear elasticity in $\hat{\Omega}_f$ (left) or Ω_f^ν (right) with no Jacobian stiffening (“0”), J_f^{-1} (“1”), $J_f^{-1} + J_f$ (“2”), $c_0 = 10^3$, $c_{\max} = 10^4$ and various $\nu_m \in (-1, 0.5)$.

a pseudo linear elasticity problem is either formulated on $\hat{\Omega}_f$ or Ω_f^ν to update the fluid domain. Fig. 8.32 contains graphs very much similar to previously presented ones, but now $c_0 = 10^3$ and $c_{\max} = 10^4$ are fixed, while ν_m and the type of Jacobian stiffening are varied. Here, the possible variants are referred to as “0” for no Jacobian stiffening, “1” for J_f^{-1} and “2” for $J_f^{-1} + J_f$, respectively. As can be seen from Fig. 8.32, mapping from $\hat{\Omega}_f$ profits from $\nu_m \rightarrow 0.5$, approaching the incompressible limit, and different types of Jacobian stiffening actually worsen results, where $J_f^{-1} + J_f$ (“2”) is better than J_f^{-1} (“1”) in this test.

Constructing \mathbf{d}_f^{n+1} from Ω_f^ν given \mathbf{d}_s^* , on the other side, Jacobian stiffening has a quite heterogeneous effect, almost doubling $\sum t_{\text{end}}$ with J_f^{-1} (“1”) and increasing reached t_{end} also for the other variant $J_f^{-1} + J_f$ (“2”). With Jacobian stiffening, the target time $t = 1$ is reached with various ν_m , where a clear trend as when mapping from $\hat{\Omega}_f$ cannot be observed. Employing J_f^{-1} , good results are achieved with $\nu_m \in (-1, -0.5]$, alternatively $\nu_m \in [-0.5, 0]$ or $\nu_m \in [0.2, 0.4]$, but a definite recommendation cannot be given. Using Jacobian stiffening of the form $J_f^{-1} + J_f$ (“2”), $\nu_m \in (-1, -0.5]$ or $\nu_m \in [0.2, 0.4]$ might be good choices, but again, such recommendations cannot be generalised based on these numerical tests.

In an attempt to improve these rather underwhelming results, a spatially variable Poisson’s ratio in pseudo linear elasticity is considered as introduced by Jendoubi et al. [467], computing an element-local $\nu_{m,e}(h_e)$ dependent on h_e as given by Eqn. (7.93),

$$\nu_{m,e}(h_e) := \nu_0 - \nu_1 \frac{c_{\min}}{c_h(h_e)},$$

such that $h_e = h_{\max}$ gives $\nu_{m,e} = \nu_0 - \nu_1$, while for the smallest element, we have $\nu_{m,e} \approx \nu_0$ (for $c_{\max} \gg c_{\min}$). With this, the element size is linked to the Poisson’s ratio, targeting $\nu \approx 0.5$ in small elements to preserve $J_f \approx 1$, while still allowing for volume changes in grid areas with larger elements.

Now, we employ the stricter quality criterion $Q_J > 0.1$ again and keep $c_0 = 10^3$ and $c_{\max} = 10^4$ fixed. Since Fig. 8.32 indicates that when mapping from $\hat{\Omega}_f$, the Jacobian stiffening has little effect, we skip it, while when mapping from Ω_f^ν , we choose $J_f^{-1} + J_f$




Fig. 8.33: Time $t \leq 1$ reached with $Q_J > 0.1$ using pseudo linear elasticity in $\hat{\Omega}_f$ with no Jacobian stiffening (left) or Ω_f^n with $J_f^{-1} + J_f$ (right), $c_0 = 10^3$, $c_{\max} = 10^4$ and variable $\nu_{m,e} = \nu_0 - \nu_1 c_{\min}/c_h(h_e)$.

and vary $\nu_0 \in (0, 1/2)$ and $\nu_1 \in (1/20, 1/2)$. The results depicted in Fig. 8.33 indicate that (i) when mapping from $\hat{\Omega}_f$, no additional gains are seen when considering a Poisson ratio dependent on the element size, and (ii) when mapping from Ω_f^n , improvements result for some parameter combinations, with $\nu_0 \approx 0.2$ being a good choice for the present grid, which together with $\nu_1 = 0.1$ or $\nu_1 = 0.5$ yields the best results here. Note that interestingly, the second combination leads to elements with $\nu < 0$ lying in the auxetic range. Results for J_f^{-1} are similar to the second case, hence do not yield additional insights and are thus omitted. From these results, we conclude (limited to this numerical test) that adapting the Poisson's ratio in the proposed form does not yield reliable improvements in addition to the stiffness scaling based on the element size but might still improve results in different scenarios.

Regarding investigations on the time-dependent extension types, we first fix $c_0 = 10^3$ and $c_{\max} = 10^4$ and compute extensions of the interface displacement with pseudo time-dependent approaches involving the first or second material time derivative and compare results to the standard stationary counterpart. Further, we take some of the most promising combinations according to the previous tests, that is, harmonic extension from Ω_f^n and $J_f + J_f^{-1}$, pseudo linear elasticity with $\nu_0 = 0.49$ and $\nu_1 = 0.1$ from either $\hat{\Omega}_f$ with no Jacobian stiffening or from Ω_f^n with $J_f + J_f^{-1}$ and finally the pseudo hyperelastic extension from $\hat{\Omega}_f$ with $J_f + J_f^{-1}$ and identical pseudo Poisson ratio. In any time-dependent scenario, a total of 3 time steps are executed when performing the mesh update, since according to Jendoubi et al. [467], mesh quality is likely further improved in this case. Taking a quality criterion of $Q_J > 0.1$ and varying the pseudo time step size relative to the FSI time step of 1 ms, the maximum reached t_{end} is increased for certain combinations compared to the stationary case as can be seen from Fig. 8.34. For the pseudo hyperelastic extension, a gain in mesh quality is obvious for both the first and second-order material time derivative, while a further increase of the quality criterion to $Q_J > 0.6$ reveals improvements also for the harmonic extension and pseudo linear elasticity as shown in Fig. 8.35.

For the nonlinear problem of pseudo hyperelasticity, taking three time steps in the time-dependent variants is enough to ensure convergence of the Newton solver, when exiting the nonlinear solver after a single Newton update similar to Shamanskiy and Simeon [471]

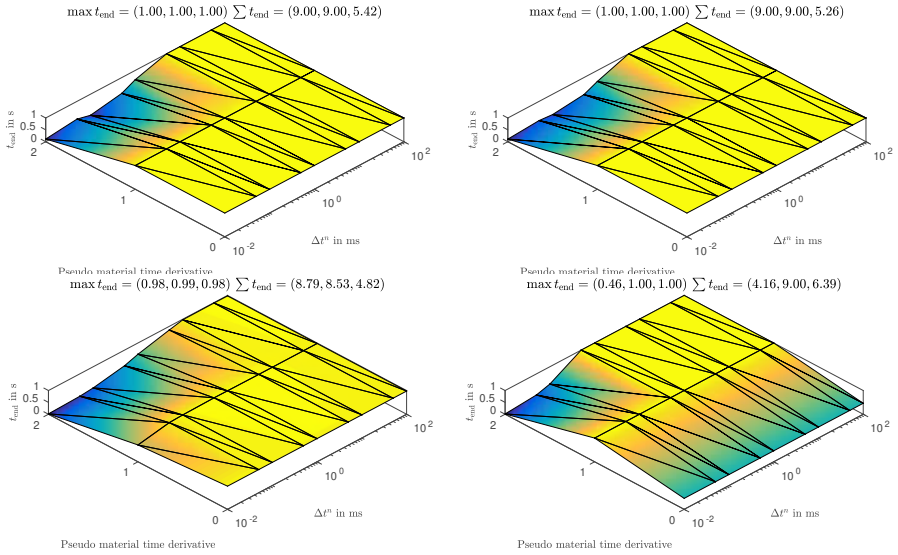


Fig. 8.34: Time $t \leq 1$ reached with $Q_J > 0.1$ using harmonic extension ($\Omega_f^n, J_f + J_f^{-1}$: top left), pseudo linear elasticity ($\hat{\Omega}_f$: top right; $\Omega_f^n, J_f + J_f^{-1}$: bottom left) and pseudo hyperelasticity ($\hat{\Omega}_f, J_f + J_f^{-1}$: bottom right).

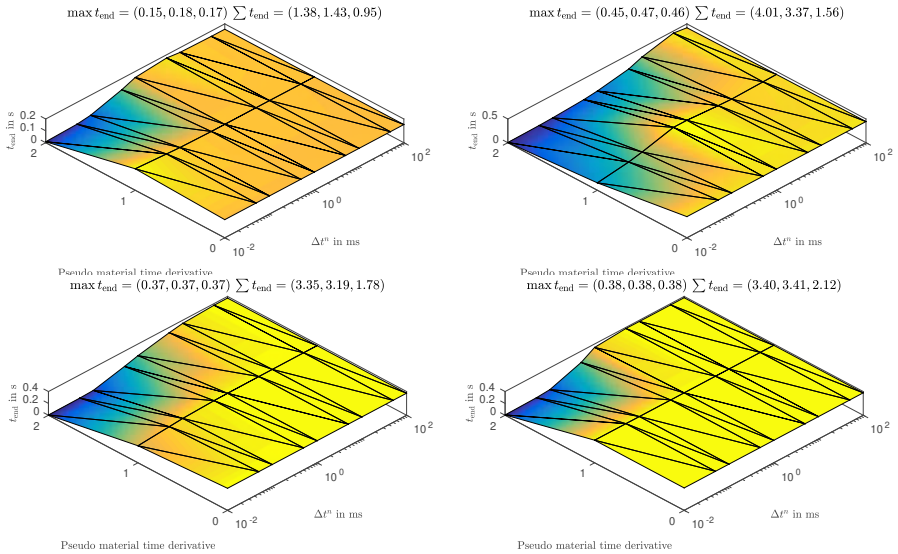


Fig. 8.35: Time $t \leq 1$ reached with $Q_J > 0.6$ using harmonic extension ($\Omega_f^n, J_f + J_f^{-1}$: top left), pseudo linear elasticity ($\hat{\Omega}_f$: top right; $\Omega_f^n, J_f + J_f^{-1}$: bottom left) and pseudo hyperelasticity ($\hat{\Omega}_f, J_f + J_f^{-1}$: bottom right).




Fig. 8.36: Time $t \leq 1$ reached with $Q_J > 0.1$ using pseudo hyperelasticity ($\hat{\Omega}_f, J_f + J_f^{-1}$) with three pseudo time steps and a full Newton solve reducing the residual by a factor of 10^3 (left) or only a single Newton step (right).

and still yields satisfactory results comparing to a setup driving the relative residual down by a factor of 10^3 in each Newton solver call. Even for $Q_J > 0.1$, the mesh update fails after a few steps due to non-convergence of the Newton solver despite the deformation being applied incrementally over 1000 steps. Fig. 8.36 highlights the difficulties arising when the nonlinear solver is not converged properly, whereas interestingly, the mesh quality and Newton convergence are ensured taking three pseudo time steps per mesh update step in the pseudo time-dependent variants.

To shed some more light on the effects of certain parameter combinations for the pseudo time-dependent extension types, we conduct a final study using pseudo hyperelasticity with $c_0 = 10^3$, $c_{\max} = 10^4$ and varying the pseudo time step size in the variants involving the first or second pseudo material time derivative. In Fig. 8.37 the reached end time t_{end} such that $Q_J > 0.6$ is shown, where we differentiate results obtained with no Jacobian stiffening (–) or $J_f + J_f^{-1}$ (–), constant $\nu_m \approx \nu_0$ (blue) or $\nu_1 = 0.1$ (orange) and performing three pseudo time steps combined with multiple (>) or only one (+) Newton step(s) or just a single pseudo time step with multiple Newton steps (×).

Here, we can see that $J_f + J_f^{-1}$ is always slightly better than no Jacobian stiffening, since all orange lines are always over the corresponding blue lines. A variable Poisson ratio with $\nu_1 = 0.1$ does not lead to reliable *and* notable improvements, especially in the well performing combinations—the dashed (–) lines are mostly above or approximately equal to their solid counterparts (–). The most important difference between first and second-order pseudo material time derivative is that with the former, maximal achieved end times t_{end} increase. This effect is observed when multiple time steps are performed, but the number of Newton steps does not influence the quality of the mapping substantially. Incorporating the second pseudo time derivative, similar slight improvements are seen, before pseudo time step size and reached end time decrease simultaneously. Connecting back to the discussion on initial conditions for the pseudo time-dependent problem with $d_t d_f$ in Sec. 7.5, we note here that choosing $\dot{d}_f = \ddot{d}_f = \mathbf{0}$ only at the beginning of the simulation at $t = 0$ or at each call of the mesh extension routine does not alter the obtained results significantly.

The pseudo time-dependent extension types may thus yield improvements, but introduce yet another layer of complexity, such that determining appropriate parameters becomes

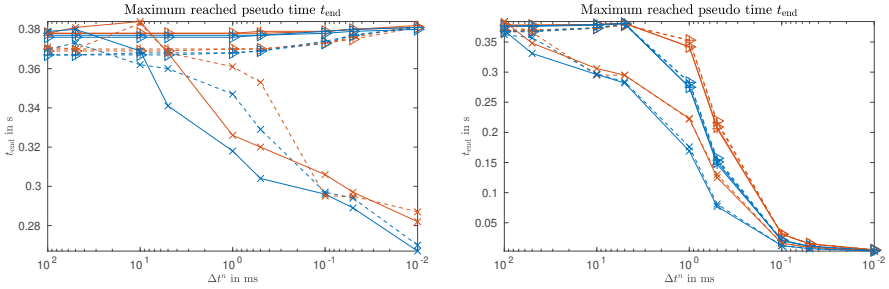


Fig. 8.37: Time $t \leq 1$ reached with $Q_J > 0.6$ using pseudo hyperelasticity ($\hat{\Omega}_f, J_f + J_f^{-1}$) combined with the first (left) or second (right) time derivative. Comparing no Jacobian stiffening (—) or $J_f + J_f^{-1}$ (---), constant $\nu_m \approx \nu_0$ (blue) or $\nu_1 = 0.1$ (orange) and performing three pseudo time steps with multiple (\triangleright) or one (\times) Newton step(s) or a single pseudo time step with multiple Newton steps (\times).

even more convoluted. Incorporating the first time derivative seems a viable alternative to the operators previously considered, since they yield better mesh quality and might be combined with nonlinear solution schemes executing only single Newton steps per pseudo time step. In addition to the favourable properties of the resulting ALE map, including a scaled mass matrix further improves conditioning of the linear system and thus also bears potential for improving linear solver convergence. Comparing the numerical effort and gains in mesh quality, linear and time-dependent variants doing only single pseudo time steps are hence a particularly attractive option.

Let us close the investigations on artificial (initial-)boundary value problems to construct the ALE extensions by some concluding remarks summarising our findings:

1. All considered extension types profit from scaling coefficients based on the element size, if suitable (problem dependent) parameters are found. At the same time, coefficients selected too aggressively hamper convergence of (non-)linear solvers.
2. Jacobian-based stiffening through J_f , $J_f + J_f^{-1}$ or J_f^{-2} might further robustify the mesh update step, but can also lead to severe ill-conditioning for similar reasons. Out of the three options tested, variants involving only first powers of J_f yield the best results.
3. In the considered example, extensions based on Ω_f^n were found more robust than their fully Lagrangian counterparts based on $\hat{\Omega}_f$. The former incorporate updated stiffness parameters based on the current mesh, while the latter are less prone to accumulate mesh distortions in simulations involving periodic solutions (not shown within this work, see, e.g. [471] for tests in this regard).
4. In pseudo elasticity and similar extension types, variable Poisson's ratio did have a beneficial effect in some cases, but did not lead to improvements consistently. Penalising $J_f \neq 1$ via nearly incompressible material laws in pseudo hyperelasticity when mapping from the Lagrangian configuration $\hat{\Omega}_f$ conserves mesh quality well, but the (non-)linear solvers needed increase the effort in constructing the extension.

5. Time dependent extensions, namely parabolic extension [467], instationary linear- or hyperelasticity [470] increase the numerical effort due to multiple pseudo time steps being performed per mesh update, but linear systems might be better conditioned (depending on the pseudo time step size) and favourable properties of time-dependent problems may render the extension more robust.
6. Out of the variants considered, good choices are (i) harmonic extension from Ω_f^n with Jacobian stiffening for its simplicity and surprising performance, (ii) pseudo linear elasticity from Ω_f^n with $\nu_0 = 0.20$ and $\nu_1 = 0.1$ or $\nu_1 = 0.5$, (iii) pseudo linear elasticity from $\hat{\Omega}_f$ for it being linear and independent of past time step's mesh deformation fields, (iv) pseudo hyperelasticity with $\nu_0 = 0.49$ and $\nu_1 = 0.1$ preserving $J_f \approx 1$ long, but suffering fast quality drop-off under extensive deformations and lastly all of (i)–(iv) combined with the first pseudo time derivative, performing a single or multiple pseudo time steps per mesh update.

With this, some first numerical examples and investigations on the PPE-based semi-implicit split-step scheme's behaviour are completed and we proceed with larger problems, tackling more challenging scenarios in the cardiovascular context and aeroelasticity involving all the modelling aspects covered within this work.

The numerical experiments summarised in the following sections are selected examples from the author's works in [5, 6, 10–12, 512] and document the development of the split-step framework over the past years. Naturally, certain algorithms covered in this thesis were added over time or further developed, such that various submodules of the software might in fact perform better than reported. Since the focus herein does not lie on performance in the sense of throughput, i.e., DoFs/s core, but rather on coupling iterations and maximum admissible time step while preserving temporal stability and robustness, this fact is ignored, noting that the most recent implementation is employed for the patient-specific case of aortic dissection presented in Sec. 8.7.

That being said, before diving deeper into the detailed test case descriptions and interpretation of the results we also want to mention that most of the presented results still inherit a certain academic character, meaning that they reflect patient-specific or clinical cases to a certain extent, but were not derived from complete sets of medical data. However, all presented results are in the physiological range, prescribing realistic flow rates and pressure levels and employing realistic material parameters, and hence, patient-specific results can be achieved, given that real medical data is available and suitably processed.

As already mentioned at various places in this work, the implementation is based on the finite element framework `deal.II` [96–99] interfacing Trilinos' ML [111] package to construct preconditioners via a variety of algebraic multigrid methods. Some of the finite element meshes are constructed via `Gmsh` [511], in some cases taking interface representations as input modified manually in `Meshmixer` [513].

8.5 Abdominal aortic aneurysm

We present the first results in a practical application of the semi-implicit PPE-based FSI coupling scheme to an idealised abdominal aortic aneurysm (AAA) geometry as considered in [5]. In silico studies on AAAs are conducted to investigate flow characteristics, quantify rupture risk or compare different treatment options in clinical support [514–516]. Given this high relevance and great value for medicine, various computational models of varying complexity ranging from benchmark-character studies all the way to patient-specific simulations have been presented (see, e.g., [353, 370, 422, 517–521]). Herein, a physiologically inspired setup is designed similar to [517], based on a prototypical geometry with dimensions as recommended by surgeons [522, 523] with volumetric flow rates taken from [524].

The AAA geometry is constructed from a straight cylinder undergoing a series of mappings to retrieve the geometry as shown in Fig. 8.38. The vessel has a length of $l = 20$ cm and a lumen radius of $r_i = 1$ cm at the inlet and outlet sections, while between, the aneurysm expands up to ≈ 6.5 cm in lateral direction and ≈ 5.5 cm in anterior-posterior direction. The finite element mesh consists of $\approx 1.3 \times 10^5$ hexahedra, which results in $\approx 1.0 \times 10^5$ nodes.

The tissue is subdivided into medial and adventitial layers of the aorta, $\hat{\Omega}_{s,1}$ and $\hat{\Omega}_{s,2}$, which have a uniform thickness of $h_s = 0.75$ mm. For this simplistic geometry with only a single vessel, the standard algorithm for fibre orientation using two Laplace solves is applicable and yields satisfactory results as shown in Fig. 8.38. The mean fibre direction can thus be computed easily depending on the parameter α_c , which inclines mean fibre directions symmetrically from the circumferential direction vector, see Sec. 6.1. Despite the small problem size, the computational domain is distributed to 8 processors as shown in Fig. 8.39(a).

Regarding boundary conditions, we fix the inlet and outlet faces and incorporate viscoelastic external tissue support with $k_e = 10^7$ N/m³, $c_e = 10^5$ Ns/m³ and $p_e = 0$ Pa chosen in a similar range as [266, 359, 360, 363]. A target volumetric flow rate is enforced at the inlet via a periodic mean inlet velocity \bar{u}_{in} depicted in Fig. 8.39(b), setting the normal inlet velocity component u_1 in terms of the vessel radius r as

$$u_1 = 2\bar{u}_{\text{in}} \left(1 - \frac{r^2}{r_i^2} \right) \eta_t(t), \quad \text{with} \quad \eta_t(t) = \begin{cases} \sin^2 \left(\frac{\pi t}{2T_r} \right) & \text{for } t \leq T_r = 0.2 \text{ s,} \\ 1 & \text{otherwise,} \end{cases}$$

to match the flow rate computed from \bar{u}_{in} with a parabolic inflow profile. At the outlet, we employ backflow stabilisation and a Windkessel model (see Sec. 4.6) with a capacitance of $C = 1.25 \times 10^{-9}$ m³/Pa and proximal and distal resistances of $R_p = 266.66 \times 10^5$ Pa s/m³ and $R_d = 6.8 \times 10^8$ Pa s/m³, respectively. In addition to that, GLS stabilisation is considered to stabilise convective terms (see Sec. 4.3).

Concerning material properties, we consider physiological parameters as introduced in Sec. 8.3 for the two-layered straight vessel; that is, fluid and structural densities of

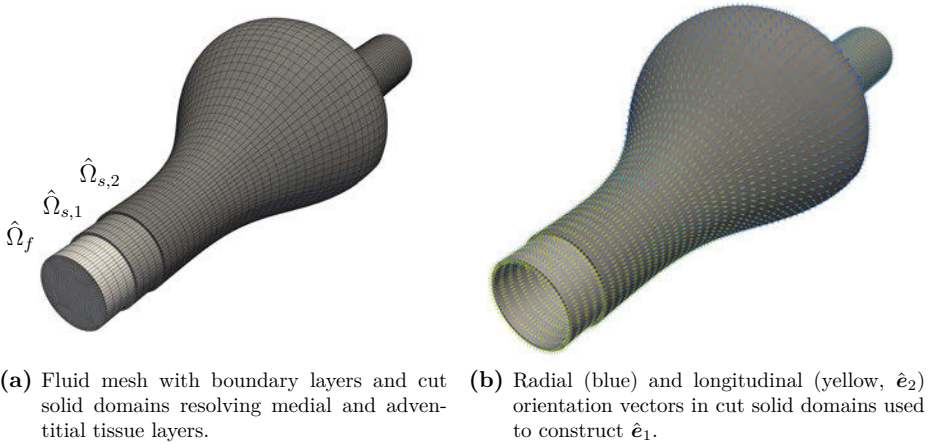


Fig. 8.38: Blood flow in an idealised AAA: lumen mesh and cut tissue layers (a) and radial (blue) and longitudinal (yellow, \hat{e}_2) orientation vectors in cut tissue layers (b).

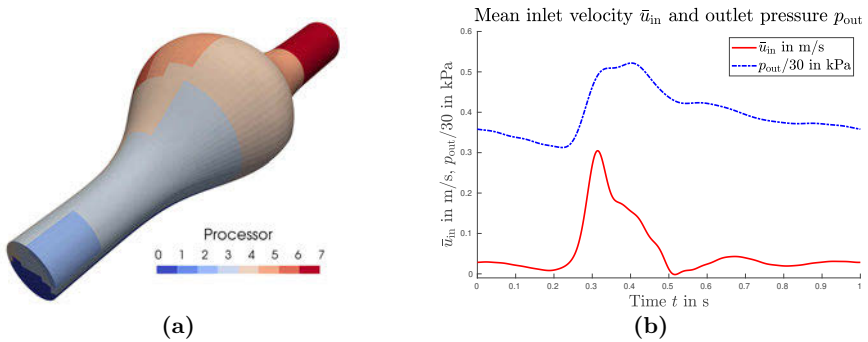


Fig. 8.39: Blood flow in an idealised AAA: mesh distributed to 8 processors (a) and flow data (b) [517, 522, 523].

$\rho_f = 1060 \text{ kg/m}^3$ and $\rho_s = 1200 \text{ kg/m}^3$, respectively, either a Newtonian fluid with viscosity $\mu_f = 3.5 \text{ mPa s}$ or a Carreau fluid with $\eta_0 = 56 \text{ mPa s}$, $\eta_\infty = 3.45 \text{ mPa s}$, $\lambda_f = 3.313 \text{ s}$ and $b = 0.3568$ [508], and the tissue is modeled as a St. Venant–Kirchhoff or linear elastic solid ($E_s = 300 \text{ kPa}$, $\nu_s = 0.3$) or as a layered neo-Hookean solid with $\nu_s = 0.499$ and shear rates $\mu_{s,1} = 62.1 \text{ kPa}$ and $\mu_{s,2} = 21.6 \text{ kPa}$, respectively. Optionally, fibre reinforcement is taken into account applying the HGO model with $k_1 = 1.4 \text{ kPa}$, $k_2 = 22.1$ and $\kappa_{c,1} = 0.12$, $\alpha_{c,1} = 27.47^\circ$ and $\kappa_{c,2} = 0.25$, $\alpha_{c,2} = 52.88^\circ$ [338, 509]. As this numerical example is taken from [5], prestress is not accounted for at this point. This is a quite drastic simplification, since the stress state in the tissue greatly depends on considering the prestress \mathbf{S}_0 present at the time of image acquisition. Nonetheless, we neglect this prestress and interpret the geometry as shown in Fig. 8.38 as a “deflated” geometry, resulting from a prestress algorithm recovering an appropriate initial zero stress state yielding the current domain under diastolic flow conditions [355, 372, 525].

The time interval considered ranges from $t = 0$ to $t = 3 \text{ s}$, such that a total of three cardiac cycles are computed using a uniform time step $\Delta t = 1 \text{ ms}$ in the fully second-order accurate scheme, i.e., BDF2 and linear extrapolation for all involved fields, including the fluid velocity ($m_u = 2$). The CH- α timestepping scheme with $\rho_\infty = 0$ is used to suppress pressure fluctuations stemming from the Windkessel model (for a discussion on this topic, see, e.g., the excellent work by Grinberg and Karniadakis [272]). Aitken’s acceleration is adopted with an initial $\omega_0 = 0.001$, coupling the PPE and solid momentum subproblems in a semi-implicit fashion with a Dirichlet–Neumann scheme (SIDN) to reach $\epsilon_{\text{abs}}^d = \epsilon_{\text{abs}}^p = 10^{-6}$ or $\epsilon_{\text{rel}}^d = \epsilon_{\text{rel}}^p = 5 \times 10^{-4}$ in Eqns. (7.63) and (7.65). If a nonlinear constitutive law for the structure is selected, Newton’s method is executed using absolute and relative convergence criteria $\epsilon_{\text{abs}}^{\text{nl}} = 10^{-10}$ and $\epsilon_{\text{rel}}^{\text{nl}} = 10^{-4}$ in (5.70).

Now, inspecting the solution’s final third cycle at late diastole ($t = 2.2 \text{ s}$), mid systole at $t = 2.4 \text{ s}$ and early diastole ($t = 2.6 \text{ s}$), one observes strong recirculatory flow in the AAA, triggered by the rapid reduction of the inflow rate. Fig 8.40 shows selected streamlines with $\|\mathbf{u}_f\|$ indicating areas of large velocity gradients and resulting large shear rates, ultimately leading to strong gradients in the viscosity field. As a consequence, the viscosity field spans the whole admissible range at almost any point in time. The fluid pressure is rather uniform in space, but is subject to large variations in time due to the Windkessel model. This pressure difference dominates the tissue’s deformation rather than the velocity acting on the vessel. Depicting the displacement and pressure at the reference apex point at $\hat{\mathbf{x}} = (0.1, 0.039, 0)^T$ in Fig. 8.41, it is obvious that a periodic state is not yet reached within three cardiac cycles using these settings. To a large extent, this is caused by the viscoelastic tissue support and the Windkessel model with $C = 1.25 \times 10^{-9} \text{ m}^3/\text{Pa}$ and $p_d = p_c^0 = 0 \text{ Pa}$, only gradually increasing the pressure level in the AAA from zero initial pressure. The deliberately chosen low proximal resistance value of $R_p = 266.66 \times 10^5 \text{ Pa s/m}^3$ yields rather small fluctuations in the pressure amplitude over a pulse cycle and is used here to showcase a scenario with a fully second-order accurate scheme. Choosing larger R_p (to give physiological pressure fluctuations), the scheme becomes increasingly unstable as elaborated in Sec. 8.3, which is for the most part caused by the Windkessel model amplifying pressure fluctuations [272].

This numerical test from [5], shows, that the fully second-order accurate semi-implicit Dirichlet–Neumann (SIDN) scheme can indeed be used in the cardiovascular context.

Doing so, temporal stability is limited, such that (i) adaptive timestepping is found unstable for reasonable target CFL numbers and (ii) Windkessel models with increased capacitance and lowered resistance have to be used to recover a stable scheme. Both of these restrictions are lifted in the following examples in Secs. 8.6 and 8.7, combining Robin interface conditions with IQN-ILS acceleration and considering a first-order accurate linearisation $\mathbf{u}_f^* \approx \mathbf{u}_f^n$ ($m_u = 1$). Nonetheless, we stick to the current choice for the sake of presentation, highlighting the improvements achievable when incorporating such techniques in the following examples.

Judging from the displacement and pressure in the reference point depicted in Fig. 8.41, we may argue that applying more advanced constitutive models does not alter the solution much due to small strains present. A difference in the system's response is only seen between nearly incompressible and compressible variants, which is also in parts caused by the different tissue parameters. Regarding differences between Newtonian and Carreau models, Fig. 8.41 does not yield any insights, since relevant quantities with respect to haemodynamics are, e.g., the time-averaged shear rate and time-averaged shear stress. To investigate effects on these quantities, we further define the time average $\bar{f}(\mathbf{x})$ of $f(\mathbf{x}, t)$ over a period T_p by

$$\bar{f} = 1/T_p \int_{i T_p}^{(i+1) T_p} f(\mathcal{A}_t(\hat{\mathbf{x}}, t), t) dt, \quad (8.4)$$

to be considered in the third cycle, i.e., $i = T_p = 1$, directly yielding the time-averaged shear rate $\bar{\dot{\gamma}}$ from $\dot{\gamma} := \sqrt{1/2 \nabla^s \mathbf{u} : \nabla^s \mathbf{u}}$ introduced in Eqn. (2.7). Similarly treating the shear stress $\boldsymbol{\tau}$ on Σ^t , defined as [17]

$$\boldsymbol{\tau} := \boldsymbol{\sigma}_f \mathbf{n}_f - [(\boldsymbol{\sigma}_f \mathbf{n}_f) \cdot \mathbf{n}_f] \mathbf{n}_f, \quad (8.5)$$

results in a time-averaged wall shear stress $\bar{\boldsymbol{\tau}}$. These haemodynamic indicators are computed for the final period, and as can be seen from Fig. 8.42, a striking difference is observed, which is a first hint towards the importance of rheological modelling in haemodynamics. Herein, however, focus lies on the coupling scheme rather than phenomenological investigations, only showcasing these results to underline the versatility and applicability of the framework.

Now, let us focus on the SIDN coupling scheme's performance and the subproblem solvers. Throughout the entire simulation, less than 30 coupling steps are needed to drive the residual down to $\epsilon_{\text{abs}}^d = \epsilon_{\text{abs}}^p = 10^{-6}$ or $\epsilon_{\text{rel}}^d = \epsilon_{\text{rel}}^p = 5 \times 10^{-4}$ with standard Aitken's relaxation as shown in Fig. 8.43. A dependence on the flow field and tissue deformation is observed. Interestingly, the accumulated FSI iteration counts in Fig. 8.43 show the less stiff NH and HGO models resulting in fewer FSI iterations, whereas the stiffer E and SVK models lead to slightly higher pressure levels and thereby to increased pressure fluctuations given the Windkessel outlet. For comparison, Fig. 8.43(b) also contains accumulated iteration counts for a fully implicit Dirichlet–Neumann (IDN) scheme and a

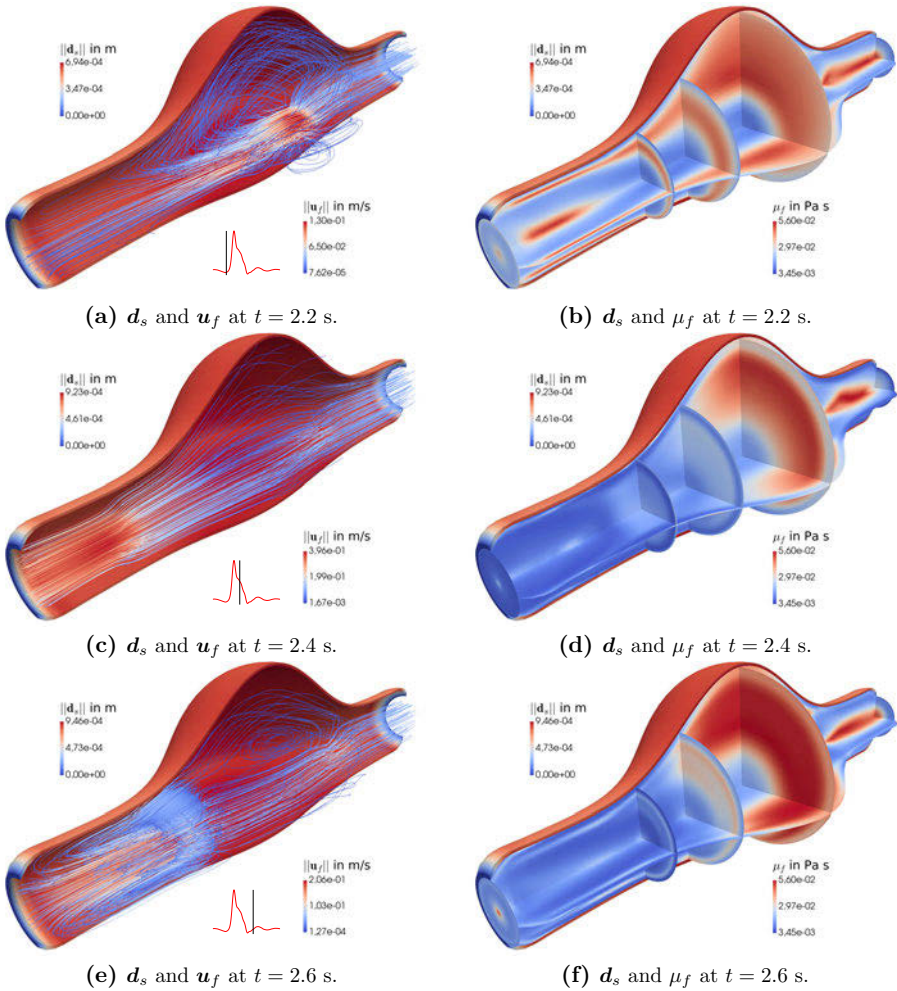


Fig. 8.40: Blood flow in an idealised AAA: solution at $t = 2.2, 2.4, 2.6$ s using a Carreau fluid and HGO solid (deformation scaled by 5). Solid displacement d_s and fluid velocity streamlines (left) and viscosity μ_f in selected slices (right).

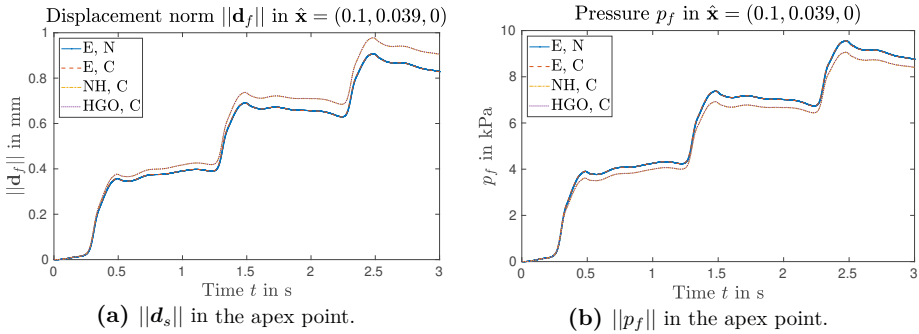


Fig. 8.41: Blood flow in an idealised AAA: displacement and pressure in the reference point using linear elasticity (E), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) solids and Newtonian (N) or Carreau (C) fluids. The periodic state is not yet reached due to the Windkessel model, and nonlinear effects are small due to small strains.

geometry explicit Dirichlet–Neumann (GEDN) approach, treating only the mesh update step explicitly. Here, only a decrease of up to 31% in the total count of FSI coupling steps is observed when using the semi-implicit variant with Aitken’s acceleration.

Conservatively applying an AMG-preconditioned [111] FGMRES solver with Chebyshev smoother leads to ≈ 3 iterations per time step to reduce the residual by 10^8 in the fluid’s momentum balance step given the small time steps as Fig. 8.44 shows. For the mesh motion equation, i.e., simple harmonic extension with Jacobian stiffening and mesh-dependent diffusion coefficient, an AMG-preconditioned CG solver is applied to yield a similar residual reduction within ≈ 15 iterations, while the viscosity projection’s mass matrix is lumped and inverted by vector operations only. The Leray projection step’s CG solver, executed only once per time step, converges within approximately 35 iterations, very much similar to the PPE step’s linear solver given the similarity of the problems.

For the overall computing time, the implicitly treated steps, i.e., the PPE and solid momentum balance steps are particularly relevant. The good performance of the AMG preconditioners by Trilinos’ ML package is crucial here, leading to nicely bounded iteration counts in the structure’s CG or FGMRES solves and the PPE’s CG solve as can be seen from Fig. 8.45.

Slight dependencies on the flow field/pressure/displacement are seen, owing to varying right-hand sides and initial guesses. As for the Newton solver, convergence is reached for all time steps within 3 iterations, again caused by small time steps and high-quality initial guesses. Comparing constitutive models, we see little difference using Carreau or Newtonian fluids, whereas the solid constitutive model influences the iteration count in the linear system solve quite drastically. Complex contributions to the linear systems and near-incompressibility render neo-Hookean and HGO models harder to solve.

In summary, all subproblem’s linear systems are solved within a satisfactory number of iterations using algebraic multigrid methods provided by Trilinos’ ML package [111], since all systems resulting from the partitioned PPE-based split-step scheme are eas-

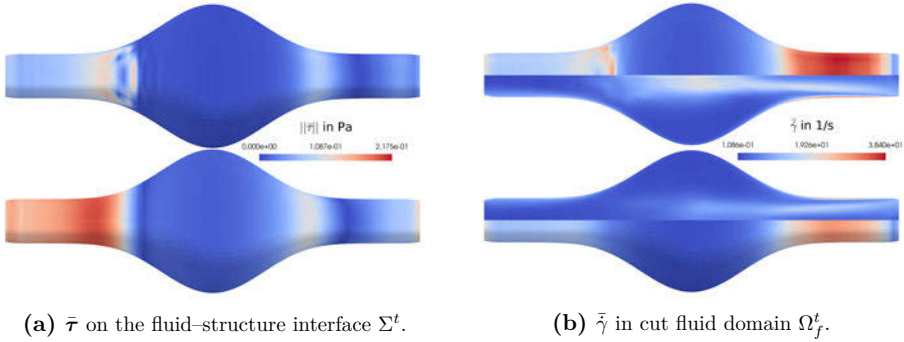


Fig. 8.42: Time-averaged shear stress $\bar{\tau}$ and shear rate $\bar{\gamma}$ in anterior-posterior view, inlet on the right: symmetric solutions show vast differences using Newtonian (top row) and Carreau (bottom row) models.

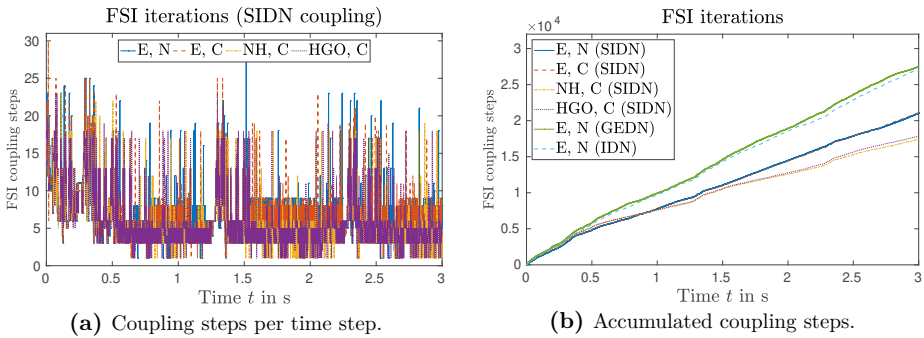


Fig. 8.43: FSI coupling steps per time step (left) and accumulated (right) employing semi-implicit (SIDN), implicit (IDN) or geometry explicit (GEDN) Dirichlet–Neumann schemes and linear elasticity (E), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) material models and Newtonian (N) or Carreau (C) fluids.

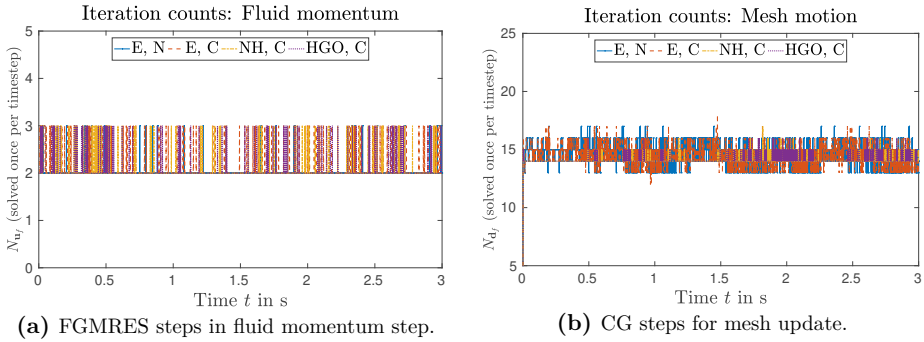


Fig. 8.44: Iterations needed in the fluid momentum balance (left) and mesh motion (right) employing linear elasticity (E), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) solids and Newtonian (N) or Carreau (C) fluids in the SIDN scheme.

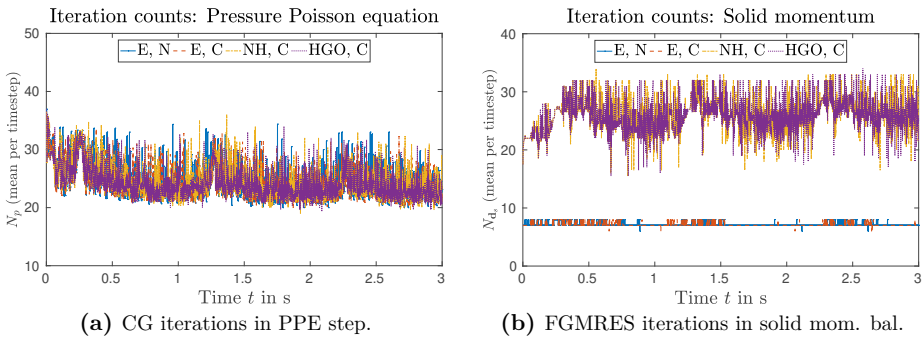


Fig. 8.45: Iteration counts in the PPE (left) and solid momentum balance (right) solvers using linear elasticity (E), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) solids, Newtonian (N) or Carreau (C) fluids in the SIDN scheme.

Tab. 8.5: Absolute (relative) computing times using linear elasticity (E), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) solid models, Newtonian (N) or Carreau (C) fluids and semi-implicit (SIDN), implicit (IDN) or geometry explicit (GEDN) Dirichlet–Neumann coupling schemes.

	computing time	FSI steps	time/FSI step
E, N (SIDN)	5.53×10^5 s (100.0%)	21046 (121.0%)	26.3 s (100.0%)
E, C (SIDN)	5.57×10^5 s (100.1%)	21118 (121.4%)	26.4 s (100.4%)
NH, C (SIDN)	9.22×10^5 s (166.7%)	17395 (100.0%)	53.0 s (201.5%)
HGO, C (SIDN)	1.09×10^6 s (196.4%)	17846 (102.6%)	56.5 s (214.8%)
E, N (GEDN)	1.07×10^6 s (193.5%)	27532 (158.3%)	38.9 s (147.9%)
E, N (IDN)	1.41×10^6 s (255.0%)	27045 (155.5%)	52.3 s (198.9%)

Tab. 8.6: Relative timings in % per time step for computationally relevant steps in the SIDN scheme indicated by the unknown solved for. Contributions split into (assembly+AMG setup+linear solve) for PPE and solid momentum step using linear elasticity (E), neo-Hookean (NH) or Holzapfel–Gasser–Ogden (HGO) material models and Newtonian (N) or Carreau (C) fluids.

	\mathbf{d}_f	\mathbf{u}_f	ψ	ζ	p_f	\mathbf{d}_s
E, N	6.0	8.0	4.8	2.6	(6.4+4.2+19.7)	(20.3+4.9+16.2)
E, C	5.9	8.0	4.9	2.6	(6.3+4.2+19.5)	(20.2+4.9+16.2)
NH, C	2.3	2.2	3.3	0.8	(1.6+2.6+14.5)	(29.5+5.1+35.7)
HGO, C	1.6	1.7	2.2	0.6	(1.3+1.9+10.4)	(47.6+3.7+26.9)

ily categorised into common classes frequently encountered in science and engineering. Lastly, let us give absolute and relative computing times per time step for the different constitutive models in Tab. 8.5. As noted before, stiffer material parameters used for the E and SVK simulation lead to increased pressure levels and an increased number of coupling steps needed as compared to the nearly incompressible hyperelastic models. The viscosity projection step is negligible in terms of numerical effort, and the constitutive models make up for large parts of the computing time per time step due to increased complexity in the element integration using nonlinear models and iteration counts $\approx 2.5\times$ higher than linear elasticity, see, again Fig. 8.45. As expected, the semi-implicit scheme outperforms the GEDN and IDN variants by a factor of ≈ 2 or more. Using the GEDN scheme, the number of coupling steps even increases slightly compared to the fully implicit scheme, but the individual coupling iterations being cheaper to execute, the overall computing time of the IDN scheme is reduced by $\approx 32\%$.

To conclude the analysis of computing times in the semi-implicit scheme, we list timings for the relevant steps in Tab. 8.6. Naturally, the implicitly coupled and hence most often

executed steps in the scheme contribute most to the time spent. This comparison shows, how little effort the explicitly treated steps require, since the mesh update step (\mathbf{d}_f), the fluid's momentum balance (\mathbf{u}_f), Leray projection (ψ) and viscosity projection are only executed once per time step. The lumped viscosity projection step is even omitted from Tab. 8.6, since it amounts for less than 0.5% of the computing time per time step. Ignoring the fact that the pressure boundary projection step (ζ) is only needed once per time step and repeatedly executing it prior to each PPE step, we see that it still does not contribute substantially to the overall computing time, even though the mass matrices for the boundary projection are assembled and solved without any simplifications. The central parts of the SIDN scheme are naturally the PPE (p_f) and structure's momentum balance (\mathbf{d}_s) steps, responsible for the largest share of execution time. Thanks to the great performance of the algebraic multigrid methods, iteration counts are nicely bounded and parallel scalability given large enough problems can be expected. Herein, however, corresponding tests regarding scalability are not considered.

8.6 Iliac bifurcation

In this second practical application of the split-step scheme in the cardiovascular setting, we consider blood flowing through an iliac bifurcation, driven by flow data in the physiological range. This scenario is the first one considering the bulk of techniques introduced in Chs. 3–7. We consider incompressible flow of a Carreau fluid, GLS and backflow stabilisations, multiple Windkessel outlets and a flow-rate-equivalent, mapped inflow profile on the fluid side. On the solid side, we combine a prestressed reference configuration based on diastolic flow conditions, a nearly incompressible, hyperelastic and fibre-reinforced tissue model with proper material orientation, viscoelastic support and fixed inlet and outlets. The semi-implicit Dirichlet–Robin variant with IQN-ILS acceleration is employed, being the most robust and fastest coupling method identified in Sec. 8.3 and further numerical tests not included for brevity.

As a starting point for the volumetric grid generation, an interface representation is reconstructed from computed tomography images, which is provided by [526–528]. This geometry of length ≈ 11 cm has a mean inlet radius of $r_M \approx 7.1$ mm, while the two outlets have mean radii of ≈ 7.8 and 8.6 mm. The fluid–structure interface is uniformly extended by 1 mm in unit normal direction with `Meshmixer` [513], yielding the tissue layer enclosing the lumen. Then, inlet and outlet sections are closed and volumetric mesh generation is carried out with `Gmsh` [511], resulting in the grid depicted in Fig. 8.46. Spatial discretisation leads to $\approx 585 \times 10^3$ tetrahedra with $\approx 108 \times 10^3$ nodes, such that employing the PPE-based split-step scheme yields $\approx 1.19 \times 10^6$ degrees of freedom in total. Here, we use only simplicial finite elements to construct ansatz spaces for the involved unknowns \mathbf{d}_s , \mathbf{u}_f , \mathbf{d}_f , p_f and the projection variable ψ . For later reference, a point on the fluid–structure interface at the bifurcation site is marked in Fig. 8.46, since we compare the solution in this point using various parameter combinations.

Constructing the material orientation in this geometry combining the longitudinal direction obtained from solving a Laplace equation and the extrapolated and conditionally-averaged interface normal as detailed in Sec. 6.1 with the proposed standard parameters




Fig. 8.46: Iliac bifurcation: cut tissue layer (grey) enclosing the fluid domain (red) and reference point (\circ) on the fluid–structure interface (top). Material orientations e_1 (blue) in circumferential direction and e_2 (yellow) in longitudinal direction (bottom).

$h_{l,i} = 100$, $i = 1, 2$, $\alpha_{\text{tol}} = 120^\circ$ and $N_{\text{avg}} = 5$ yields satisfactory material orientation vector fields as depicted in Fig. 8.46.

Material parameters are set in the physiological range, taking $\rho_f = 1060 \text{ kg/m}^3$ and $\rho_s = 1200 \text{ kg/m}^3$ for the fluid and solid densities and a Carreau fluid ($\kappa_f = 1$, $a = 2$) with $\eta_0 = 56 \text{ mPa s}$, $\eta_\infty = 3.45 \text{ mPa s}$, $\lambda_f = 3.313 \text{ s}$ and $b = 0.3568$ [508]. The tissue model is of HGO-type [335] with parameters taken from the literature [338, 509]. Here, a single layer is accounted for, neglecting intima and adventitia to reduce complexity in the mesh construction step. Regarding the structure’s material parameters we use $\nu_s = 0.499$, $\mu_s = 62.1 \text{ kPa}$, $\kappa_b = 2\mu_s(1 + \nu_s)/(3 - 6\nu_s) \approx 3.1 \text{ MPa}$ with fiber parameters chosen as $k_1 = 1.4 \text{ kPa}$, $k_2 = 22.1$ and $\kappa_c = 0.12$ and $\alpha_c = 27.47^\circ$.

The boundary conditions are set as follows. The tissue and lumen are fixed at the inlet and outlets and viscoelastic support with $k_e = 10^7 \text{ N/m}^3$ and $c_e = 10^5 \text{ Ns/m}^3$ is considered. A parabolic inflow profile is mapped onto the non-circular inlet as laid out in Sec. 4.4 to match volumetric flow rates given by Mills et al. [524]. The scaling has to be computed each time the inlet position changes, that is, only once per simulation when employing a fixed inlet as in the present case. From an initial quiescent state




Fig. 8.47: Physiological inlet velocity scale \bar{u}_{in} and pressure \bar{p}_f [524] and computed pressure p_f in the reference point marked in Fig. 8.46.

$\mathbf{u}_f = \mathbf{0}$ at $t = 0$, a smooth ramping (8.3) to $T_r = 0.5$ s is used. The outlet pressures are determined via Windkessel models with $C_i = 1.5 \times 10^{-8}$ m³/Pa, $R_{p,i} = 2 \times 10^8$ Pa s/m³, $R_{d,i} = 10^9$ Pa s/m³, $p_{p,i}|_{t=0} = 0$ Pa and $p_{d,i} = 7$ kPa for both outlets $i = 1, 2$, such that pressures in the physiological range [524] are obtained.

In total, six cardiac cycles are computed, and the chosen parameters lead to the pressure in the reference point marked in Fig. 8.46 matching the reference pressure from literature [524], denoted as \bar{p}_f , after completion of the prestress phase. A comparison is shown in Fig. 8.47, which additionally contains the temporal scaling of the inlet profile \bar{u}_{in} . Naturally, a better fit might be achieved fine-tuning the Windkessel parameters or employing more elaborate models (e.g., [272]), but since this numerical example does not consider patient-specific data, the remaining discrepancy is neglected. Note also that employing Windkessel models might not yield smaller discrepancies at all, given the model’s restrictions and limited tuning possibilities.

Prestress in the tissue present at time of image acquisition is accounted for by a precursor flow simulation up to $t = 2.5$ s, where the diastolic phase of the isolated, periodic flow solution is reached. As detailed in Sec. 6.4, the diastolic fluid load on the tissue is kept fixed and the pseudo timestepping and continuation method is started with $\Delta t^0 = 1$ ms, scaling ρ_s , c_e and η_c over 100 increments. A quasi-static solution with tolerances of $\epsilon_{abs}^{d_0} = 10^{-20}$, $\epsilon_{rel}^{d_0} = 10^{-10}$ and $\epsilon_{rel}^{\dot{d}_0} = 10^{-3}$ is reached.

The vector field \mathbf{d}_0 entering the structure’s momentum balance equation via a constant contribution $\mathbf{S}_0(\mathbf{d}_0)$ is shown in Fig. 8.48. The vector field \mathbf{d}_0 resembles simple “inflation” caused mainly by the diastolic pressure. Note, however, that this vector field \mathbf{d}_0 is *not* to be understood as the initial condition for \mathbf{d}_s , \mathbf{d}_s^0 , at the start of the FSI simulation. In this sense, Fig. 8.48 merely shows the vector field \mathbf{d}_0 yielding $\mathbf{S}_0(\mathbf{d}_0)$ solving $-\hat{\nabla} \cdot \mathbf{S}_0 = \mathbf{0}$ in $\hat{\Omega}_s$ (for details see Sec. 6.4), while the structure’s displacement, velocity and acceleration are initialised as $\mathbf{d}_s^0 = \dot{\mathbf{d}}_s^0 = \ddot{\mathbf{d}}_s^0 = \mathbf{0}$.

Let us also note here that the prestress-incorporating \mathbf{d}_0 is constructed at $t = 2.5$ s, which lies in the very beginning of the diastolic phase. This time is deliberately chosen to highlight the robustness of the scheme given higher pressure levels. Due to the prestress algorithm (see Sec. 6.4), the fluid load from a pure flow simulation is counteracted by \mathbf{d}_0 . However, as can be seen from Figs. 8.47 and 8.50, accounting for FSI leads to a




Fig. 8.48: Vector field \mathbf{d}_0 counteracting diastolic fluid loads via $\mathbf{S}_0(\mathbf{d}_0)$ generated at $t = 2.5$ s.

vastly different pressure wave, such that incorporating prestress does not lead to zero displacements at any point in the cardiac cycle as one might expect. This effect is amplified by (i) the flow field never returning to the precursor pure fluid simulation due to the deformation of the computational domain and (ii) the adopted method by Hsu and Bazilevs [318] greatly simplifying incorporation of non-infinitesimal \mathbf{d}_0 by performing an additive split $\mathbf{P} \leftarrow \mathbf{F}(\mathbf{S} + \mathbf{S}_0)$ to counteract the prestress load. That being said, we additionally want to emphasise that the present choice of prestress algorithm can be exchanged easily, not adapting any of the remaining steps within the partitioned FSI algorithm and only serves as one out of many possible options.

As the FSI simulation is started from $t = 2.5$ s with nonzero pressure, fluctuations are seen in the reference point’s pressure shown in Fig. 8.47 for a short period of time. Thus, the time instant of allowing for structural displacements is a critical one, since small variations of the overall volume can lead to drastic pressure changes due to the Windkessel outlets. Here, the IQN-ILS scheme and even more so high Robin parameters η_s^R robustify the coupling scheme. Judging from Fig. 8.47, the periodic state might be reached after a single FSI cycle, but nonetheless, solution snapshots presented in Fig. 8.49 consider the solution in the sixth cardiac cycle at $t = 5.31$ and 5.49 s. At these distinct points, either the inlet velocity or the pressure and displacement are maximal. Owing to the rapid reduction of inflow and the slightly curved geometry, strong recirculations are observed, resulting in large velocity gradients and hence large differences in the apparent viscosity. Despite the fact that studying such a scenario in detail might be of high interest, we herein focus on the numerical performance of the coupling schemes.

In analogy to previous investigations, the second-order accurate scheme is found rather unstable, especially when accounting for Windkessel outlets and adaptive timestepping. Therefore, a trivial “extrapolation” of the velocity $\mathbf{u}_f^* = \mathbf{u}_f^n$ ($m_u = 1$) is considered. Let us note here, that this effect is again much less relevant for uniform time steps or when enforcing constant tractions at the outlets—which we, however, do not aim for. Instead, we start from $\Delta t^0 = 1$ ms, and select the time step size according to the CFL condition, ensure $\Delta t \leq 5$ ms and allow for a maximum growth of 5% per time step. Varying the Robin parameter via α_R in $\eta_s^R = \alpha_R \rho_f / \Delta t^n$, the solutions obtained via Dirichlet–Robin coupling clearly converge despite keeping $\eta_s^R > 0$ for increased stability. As Fig. 8.50 indicates, setting $\text{CFL}_{\max} = 0.6$ and computing the solution for $\alpha_R = 10$ and $\alpha_R = 1$ already agree quite well in amplitude and phase of the pressure and displacement in the reference point. Reducing the Robin parameter further by choosing $\alpha_R = 0.1$, oscillations in the pressure increase, leading to a less stable scheme. Naturally, decreasing the time



Fig. 8.49: Iliac bifurcation: fluid velocity vectors and structural displacement in the deformed fluid and cut solid domains. Instants of maximum inlet velocity ($t = 5.31$ s, top) and maximum displacement and pressure (5.49 s, bottom) in the last cardiac cycle. Rapidly decreasing inflow causes a flow redirection.

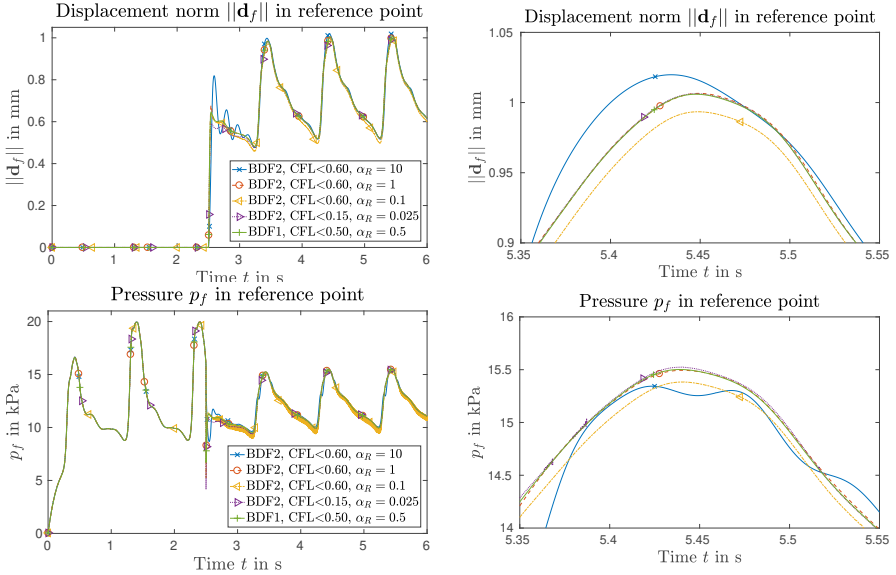


Fig. 8.50: Iliac bifurcation: structural displacement \mathbf{d}_s (top) and fluid pressure p_f (bottom) in the reference point. Overall view including the prestress phase from $t = 0$ to $t = 2.5$ s (left) and zoom-in on $t = 5.35$ to 5.55 s (right).

step size and keeping the Robin parameter scale in the same range by an identical CFL_{\max} to α_R ratio, we obtain a stable solution, which parallels the observations in Sec. 8.3 on the pulsatile flow through a straight vessel. In comparison to that, the IQN-ILS-accelerated SIDN scheme cannot be used for the full duration of the simulation, and is frequently found diverging especially in the initial moments after the prestress phase when employing Windkessel models and adaptive timestepping.

For the sake of completeness, a fully first-order accurate scheme (“BDF1”) is included in Fig. 8.50 as well. This scheme delivers similar results to the former presented second-order scheme with $m_u = 1$. Strictly speaking, both schemes are first-order accurate in time due to the low-order linearisation of the fluid velocity.

Inspecting now the pressure Poisson solves per time step over the course of the entire simulation, an increase from 1 initially in the precursor flow simulation to ≈ 3 per time step at $t = 2.5$ s is seen. Fig. 8.51 shows the mean number of coupling iterations over time t for the six computed cardiac cycles. Naturally, the targeted CFL_{\max} and resulting number of time steps influence the amount of accumulated PPE solves. Decreasing CFL_{\max} while keeping the Robin parameter η_s^R the same leads to fluctuating Δt^n , which might suggest that further tuning the time step selection would be beneficial or that the added-mass effect’s influence is already increased noticeably. Given these fluctuations, the first-order initial guess might in fact be preferred over the chosen second-order initial guess. It is thus not surprising that $\Delta t \rightarrow 0$ leads to a slight increase in coupling iterations in this scenario as well.

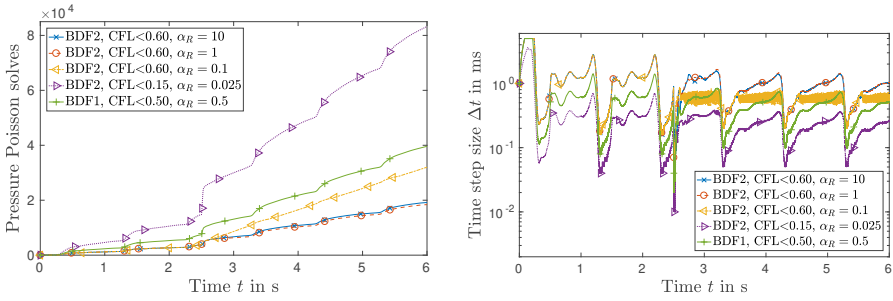


Fig. 8.51: Blood flow through an iliac bifurcation: accumulated pressure Poisson solves (left) and time step size (right). At $t_0 = 2.5$ s the prestress is calculated based on a pure CFD simulation, after which the fluid–structure coupling is accounted for.

Tab. 8.7: Iliac bifurcation: absolute (relative) time step number N_t and coupling iterations N_{FSI} in the FSI phase with different targeted CFL_{max} and scaling parameters α_R using $m_u = 1$, i.e., $\mathbf{u}_f^* = \mathbf{u}_f^u$.

BDF m	CFL_{max}	α_R	$\alpha_R/\text{CFL}_{\text{max}}$	N_t	N_{FSI}	N_{FSI}/N_t
2	0.60	10.0	16.67	5445 (1.02)	15442 (1.05)	2.84 (1.03)
2	0.60	1.0	1.67	5356 (1.00)	14679 (1.00)	2.74 (1.00)
2	0.60	1.0	0.167	6985 (1.30)	28290 (1.93)	4.05 (1.48)
2	0.15	0.025	0.167	22143 (4.13)	68469 (4.66)	3.09 (1.13)
1	0.50	0.5	1.00	10703 (2.00)	32206 (2.19)	3.00 (1.10)

From the presented cases, $\text{CFL}_{\text{max}} = 0.6$ yields the lowest overall number of time steps and coupling iterations as listed in Tab. 8.7. As turns out, this solution with $\alpha_R = 1$ is faster compared to $\alpha_R = 10$, while delivering results very much similar to the reference solution computed with $\alpha_R = 0.025$ and $\text{CFL}_{\text{max}} = 0.15$. The linear solvers of each of the subproblems converge nicely as in the previous Sec. 8.5 and are omitted here for brevity. A change in the convergence criteria of the linear solvers, the tolerance used in Newton’s method, or the parameters for the adaptive time step selection might alter results obtained here, but a large potential for performance gains is evident. Additionally, the Robin coupling scheme is found much more robust when accounting for prestress in the reference geometry following the strategy in Sec. 6.4.

8.7 Aortic dissection

In this final numerical example in the cardiovascular context, simulations of aortic dissection (AD) based on patient-specific data provided by Bäumlér et al. [266] are carried out. This cardiovascular disease motivated some specific innovations by the authors and is thus presented in various occasions within this work, but summarised here again for

the convenience of the reader. AD is characterised by delamination and rupture of the aortic wall, such that blood can penetrate the dissected vessel layers and thereby creates the so-called false lumen. Acute cases of AD often demand immediate surgical intervention, whereas chronic AD demands continuous surveillance and long-term medical treatment. We refer to [356, 357] for a rigorous discussion on diagnosis, pathology and treatment options of AD and focus on the computational modelling instead. Further, we neglect rupture processes to reduce complexity, acknowledging their importance for the initiation and progression of aortic dissection [9, 332]. In chronic cases of AD, capturing the dissection flap's motion, flow splits between outlets and stresses in the tissue are of interest [10–12, 266]. Digital twinning and virtual surgery or in-silico studies regarding suitable long-term medical treatment are possible future applications of a software framework capable of reproducing in-vivo measurements. The aim within this work is to demonstrate applicability of the PPE-based split-step framework adopting patient-specific data and parameters in the physiological range wherever available. Therefore, we present here some first tests being part of ongoing investigations on a Stanford Type B AD case of a 52-year old male.

Medical images were processed by Bäumler et al. [266], based on which the spatial discretisation shown in Fig. 8.52 is constructed by extending the interface normal in two steps by 0.8 mm and 1.2 mm with *Meshmixer* [513], sweeping the medial and adventitial tissue layers of uniform thickness, depicted in Fig. 8.52 using red and peach colour tones, respectively. Towards the vessel's centerline, a similar approach leads to a boundary layer in the fluid mesh, shown in grey in Fig. 8.52, to better capture large gradients in the fluid velocity field. The dissection flap, shown in light blue in Fig. 8.52, is detected by simply marking the structure's elements on the exterior and subtracting them from the element set followed by subtraction of the first 5 layers of elements touching the previously subtracted layer. A close-up of the finite element mesh respecting tissue layers and the dissection flap is given in Fig. 8.53, where different colours indicate these element sets. Overall, a finite element mesh with $\approx 2.42 \times 10^6$ tetrahedra and 4.35×10^5 nodes results, leading to $\approx 3.27 \times 10^6$ nodal DoFs in the PPE-based split-step scheme, which might be considered the lower limit in this complex topology using a conforming boundary layer tetrahedral mesh and rather coarse meshing via *Gmsh* [511] of the remaining bulk fluid and solid volumes.

As can be seen from Fig. 8.52, the proximal intimal tear—the point of layer separation closest to the heart—is located at the left subclavian artery in the ascending arch and extends all the way down to the iliac arteries. There, re-entry tears connect true and false lumina, such that merely the left subclavian artery is supplied by both lumina as can be seen from Figs. 8.52 and 8.53. The remaining branch vessels are supplied by the true lumen only, but large parts of the aorta are dissected with a false lumen and hence the thin dissection flap originating at the aortic arch and ending at the lower end of the considered geometry posterior to the iliac bifurcation.

Zooming in on the arch region now, Fig. 8.53 depicts the spatial discretisation including all tissue layers and the lumen in Fig. 8.53(a), where cutting parts of the media and adventitia from the model reveals the conforming fluid mesh (grey) underneath as shown in Fig. 8.53(b). Omitting the fluid from Fig. 8.53(b), the dissection flap coloured in light blue is visible in Fig. 8.53(c). In the arch, the lumen positioned anterior (bigger lumen




Fig. 8.52: Geometric discretisation of the aortic dissection case: partitioned finite element mesh (a) composed of adventitia (b) and media (c) tissue layers, dissection flap (d) separating the true and false lumina and fluid mesh (e).

in the front) corresponds to the false lumen, whereas the posterior lumen (smaller lumen in the back) under close inspection turns out to be the true lumen, which is mostly occupied by the dissection flap in this anterior-posterior view.

The orientation algorithm with standard settings ($\alpha_{\text{tol}} = 120^\circ$, $h_{l,i} = 100$ and $N_{\text{avg}} = 5$) leads to a satisfactory material orientation, following the tissue's longitudinal and circumferential directions as can be seen in Fig. 8.54. In this scenario involving a tissue layer wetted from both sides, the algorithm laid out in Sec. 6.1 performs significantly better than a standard approach involving the consecutive solution of two artificial Laplace problems as demonstrated in Sec. 6.1.

The physiological flow rate as measured in-vivo via 4D flow MRI [266] is prescribed at the aortic root by (i) mapping a quadratic velocity profile defined over a circular inlet onto the present non-circular inlet employing the technique described in Sec. 4.4, yielding the spatial scale and (ii) temporal scaling to achieve a mean inlet velocity \bar{u}_{in} , both depicted in Fig. 8.55, such that the inlet profile is given by

$$\mathbf{u}_{\text{in}}(\mathbf{x}, t) = -\mathbf{n}_f \eta_r(t) \bar{u}_{\text{in}}(t) \frac{Q_{\text{circ}}(t)}{\bar{Q}_{\text{in}}(t)} \tilde{u}_{\text{in}}(\tilde{r}(\mathbf{x}, t)). \quad (8.6)$$

The inlet $\Gamma_{\text{in},f}^t$ has an area of $|\Gamma_{\text{in},f}^t| \approx 5.84 \text{ cm}^2$, such that with the temporal scale $\bar{u}_{\text{in}}(t)$, a maximum volumetric flow rate of $Q_{\text{max}} \approx 481.6 \text{ cm}^3/\text{s}$, which is equivalent to 28.9 l/min, results from enforcing the rescaled $\mathbf{u}_{\text{in}}(\mathbf{x}, t)$. The negative normal vector $-\mathbf{n}_f$ enforces inflow normal to $\Gamma_{\text{in},f}^t$ and η_r is a smooth ramp from $t = 0$ to $t = T_r = T_p/2$

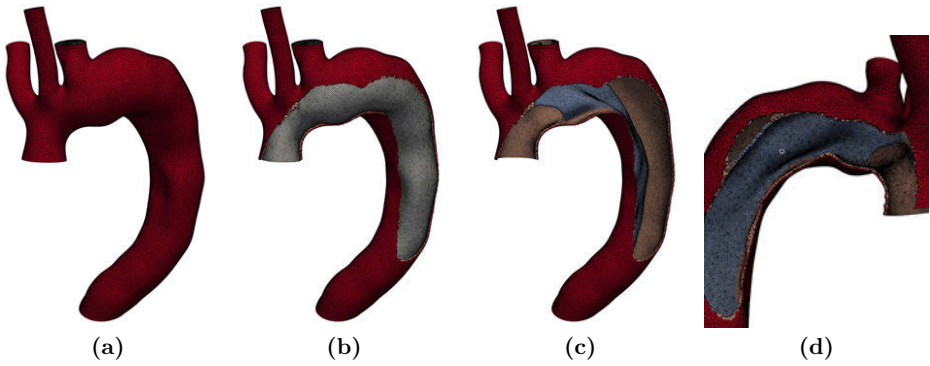


Fig. 8.53: Finite element discretisation of the dissected aortic arch: adventitia (red), media (peach) and flap (light blue) regions of the tissue and lumen (grey). Complete mesh (a), cut aortic arch revealing lumen mesh (b) and cut aortic arch showing tissue elements only (c), all in anterior-posterior view, and reference point A (o) on the true lumen side of the dissection flap in posterior-anterior view (d).

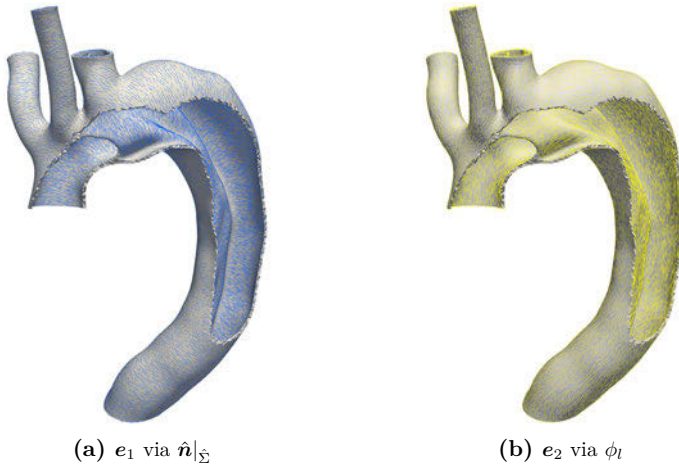


Fig. 8.54: Material orientation in the aortic arch region: circumferential directions e_1 computed from extrapolated, conditionally-averaged interface normal $\hat{\mathbf{n}}|_{\Sigma}$ (a) and longitudinal direction e_2 computed from a Laplace equation in ϕ_l (b).




Fig. 8.55: Ascending aortic arch with media (peach) and adventitia (red) tissue layers and spatial velocity scale (left, high values in red) and mean temporal velocity scale (right) on the inlet $\Gamma_{in,f}^t$, where $|\Gamma_{in,f}^t| \approx 5.84 \text{ cm}^2$, such that $Q_{\max} \approx 481.6 \text{ cm}^3/\text{s}$ equivalent to 28.9 l/min.

with $T_p = 0.78 \text{ s}$, given as

$$\eta_r(t) = \begin{cases} \sin^2\left(\frac{\pi t}{2T_r}\right) & \text{for } t \leq T_r = T_p/2 = 0.39 \text{ s,} \\ 1 & \text{otherwise.} \end{cases} \quad (8.7)$$

Tuning the Windkessel parameters to recover measured flow rates and realistic pressure levels at each of the outlets is a time-consuming, tedious process. In fact, tuning the parameters by hand in the present scenario with a total number of 12 Windkessel outlets, an error inevitably remains. Within this work, we fit the parameters of each Windkessel outlet to in-silico flow rates provided by Bäumlér et al. [266] in a one-by-one manner, simply ignoring the necessary coupling. Thus, the Windkessel parameters provided in Tab. 8.8 suffice for a first numerical tests, but a more rigorous data fitting is necessary. However, given that realistic pressure levels and flow rates are obtained, we can still argue that general applicability of the scheme can be shown and statements regarding the scheme’s performance translate to a setting with “perfect” parameter fit. Regarding the remaining boundary conditions, we fix the inlet and outlet faces and incorporate viscoelastic support with $k_e = 10^7 \text{ N/m}^3$, $c_e = 10^5 \text{ Ns/m}^3$ and $p_e = 0 \text{ Pa}$ chosen in a similar range as [266, 359, 360, 363].

Concerning material properties, we consider physiological parameters as introduced in Sec. 8.3 for the two-layered straight vessel; that is, fluid and structural densities of $\rho_f = 1060 \text{ kg/m}^3$ and $\rho_s = 1200 \text{ kg/m}^3$, respectively, and a Carreau fluid with viscosity limits $\eta_0 = 39.13 \text{ mPa s}$ and $\eta_\infty = 5.13 \text{ mPa s}$ and fitting parameters $\lambda_f = 0.9003 \text{ s}$, $b = 0.3224$ provided by Ranftl et al. [18], which corresponds to 60% haematocrit. The tissue is modeled as a layered HGO solid with $\nu_s \approx 0.499$ and shear rates $\mu_{s,1} = 120.0 \text{ kPa}$, $\mu_{s,2} = 62.1 \text{ kPa}$, and $\mu_{s,3} = 21.6 \text{ kPa}$ for the dissection flap (1), media (2) and adventitia (3). Fiber reinforcement is taken into account with $k_1 = 1.4 \text{ kPa}$, $k_2 = 22.1$ and $\kappa_{c,1} = \kappa_{c,2} = 0.12$, $\alpha_{c,1} = \alpha_{c,2} = 27.47^\circ$ and $\kappa_{c,3} = 0.25$, $\alpha_{c,3} = 52.88^\circ$ [338, 509], with identical parameters for the dissection flap (1) and medial layer (2). The dissection flap’s stiffness is increased to prevent excessive motion, which otherwise leads to contact of the

Tab. 8.8: Relative flow rates and Windkessel parameters per outlet approximating flow and pressure data from [266] (a.=artery, l.=left, r.=right, ext.=external, int.=internal, true and false refer to true and false lumen outlets).

	flow fraction %	C_i $10^{-10} \text{ m}^3/\text{Pa}$	$R_{p,i}$ $10^8 \text{ Pa s}/\text{m}^3$	$R_{d,i}$ $10^7 \text{ Pa s}/\text{m}^3$
Brachiocephalic trunk	20.2	14.43	4.16	3.74
L. common carotid a.	5.8	4.12	14.48	13.03
L. subclavian a. (true)	7.6	5.40	11.05	9.95
L. subclavian a. (false)	5.9	4.19	14.24	12.81
celiac trunk	12.6	8.95	6.67	6.00
Superior mesenteric a.	4.1	2.91	20.49	18.44
R. renal a.	7.2	5.11	11.67	10.50
L. renal a.	8.4	5.96	10.00	9.00
R. ext. iliac a.	10.9	7.74	7.71	6.94
R. int. iliac a.	4.5	3.20	18.67	16.80
L. int. iliac a.	9.0	2.63	22.70	20.43
L. ext. iliac a.	3.7	6.39	9.33	8.40

dissection flap and the media and non-invertibility of the ALE map, i.e., breakdown of the algorithm.

The time interval considered ranges from $t = 0$ to $t = 3.9$ s, such that a total of five cardiac cycles are computed using adaptive timestepping with initial step length of $\Delta t^0 = 1$ ms and different CFL numbers targeted. The BDF2 and CH- α ($\rho_\infty = 0$) schemes are combined with linear extrapolation of all involved fields but the fluid velocity, which is considered with first-order accurate $m_u = 1$ ($\mathbf{u}_f^* = \mathbf{u}_f^n$) to enhance temporal stability in case of adaptive timestepping and Windkessel models applied. Doing so, excessive pressure fluctuations are suppressed.

Convergence criteria are set to $\epsilon_{\text{abs}}^p = \epsilon_{\text{abs}}^d = 10^{-7}$ and $\epsilon_{\text{rel}}^p = \epsilon_{\text{rel}}^d = 10^{-3}$, while we focus on the SIDR scheme with scaled Robin parameter $\eta_s^R = \alpha_{R^{\rho f}}/\Delta t^n$. To accelerate coupling convergence, we employ IQN-ILS with settings almost unchanged from the previous examples being $\omega_0 = 0.01$, $q = 10$ and $\epsilon_{QR} = 10^{-16}$. Newton's method is executed using absolute and relative convergence criteria $\epsilon_{\text{abs}}^{\text{nl}} = 10^{-10}$ or $\epsilon_{\text{rel}}^{\text{nl}} = 10^{-4}$ in (5.70) and GLS stabilisation is considered to stabilise convective terms.

The prestress algorithm adopts flow data from the second cardiac cycle's diastolic phase at $t = 1.5$ s, applying the interface traction over 50 load steps with increased viscoelastic bedding $c_e = 10^9$, reaching a stationary state with relative and absolute convergence criteria of $\epsilon_{\text{rel}}^{\dot{d}_0} = 10^{-5}$ and $\epsilon_{\text{rel}}^{d_0} = 10^{-9}$ within a total of 220 pseudo time steps. Note here that in contrast to [266], we do incorporate the prestress tensor \mathbf{S}_0 in the dissection flap region, as otherwise \mathbf{d}_0 cannot yield an \mathbf{S}_0 solving the stationary momentum balance

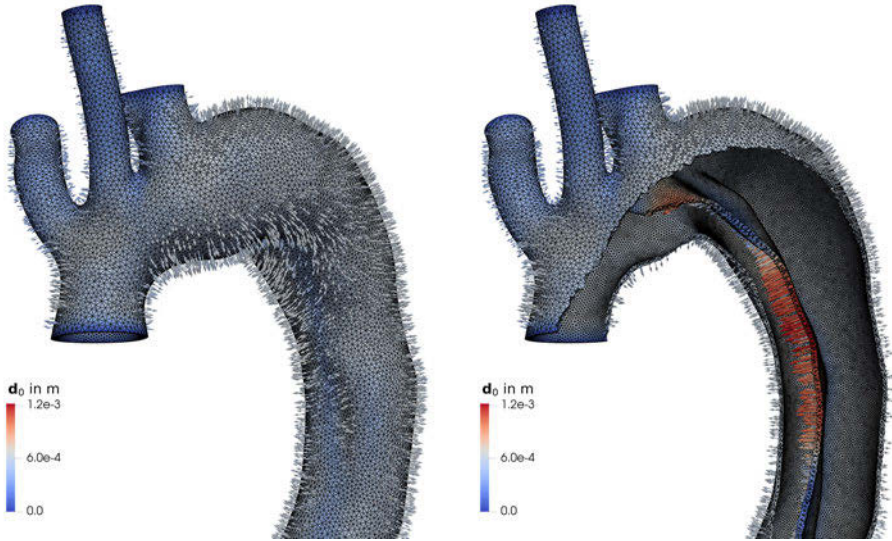


Fig. 8.56: Vector field \mathbf{d}_0 counteracting the diastolic fluid load via $\mathbf{S}_0(\mathbf{d}_0)$ generated in the second cardiac cycle’s late diastole at $t = 1.5$ s. Close-up of the aortic arch region in anterior-posterior view: exterior and cut revealing the dissection flap.

equation in $\hat{\Omega}_f$. However, the prestress-inducing vector field \mathbf{d}_0 shows the largest values in the dissection flap region, as can be seen from Fig. 8.56. Similar to the numerical example of an iliac bifurcation in Sec. 8.6, \mathbf{d}_0 resembles “inflating” the vessel, but here, the pressure difference between true and false lumen at $t = 1.5$ s additionally leads to a nonzero \mathbf{d}_0 in the dissection flap.

To ease convergence of the FSI coupling loop and to further robustify the prestress algorithm, we additionally add linear weak (mass-proportional) damping to the structure’s momentum balance equation, that is, a term of the form

$$c_m(t)\langle \boldsymbol{\varphi}, \dot{\mathbf{d}}_s \rangle_{\hat{\Omega}_s},$$

with a time-dependent parameter c_m being 0.1 initially, and smoothly decreased via $1 - \eta_t(t)$ (8.7) with $T_r = 0.2$ s between $t = 1.5$ s and $t = 1.6$ s to $c_m(t = 1.6) = 0.01$. Mass-proportional damping or including viscoelastic effects of the structure equations is a well-known technique for improving temporal stability, coupling convergence speed and preventing finite time blow-up [25, 529–531], which has been applied in the cardiovascular context, e.g., in [268, 479, 530, 531]. Here, we add just enough mass-proportional damping after stabilising the initial “release” after the prestress phase, such that the periodic solution is almost unaffected. Radtke et al. [479] show, that for physiological blood flow and pressure levels in a bent vessel of HGO material, $c_m = 0.01$ and $c_m = 0.001$ already yield almost indistinguishable transients on the cardiac cycle’s scale, such that we can expect minor alteration of the results. Nonetheless, this parameter requires rigorous studying similar to α_r , in the Robin boundary condition or the temporal and spatial res-

olutions used in the approximation, which is why selected parameter combinations with negligible mass-proportional damping are included in Fig. 8.63, showing that indeed, $c_m = 0.01$ yields both the desired temporal stabilisation while only slightly affecting the solution.

The finite element mesh is distributed to 32 subdomains as shown in Fig. 8.52(a), completely ignoring the fluid–structure interface. With the current number of subdomains, such an approach leads for the current mesh to a partitioning mostly into (bent) vessel pieces with subdomain interfaces being perpendicular to the vessel centerline. Increasing the core count further will inevitably lead to idle processors, since some subdomains may not contain a well-balanced amount of fluid and solid subdomains. Therefore, a further increase of the spatial resolution might require a partitioning of the fluid and solid subdomains separately, which is not considered here.

Now, let us start by inspecting the flow field and tissue deformation within the fifth considered cardiac cycle from $t = 3.12$ to $t = 3.9$ s, obtained via the SIDR scheme with $\text{CFL}_{\max} = 0.6$ and $\alpha_R = 0.025$. As can be seen from Fig. 8.57, an approximately periodic flow solution is reached after two cardiac cycles (see curve “CFD”). Therefore, the pre-stress is computed at $t = 1.5$ s, incorporating the diastolic flow field and corresponding interface tractions as load. In Figs. 8.58–8.61, three distinct points in the cardiac cycle are shown, namely, peak systole at $t = 3.29$ s, where the peak flow rate is reached, mid systole at $t = 3.35$ s with the maximum displacement and high pressure level in the reference point and mid diastole at $t = 3.71$ s. Inspecting the pressure field over time in the aortic arch (see Fig. 8.58), we observe pressure levels in the physiological range, with large fluctuations over time, but also significant differences in space, i.e., between the inlet and outlets due to the Windkessel models and between the true and the false lumina due to topology. The maximum pressure drop over the dissection flap is in the range of a few hundred Pascal, since it strongly depends on the stiffness of the dissection flap being a thin, deformable membrane only bearing little load before deforming excessively.

The viscosity changes drastically over the cardiac cycle as can be seen from Fig. 8.59. During systole, the volumetric flow rate in the true lumen (smaller diameter) is higher, leading to a viscosity close to the Newtonian limit, whereas in diastole, the low shear rates lead to an increase in viscosity, approaching the other end of the parameter range, η_0 . Moreover, large differences are seen in space, since velocities in the false lumen (larger diameter) are limited and the circulatory flow field seen in the arch region and past the intimal tear gets close to laminar quickly. This flow field is indicated by velocity vectors shown in Fig. 8.60, also showing the deformed vessel. Due to the position of the false lumen, recirculatory flow is seen in the proximal end of the false lumen, where redirected flow from the ascending aorta enters. In the true lumen, peak flow leads to well-organised, non-turbulent flow. The flow split between true and false lumina is one-sided in favour of the true lumen, since the false lumen has only a single outlet at the left subclavian artery (where the Windkessel model dictates the pressure and hence the volumetric flow rate) and otherwise connects back to the true lumen distal to the iliac bifurcation. Interestingly, the point of maximal displacement of the dissection flap coincides with a narrowing of the true lumen, such that a high flow rate leads to an increased pressure drop between false and true lumina. As depicted in Fig. 8.61, the thin

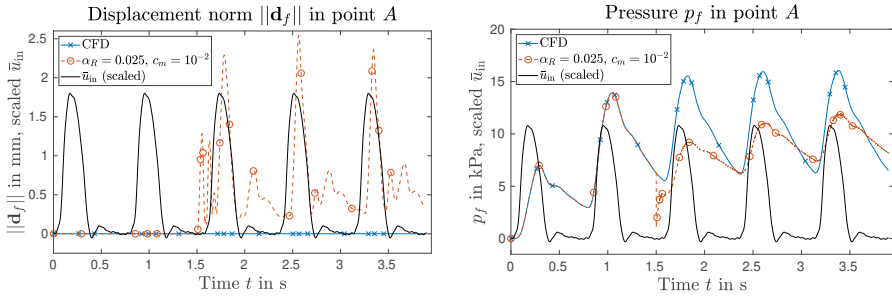


Fig. 8.57: Aortic dissection: structural displacement \mathbf{d}_s (left) and fluid pressure p_f (right) in the reference point comparing a flow simulation (“CFD”, \times) and the SDR scheme ($\alpha_R = 0.025, c_m = 10^{-2}$, \circ), both adopting BDF2 time integration and $CFL_{max} = 0.6$. The inflow scale (\bar{u}_{in} , $-$) is included for reference, showing that the flow field is approximately periodic from the second cycle on, such that at $t = 1.5$ s the prestress is computed in the FSI case.

elastic membrane thus undergoes a “flapping” motion, while the media and adventitia tissue layers on the vessel tree’s exterior are simply expanding in normal direction because of the elevated pressure compared to the prestress reference time $t = 1.5$ s. Regarding the non-zero displacement in diastole despite \mathbf{d}_0 considered, the same remarks as in the iliac bifurcation case apply (see Sec. 8.6).

The large displacements observed for $\alpha_R > 1.0$ require an increased dissection flap stiffness $\mu_{s,1} = 120$ kPa $\approx 2\mu_{s,2}$ to avoid contact of the dissection flap with the medial tissue layer for the selected numerical parameter combination. With this increased flap stiffness, the maximum dissection flap displacement in normal direction in the reference point ranges from -1 mm to $+3$ mm, see Fig. 8.61, which underestimates measurements via 4D flow MRI yielding 8.7 mm [266]. Note, however, that the selected parameters were chosen for demonstration purposes to vary CFL_{max} and α_R over a wider range. Variations of α_R and CFL_{max} suggest that a further decrease of the dissection flap stiffness might in fact be possible, since $\alpha_R \ll 1.0$ results in a maximum dissection flap displacement of ≈ 3 mm. The cases with Robin parameters chosen *too high* ($\alpha_R \geq 1$ in this specific case) yield larger displacements, which determined the tissue stiffness to prevent self-contact. Therefore, careful parameter selection to obtain temporally converged results while using a smaller dissection flap stiffness is a matter of parameter fitting and/or demands contact treatment, which we will not focus on. Instead, the increased dissection flap stiffness chosen for demonstration purposes allows us to employ larger variations in α_R and CFL_{max} close to 1, which would otherwise quickly lead to self-contact if $\mu_{s,1} = 62.1$ kPa $= \mu_{s,2}$ due to excessive weight put on the extrapolated quantities (and larger displacements due to lower stiffness, of course).

We close the presentation of exemplary results with the shear stress vector $\bar{\boldsymbol{\tau}}$ on the fluid–structure interface as defined in Eqn. (8.5) and shear rate $\bar{\dot{\gamma}}$, both time-averaged over the fourth cardiac cycle according to Eqn. (8.4). As shown in Fig. 8.62, shear in the false lumen is reduced, while increased shear can be identified in the true lumen, especially in regions with either redirecting flow (in the aortic arch and iliac arteries) or

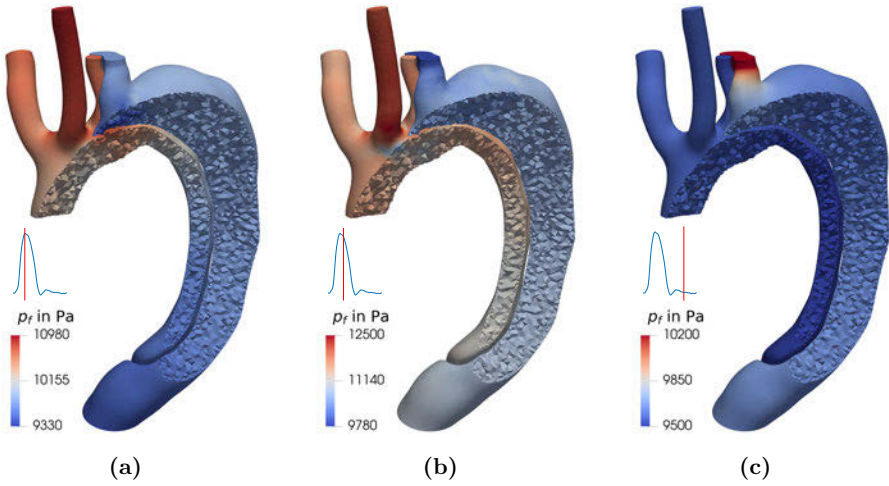


Fig. 8.58: Arch region in the AD case: pressure p_f in Ω_f^t at the fifth cardiac cycle's peak systole ($t = 3.29$ s, left, max. inflow), mid systole ($t = 3.35$ s, middle, max. displacement) and mid diastole ($t = 3.71$ s, right).

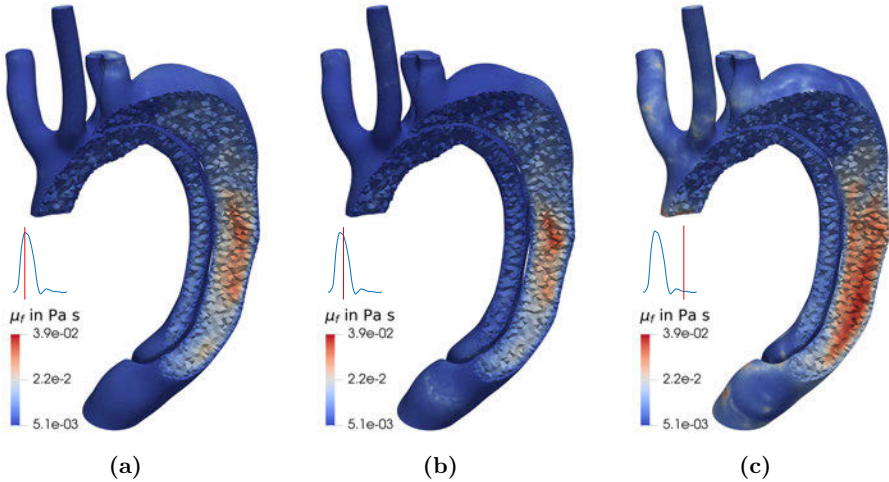


Fig. 8.59: Arch region in the AD case: viscosity μ_f in Ω_f^t at the fifth cardiac cycle's peak systole ($t = 3.29$ s, left, max. inflow), mid systole ($t = 3.35$ s, middle, max. displacement) and mid diastole ($t = 3.71$ s, right).

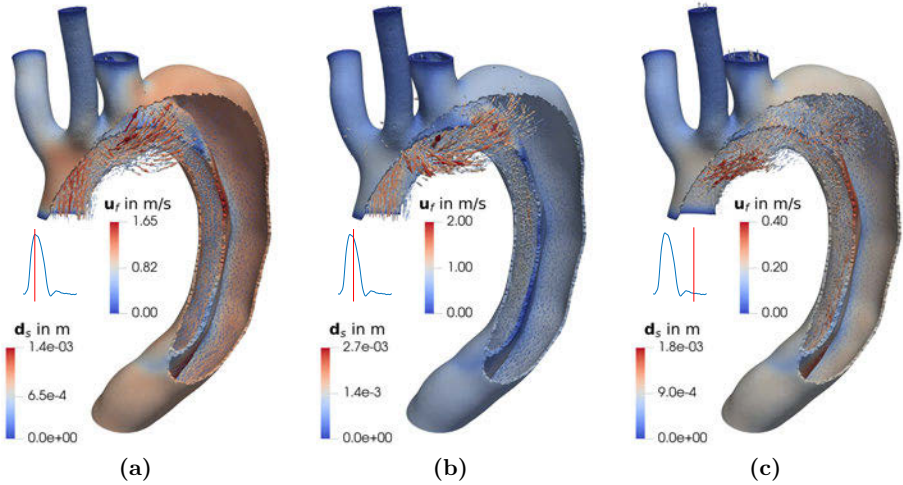


Fig. 8.60: Arch region in the AD case: velocity u_f in Ω_f^t and structural displacement d_s in Ω_s^t at the fifth cardiac cycle's peak systole ($t = 3.29$ s, left, max. inflow), mid systole ($t = 3.35$ s, middle, max. displacement) and mid diastole ($t = 3.71$ s, right).

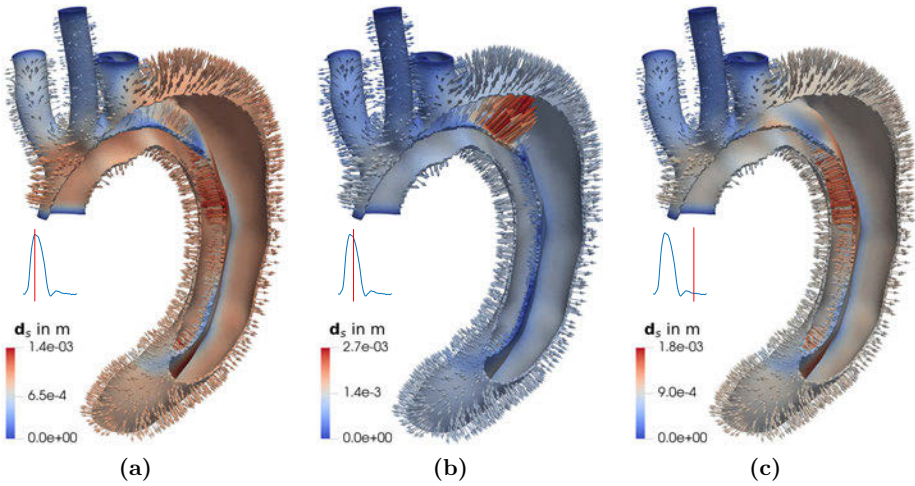


Fig. 8.61: Arch region in the AD case: structural displacement d_s in Ω_s^t at the fifth cardiac cycle's peak systole ($t = 3.29$ s, left, max. inflow), mid systole ($t = 3.35$ s, middle, max. displacement) and mid diastole ($t = 3.71$ s, right).

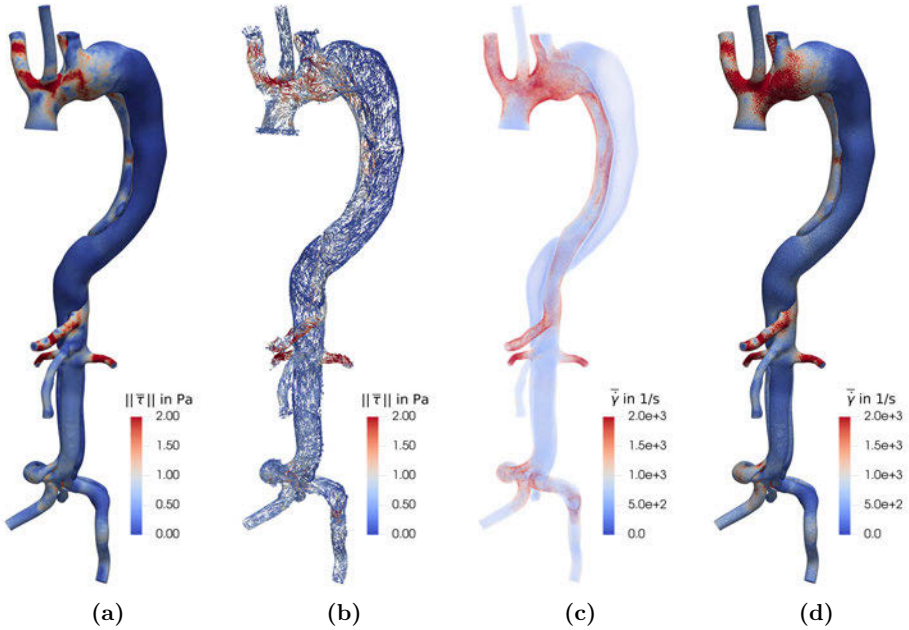


Fig. 8.62: Overall view of the AD case: area-weighted nodal average of the shear stress vector $\bar{\tau}$ on the fluid–structure interface (a,b) and shear rate $\bar{\dot{\gamma}}$ (c,d), all time-averaged over the fourth cardiac cycle.

larger variations of the fluid velocity over the vessel cross section (in the brachiocephalic trunk, left subclavian artery, narrowed true lumen, celiac trunk and the renal arteries) due to comparably large volumetric flow rates.

Now, turning our attention to the coupling scheme’s performance, we consider a reduced time interval from $t = 0$ to $t = 3$ s to reduce numerical effort. The scaling parameter α_R in the Robin interface condition on the structure side in the SIDR scheme is varied to give $\eta_s^R = \alpha_R \rho_I / \Delta t^n$ as in the previously considered examples. The tissue displacement and fluid pressure in the reference point on the dissection flap’s true lumen side highlighted in Fig. 8.53(d) are used as reference values, given that displacements are approximately the maximally observed ones.

For a first comparison, Figs. 8.63 and 8.64 also contain transients obtained via a fully first-order accurate time integration scheme (“BDF1”), a simulation with reduced mass-proportional damping used from $t = 1.6$ s onwards and results obtained with stricter convergence criteria $\epsilon_{\text{rel}}^p = \epsilon_{\text{rel}}^d$. Differences in the solution are in all cases small, while the increased quality of the initial guess in the second-order scheme does not yield lower iteration counts in comparison to BDF1 in the present scenario. Mass proportional damping does also not significantly alter the number of FSI coupling steps, but robustifies the coupling scheme by damping high frequency modes introduced by the Windkessel model and potentially worsened initial guesses when adaptively setting time step sizes. Then, higher tolerance criteria can be used, while still reaching convergence. The number

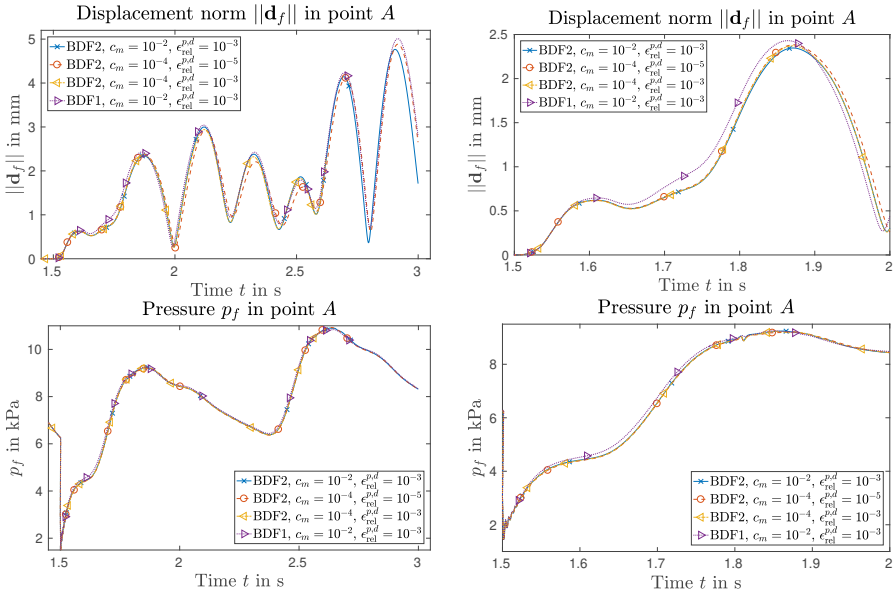


Fig. 8.63: Aortic dissection: structural displacement \mathbf{d}_s (top) and fluid pressure p_f (bottom) in the reference point using CFL < 0.6 , $\alpha_R = 1.0$ and varying mass-proportional damping c_m and convergence criteria $\epsilon_{rel}^{p,d} = \epsilon_{rel}^p = \epsilon_{rel}^d$. Overall view excluding the prestress phase from $t = 0$ to $t = 3.0$ s (left) and zoom-in on $t = 1.5$ to 2.0 s (right).

of total executed PPE solves once again largely depends on the number of time steps used as can be seen from Fig. 8.64.

Next, in a second study, we compare various α_R and target CFL numbers, where one can clearly observe from Fig. 8.65 that solutions converge as α_R decreases. However, α_R has to be chosen suitably—low enough for accuracy and high enough to counteract added-mass instabilities—to obtain accurate results within an acceptable number of coupling steps. As Tab. 8.9 highlights, the number of semi-implicit coupling steps does only mildly depend on α_R , while its use is primarily motivated by improving temporal stability. Then, higher CFL numbers can be employed, which significantly reduces the number of time steps executed. In the considered time interval, choosing the time step size based on the CFL condition, time step sizes vary by an order of magnitude between diastolic and systolic phases and total executed PPE solves once again largely depend on the number of time steps used as can be seen from Fig. 8.66. In a nutshell, the very same trends as in the numerical experiments targeting flow through a straight vessel (Sec. 8.3) and an iliac bifurcation (Sec. 8.6) can be identified.

At this point, comparing various constitutive models and material parameters is highly interesting, but requires much greater care be taken, e.g., ensuring converged solutions with respect to the spatial discretisation, time integration, tuning material and numerical parameters and more. However, this numerical experiment only serves as a proof of applicability within this work, while the rigorous comparison of constitutive models is

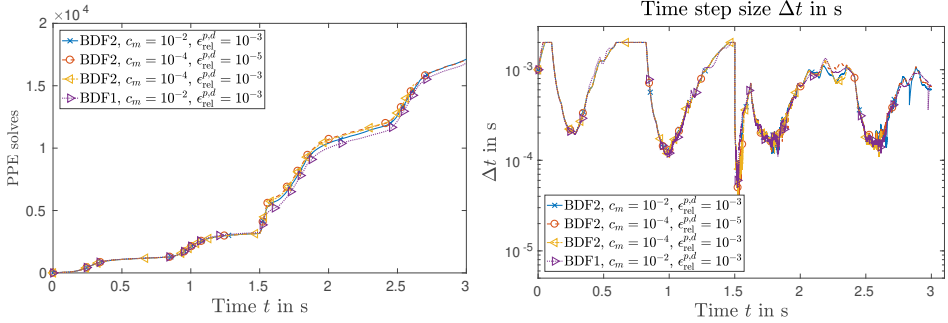


Fig. 8.64: Aortic dissection: accumulated pressure Poisson solves (left) and time step size (right) using $CFL < 0.6$, $\alpha_R = 1.0$ and varying mass-proportional damping c_m and convergence criteria $\epsilon_{rel}^p = \epsilon_{rel}^d = \epsilon_{rel}^d$. At $t_0 = 1.5$ s the prestress is calculated based on a CFD simulation, after which the fluid-structure coupling is accounted for.

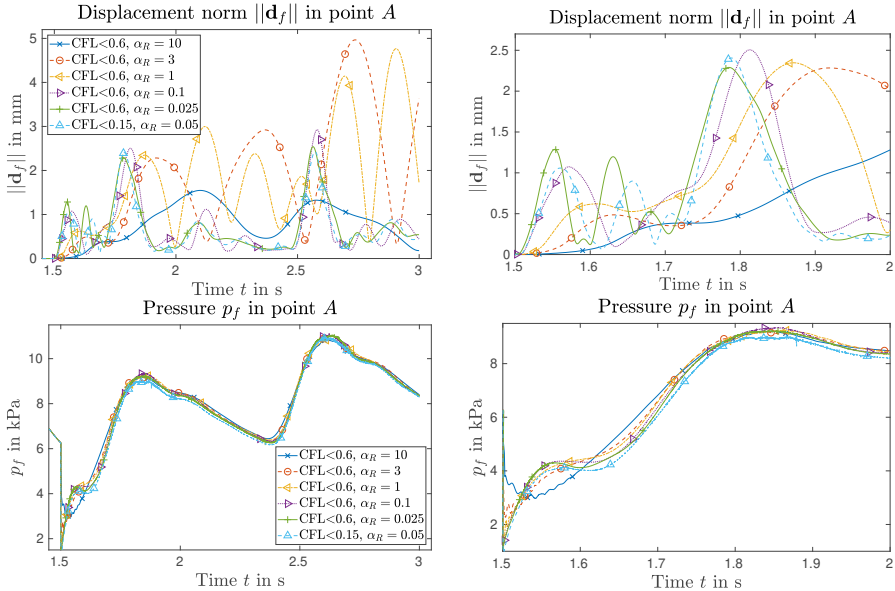


Fig. 8.65: Aortic dissection: structural displacement \mathbf{d}_s (top) and fluid pressure p_f (bottom) in the reference point using the SIDR scheme with BDF2 ($m_u = 1$) and $c_m = 10^{-2}$. Overall view excluding the prestress phase from $t = 0$ to $t = 1.5$ s (left) and zoom-in on $t = 1.5$ to 2.0 s (right).

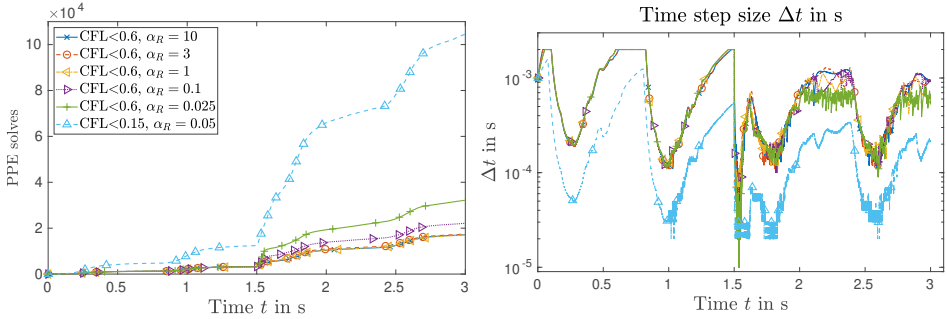


Fig. 8.66: Aortic dissection: accumulated pressure Poisson solves (left) and time step size Δt (right) using the SIDR scheme with BDF2 ($m_u = 1$) and $c_m = 10^{-2}$. At $t_0 = 1.5$ s the prestress is calculated based on a pure CFD simulation, after which the fluid–structure coupling is accounted for.

Tab. 8.9: Aortic dissection: absolute (relative) time step number N_t and coupling iterations N_{FSI} in the FSI phase with different targeted CFL_{max} and scaling parameters α_R using $m_u = 1$, i.e., $\mathbf{u}_f^* = \mathbf{u}_f^n$ and $c_m = 10^{-2}$.

BDF m	CFL_{max}	α_R	$\alpha_R/\text{CFL}_{\text{max}}$	N_t	N_{FSI}	N_{FSI}/N_t
2	0.60	10	16.67	4567 (1.02)	13856 (1.02)	3.03 (1.00)
2	0.60	3	5.00	4652 (1.04)	14164 (1.04)	3.04 (1.00)
2	0.60	1	1.67	4563 (1.02)	13963 (1.03)	3.06 (1.01)
1	0.60	1	1.67	4456 (1.00)	13562 (1.00)	3.04 (1.00)
2	0.60	0.5	0.83	4737 (1.06)	14540 (1.07)	3.07 (1.01)
2	0.60	0.1	0.167	4985 (1.12)	18961 (1.40)	3.80 (1.25)
2	0.60	0.025	0.042	5672 (1.27)	29021 (2.14)	5.12 (1.68)
2	0.15	0.05	0.333	22293 (5.00)	92107 (6.79)	4.13 (1.36)

Tab. 8.10: Relative timings in % per time step for computationally relevant steps indicated by the unknown solved for split into (assembly+AMG setup+linear solve) for PPE and solid momentum step. Comparing software versions and settings used in the AAA and AD cases, one sees an improvement in the structure element integration and relative solver timings due to adapted parameters.

	\mathbf{d}_f	\mathbf{u}_f	ψ	ζ	p_f	\mathbf{d}_s
AAA case	1.6	1.7	2.2	0.6	(1.3+1.9+10.4)	(47.6+3.7+26.9)
AD case	2.9	16.8	3.6	0.3	(7.0+0.4+0.7)	(40.5+3.6+9.39)

part of ongoing investigations. This demonstrates that the proposed PPE-based, semi-implicit and IQN-ILS-accelerated coupling scheme can be readily applied to challenging scenarios, where it bears great potential in terms of reducing the overall computing time by reducing the number of required time steps, FSI coupling steps and even reduces the cost of individual FSI coupling iterations merely requiring a PPE solve and structure update per coupling step.

The implementation of the FSI framework is currently still under improvement, such that the timings reported within this thesis can only be interpreted relative to one another, since aspects of (serious) high-performance computing are not in the main focus (while parallelisation through MPI and scalable iterative solvers via `deal.II` and Trilinos' ML or PETSc's BoomerAMG are employed). To highlight these improvements, we compare a previous version of the software used for the results presented in Sec. 8.5 [5] with the current version [6] in Tabs. 8.10 and 8.11. The two considered numerical tests both involve a Holzapfel–Gasser–Ogden material model and a Carreau fluid applied to (i) an abdominal aortic aneurysm (AAA) yielding 0.9×10^6 DoFs distributed to 8 subdomains, solved with the Aitken-accelerated SIDN scheme and a fixed time step size of $\Delta t = 1$ ms (see Sec. 8.5) and (ii) the present AD case with 3.3×10^6 DoFs distributed to 32 subdomains solved with the IQN-ILS-accelerated SIDR scheme ($\alpha_R = 0.1$) and adaptive timestepping targeting $\text{CFL}_{\max} = 0.6$. Keeping the DoFs per core roughly constant, a striking difference is seen. Thanks to adaptive timestepping, IQN-ILS-accelerated Robin coupling as well as optimisations in the structural solver's element integration and AMG settings, the AD case involving a factor of 3.67 more DoFs is solved in 17.7% of the wall time denoted as t_{wall} . It is thus clear that continued performance optimisation and more elaborated tuning of the numerical parameters might further decrease the time to solution. Let us also note here that tests regarding scalability to large problem sizes and core counts are required, but not considered herein.

With this, the applicability of the split-step scheme in the haemodynamic context has been sufficiently demonstrated in numerical experiments ranging from blood flow in a straight vessel, an idealised abdominal aortic aneurysm, a patient-specific geometry of an iliac bifurcation up to a patient-specific aortic dissection case. The semi-implicit coupling schemes are, fortunately, not necessarily limited to cardiovascular problems, but can also be applied in other flow/FSI regimes with a less dominant added-mass effect as is seen in the following section.

Tab. 8.11: Performance comparison of the AAA and AD cases shows a speed-up of 3 for a 3.67 times larger problem while keeping the DoFs per core roughly constant due to adaptive timestepping, IQN-ILS-accelerated Robin coupling, improvements in the structural element integration and AMG settings.

	computing time t_{wall}	N_{FSI}	$\frac{t_{\text{wall}}}{N_{\text{FSI}}}$	$\frac{10^6 \text{DoFs}}{\text{core}}$
AAA case	1.09×10^6 s	17846	61.1 s	0.11
AD case	1.93×10^5 s	20670	9.3 s	0.10

8.8 Human phonation

In this final numerical example, we aim to show the applicability of the accelerated semi-implicit methods outside of the cardiovascular regime, that is, in the context of aeroelasticity. Here, we investigate human phonation as yet another challenging application of FSI in biomechanics. Human phonation describes the process of sound creation by elastic tissue, the so-called vocal folds, interacting with air flow from the lungs (see, e.g., [382, 383, 532–534]). Our motivation for considering such a problem is the vast difference in the physical parameters and flow regime present. This scenario is related with high Reynolds numbers and low, but non-zero added-mass effects. A geometry from [532] is adapted marginally, while physical parameters are taken from [532, 533] in the physiological range.

The vocal folds are located within the glottis, which's symmetry plane is considered here as a simplified, straight channel from $x = 0$ to $x = L$ with the channel's length $L \approx 65.4$ mm and a height of $H = 18$ mm. Two pairs of vocal folds are placed in the glottis, the first pair, the (true) vocal folds, with a height of 8.9 mm are positioned 7.5 mm upstream the false vocal folds with a height of 6.5 mm. This configuration is depicted in Fig. 8.67, with the inlet being on the left-hand side. These four elastic bodies are subject to air flow and hence deform. This might lead to contact of the vocal folds, exchanging energy and synchronising their oscillation under healthy conditions. However, in glottal insufficiency, the two vocal folds do not come into contact during the oscillation cycle [535]. Such conditions can be modeled via a gap between the vocal folds (first, upstream pair) of 0.2 mm as shown in Fig. 8.67, such that contact of submerged structures can be neglected.

Air is modeled as a Newtonian fluid with a dynamic viscosity of $\mu_f = 0.019$ mPa s and density $\rho_f = 1.145$ kg/m³. The tissue in this scenario, i.e., the vocal folds are considered having a density of $\rho_s = 1200$ kg/m³ and linear elastic material with material parameters $E_s = 20$ kPa and $\nu_s = 0.45$. FSI is considered starting from the quiescent state, i.e., $\mathbf{u}_f = \mathbf{0}$, ramping the prescribed inlet velocity smoothly from $t = 0$ to $T_r = 0.01$ s via

$$\mathbf{u}_f|_{x=0} = (\eta_t 0.6, 0)^\top, \quad \text{with} \quad \eta_t = \begin{cases} \sin^2\left(\frac{\pi t}{2T_r}\right) & \text{for } t \leq T_r, \\ 1 & \text{otherwise,} \end{cases}$$




Fig. 8.67: Human phonation with glottal insufficiency: idealised glottis' symmetry plane, i.e., computational domain (top) and close-up of the interglottal channel's grid and the top vocal fold's trailing edge (bottom) with reference points A (\times) and B (\circ). Structural domain $\hat{\Omega}_s$ in dark grey, fluid domain $\hat{\Omega}_f$ in light grey.

Note that a *constant* inlet velocity of 0.6 m/s is reached after the initial ramp-up phase, which is distinctly different from many studies considering only fluid dynamics. Often, a periodic solution is achieved with periodically scaling the volumetric flow rate or the pressure drop, a fundamental difference from our approach here. This way, we can observe the self-excited motion of the vocal folds considering for fluid–structure interaction. This allows recovering the natural frequency of the vocal folds' oscillatory response under *constant* flow rates [383, 533, 534].

Regarding the remaining boundary conditions, we adopt zero traction conditions at $x = L$, thereby prescribing a zero reference pressure and no-slip boundary conditions at the channel's walls. The computational domain is fixed on the entire boundary, such that the linear elastic structure deforms due to air flow only. This entire setup allows capturing vibrations of the vocal folds, but neglects stress present in the vocal folds due to missing pressure. Such a simplifying assumption might be mended considering nonzero pressure at the outlet and viscoelastic support including a suitable external pressure, but within this example, the focus lies on applying the PPE-based semi-implicit split-step scheme in an aeroelastic context and not necessarily on capturing the vocal folds' stress states, as the assumption of linear elastic material behaviour might have already suggested.

Using a ramped inlet velocity profile to reach a constant inlet velocity of 0.6 m/s results in a maximum velocity of 35 m/s in the interglottal channel between the vocal folds. The corresponding Reynolds number is $\text{Re} \approx \mathcal{O}(1000)$, such that a fine mesh as shown

in Fig. 8.67 is required to yield element Reynolds numbers $Re_e \approx 250$. The considered spatial discretisation is rather coarse and can only resolve main flow features, while fine-scales are completely neglected in this preliminary study. This is accepted due to $CFL_{\max} < 1$ requiring a small time step size $\Delta t < 3.4 \times 10^{-6}$ already due to the current grid selected. In this numerical experiment, we choose $CFL_{\max} = 0.2$ conservatively to ensure temporal stability. Setting the time interval from $t = 0$ to $T = 0.15$ s, we are already forced to execute more than 4×10^5 time steps, also motivating the two-dimensional setup considered. Choosing a wider glottal channel—as often resorted to in literature—decreases peak velocities and reduces the necessary numerical effort by reducing velocity gradients in the flow field. Since we aim for a proof of concept more than a rigorous investigation employing the semi-implicit FSI solver in a demanding application in aeroelasticity, however, such simplifications not capturing all the solution’s details are simply accepted.

Concerning the time discretisation, an adaptive first-order scheme, i.e., BDF1 combined with CH- α ($\rho_\infty=0$) timestepping with $\Delta t^0 = 10^{-4}$ and $CFL_{\max} = 0.2$ is employed. GLS and backflow stabilisations are considered and convergence criteria $\epsilon_{\text{abs}}^p = \epsilon_{\text{abs}}^d = 10^{-7}$ and $\epsilon_{\text{rel}}^p = \epsilon_{\text{rel}}^d = 10^{-3}$ are set. We focus on the SIDR scheme with scaled Robin parameter $\eta_s^R = \alpha_{R^p t} / \Delta t^n$ combined with IQN-ILS acceleration and settings completely unchanged from the previous examples being $\omega_0 = 0.01$, $q = 5$ and $\epsilon_{QR} = 10^{-16}$.

The solution reaches a periodic state after an initialisation phase with distinct points shown in the snapshots in Fig. 8.68. Due to the narrowing of the channel and the resulting constrained flow through the vocal folds, pressure increases, causing the vocal folds to deform. Afterwards, the pressure is lowered, thereby causing the elastic vocal folds to return to their original position. Due to inertia, the vocal folds pass their initial, undeformed configuration, narrowing the interglottal channel further. This causes the pressure to rise yet again, deforming the vocal folds. Thus the cycle repeats, causing the onset pressure to increase periodically and the vocal folds to oscillate around their stress-free configuration.

This oscillatory motion consists of multiple modes as analysed, e.g., by Luo et al. [383]. Moreover, the vocal folds’ displacement patterns are unsymmetric past a certain time due to vortex shedding given the high Reynolds number and the missing exchange of momentum in the case of glottal insufficiency. The vocal folds’ apex points A and B (see Fig. 8.67) thus follow a complex pattern, of which the last four periods from $t \approx 0.12$ to 0.15 s are depicted in Fig. 8.69. Here, a fundamental frequency of ≈ 130 to 150 Hz can be measured, which is perfectly within the expected range [382, 532] factoring in the vast simplifications (two-dimensional setup, linear elasticity, zero reference pressure, fixed channel, pinned vocal folds, no turbulence model applied).

High velocity gradients are observed in the resulting flow field, with vortices of variable size and fluctuations in the pressure. When the vocal folds are almost in contact, reducing the glottal channel’s width, onset pressure is increased drastically, forming a strong jet. As shown in Fig. 8.70, this strong jet triggers pressure fluctuations being a key aspect in the generation of the human voice. Changing the problem setup by lowering the stiffness of the vocal folds or increasing the flow rate, the vocal folds might get in contact, which requires special treatment when employing an ALE formulation.

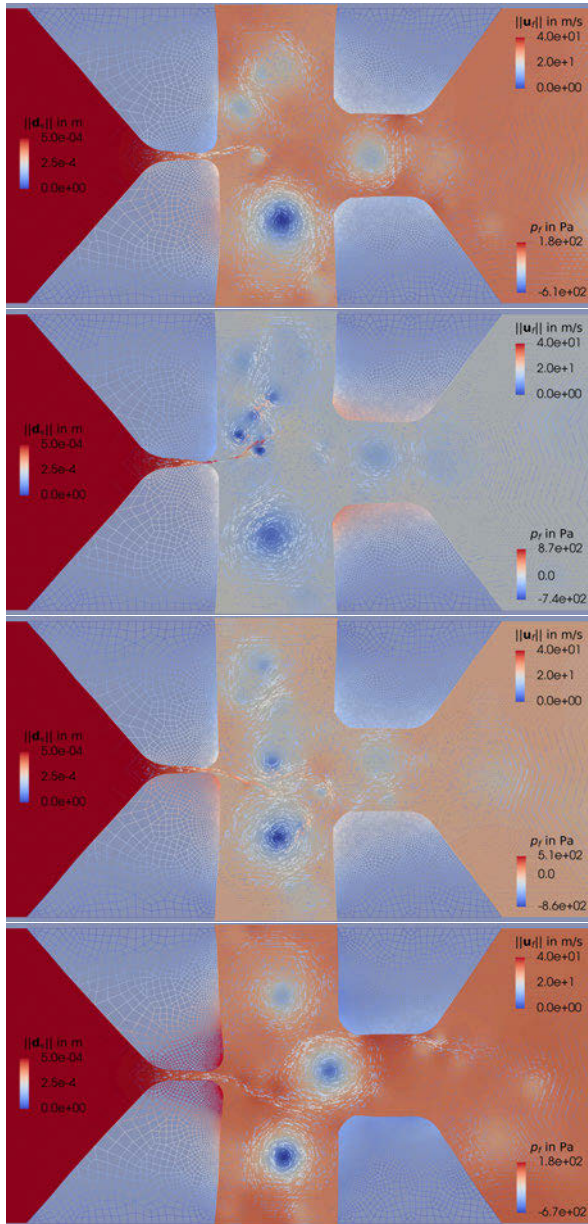


Fig. 8.68: Human phonation with glottal insufficiency: u_f and p_f in Ω_f^t and d_s on the structural grid's edges. Periodic solution (from top to bottom): (i) stress-free configuration with a 0.2 mm gap, (ii) narrowing of the glottal channel, (iii) rising onset pressure triggering strong jet and vocal fold deformation, (iv) maximum deformation and reducing pressure drop, return to (i).




Fig. 8.69: Human phonation with glottal insufficiency: displacement $\|\mathbf{d}_s\|$ in the vocal folds' apex points A (\times) and B (\circ) highlighting the multi-modal, non-symmetric motion pattern of the two vocal folds.




Fig. 8.70: Human phonation with glottal insufficiency: selected velocity vectors \mathbf{u}_f over pressure p_f and deformed solid mesh with narrowing vocal folds triggering high velocities and pressure fluctuations.

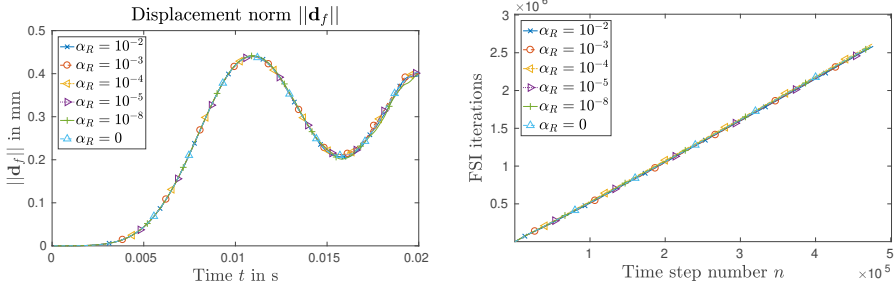


Fig. 8.71: Human phonation with glottal insufficiency: semi-implicit Dirichlet–Robin scheme with $\eta_s^R = \alpha_R \rho_f / \Delta t^n$, where α_R has little influence on d_f (left) and does not accelerate coupling convergence (right).

Tab. 8.12: Human phonation with glottal insufficiency: mean FSI iteration counts depending on Robin parameter scaling α_R .

α_R	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-8}	0
mean FSI iter./step	5.42	5.42	5.50	5.43	5.46	5.48

To examine now the coupling schemes’ performance, we compare the number of accumulated FSI coupling iterations over time in Fig. 8.71 using various α_R . Clearly, the number of required coupling iterations remains almost constant independent of the Robin scaling parameter as shown by the straight line in Fig. 8.71. A total number of 4.7×10^5 time steps are performed with an average size of 3.2×10^{-7} s aiming for $\text{CFL}_{\max} = 0.2$. An average of ≈ 5.48 FSI coupling steps are needed per time step, summing up to more than 2.56×10^6 iterations in the semi-implicit coupling loop.

Given the physical parameters in the present scenario, no significant improvements in terms of accelerating convergence with Robin interface conditions are achieved as highlighted in Tab. 8.12. A Robin parameter chosen too high leads to divergence of the solver as in previously considered examples in the haemodynamic context, since explicit contributions to the interface terms reduce temporal stability. Thus, using Robin conditions is not recommended here, even though the solution does not differ significantly when increasing the Robin parameter by several orders of magnitude as shown in Fig. 8.71. Within the interval $t = 0$ to $t = 0.02$ s, the solution compares well, but differences increase when a long time interval and hence large numbers of time steps are considered. This, however, is caused by the linearisations, loose tolerance settings and nonlinear influence of past time steps’ solutions.

In a nutshell, this last numerical experiment in the context of aeroelasticity demonstrates the versatility of the semi-implicit coupling scheme, but improvements applying Robin interface conditions can not be expected given the physiological parameter range and the added-mass effect being less dominant compared to the cardiovascular context. The PPE-based semi-implicit split-step framework with GLS and backflow stabilisations

shows great robustness in this parameter range, while the increased numerical effort connected with other coupling schemes renders them ineffective, as long as temporal stability limits cannot be circumvented. Using a split-step scheme, such a constraint limits time step sizes as demonstrated. The semi-implicit coupling combined with IQN-ILS acceleration performs well, simply using the same parameters as in the haemodynamic range. However, the gains in performance are less impressive here, since the added-mass effect reduces as $\rho_s \gg \rho_f$. A semi-implicit coupling scheme can be used nonetheless to counteract build-up of instabilities over time. In the present scenario, fully-explicit schemes performing only a single iteration were found stable for large portions of the considered time interval, but are not sufficiently robust to complete the entire simulation.

8.9 Summary and conclusion

Within chapters 7 and 8, we present a family of FSI schemes involving three-dimensional solids interacting with incompressible viscous flows, where the latter are solved via the newly proposed PPE-based split-step scheme decoupling the fluid's velocity and pressure fields. As a consequence, substantial parts, namely, all steps but the pressure Poisson equation and the structural solver may be treated in an explicit fashion, while accuracy and added-mass stability are still preserved. The mesh update step, the fluid's momentum balance equation, an optional Leray projection step for improving mass conservation properties and a viscosity projection step accounting for generalised Newtonian fluids are only solved once per time step, such that the bulk of the numerical effort associated with these steps is reduced to a minimum. Favourable properties from the split-step flow solver translate to the split-step FSI framework, that is, rheological laws of generalised Newtonian fluids are effortlessly exchanged adapting only the right-hand side of the viscosity projection step and (lower) equal-order interpolation using C^0 -continuous finite elements is admissible.

Time integration is carried out combining BDF and corresponding extrapolation schemes for the fluid phase and generalised- α schemes in the structural solver, where we observe that temporal oscillations in the fluid's pressure are loosely connected to the spectral radius in the high frequency limit ρ_∞ . Choosing $\rho_\infty = 1$ as for the Newmark- β method results in second-order accurate solid accelerations (and third-order accurate displacements and velocities), but at the same time does not introduce desired numerical damping to the high-frequency components of the structural response. However, $\rho_\infty \neq 1$ renders the structure's acceleration first order accurate in time, such that the interface term coupling the PPE and structural solver involves a first-order error. Consequently, the pressure's approximation in time is reduced in accuracy, which might be mended adopting higher-order accurate versions of the generalised- α scheme, but is not considered within this work.

The PPE-based split-step FSI framework with its partitioned design then further allows considering all of the modelling aspects and numerical techniques highly relevant for patient-specific simulations as introduced in previous sections. Owing to this setup, the structural solver is thus easily exchanged, while for the fluid subproblem solver,

projection or split-step schemes are required to enable (the anyways optional) semi-implicit coupling. We also show how to account for lumped, zero-dimensional parameter models such as resistance-based boundary conditions or the three-element Windkessel model incorporating effects of the downstream vasculature in truncated domains, while preserving consistency between fluid momentum and PPE steps when stabilising re-entrant flow.

The prestress algorithm as laid out for the solid subproblem previously is then trivially extended to the FSI setup, where a precursor flow simulation is employed to determine diastolic fluid loads present at the time of image acquisition to construct a suitable prestress tensor counteracting exactly these interface tractions. The remaining extensions to the subproblem solvers in the context of biomedical applications such as constructing suitable material orientations or inlet profiles given a certain flow rate on the inlet or accounting for viscoelastic external tissue support are employed as well, since they are carried over from the individual subproblem solvers and thus trivially extend towards the FSI case.

Another central contribution of this work is the combination of (i) semi-implicit coupling, (ii) Robin interface conditions and (iii) acceleration schemes such as Aitken's relaxation or the Interface Quasi-Newton Inverse Least-Squares method, treating each of the iteratively coupled schemes as black boxes. This mix of numerical schemes results in increased robustness and tremendous reduction of the required number of coupling iterations compared to each of these schemes on their own, while individual iterations within the FSI coupling loop are much cheaper to execute since they involve only the PPE and solid momentum steps. Extensive numerical tests moreover suggest that the influence of the Robin parameter chosen as $\eta_s^R = \alpha_R \rho_f / \Delta t^n$ on the solution vanishes as $\Delta t \rightarrow 0$, such that accelerated convergence and increased robustness can really be exploited, while the Dirichlet–Robin and Dirichlet–Neumann schemes (of course) yield the same solution when fully converged.

With this, problems of practical relevance can be tackled with low enough Robin parameters preserving the solution's transients, which becomes particularly important when comparing the proposed scheme to fully explicit FSI schemes based on Nitsche-type mortaring or Robin–Robin coupling. Furthermore, a smooth transitioning from the proposed Dirichlet–Robin to standard Dirichlet–Neumann coupling paradigms without ill-conditioning of the linear systems is given naturally, simply setting $\eta_s^R = 0$. As a consequence, selecting suitable Robin parameters is thus much less critical, especially since for most applications, a low number of coupling iterations results already with the IQN-ILS accelerated semi-implicit Dirichlet–Neumann scheme. This gained versatility is advantageous when comparing the split-step FSI scheme with explicit (mostly first-order accurate) FSI schemes, which heavily rely on a suitable Robin parameter for temporal stability.

As a last vital ingredient to bridge the gap to practically relevant simulations, adaptive timestepping solely targeting temporal stability, i.e., obeying to a standard CFL condition is applied as a placeholder for more involved techniques. This drastically cuts down the number of required time steps, while not at all affecting the coupling schemes' performance. However, numerical evidence suggests that adaptive timestepping impairs the

pressure's temporal stability via fluctuations in the flow rates entering the Windkessel models, which dominate the pressure response and hence the overall FSI problem.

We first assess the scheme's accuracy in problems of academic nature, comparing the approximations to analytical solutions both in space and in time, employing Q_2/Q_1 and Q_1/Q_1 finite element pairs and combining BDF2 with generalised- α time integration schemes. The framework is further tested in a benchmark example describing blood flow through a straight vessel, where the just described fundamental relation between time step size and differences between Dirichlet–Robin and Dirichlet–Neumann variants are investigated in the practically relevant parameter range. Here, the semi-implicit Robin–Neumann and Dirichlet–Neumann schemes outperform an Aitken-accelerated fully implicit scheme by a factor of 10 to 20, giving a first hint at the potential gains in efficiency. Incorporating adaptive time step selection, comparisons to simulations with a fixed time step in the pulsatile flow example show a possible speed-up of another factor of 5 on top, such that *compared to a naive Aitken-accelerated, fully implicit strategy with uniform timestepping, the computing time can be reduced by up to 2 orders of magnitude.*

The FSI framework is then applied to more challenging examples in the cardiovascular context and aeroelasticity, namely, blood flow through an idealised abdominal aortic aneurysm, a patient-specific iliac bifurcation geometry, a patient-specific case of aortic dissection and lastly air flow through the vocal tract in human phonation. In any of these scenarios, the applicability of the PPE-based split-step scheme is showcased, highlighting potential for application in clinical support, surgical planning or parameter studies on virtual cohorts of digital twins. However, we want to emphasise at this point once again that the presented results still retain their academic character owing to the fact that this work's focus lies on the development of incompressible flow and FSI solvers for biomedical applications, but not on actual patient-specific simulations themselves.

In a nutshell, we have developed and tested an *efficient* FSI framework in terms of required coupling iterations, total time steps and their related cost, which is due to its partitioned design *highly flexible* and can be applied to cardiovascular applications, where high added-mass effects are prevalent, but can also be used in scenarios where the pressure-structure coupling is less dominant. The FSI algorithm incorporates all modelling aspects and numerical techniques necessary to tackle realistic, patient-specific problems, which is successfully demonstrated in a series of challenging problems ranging from aortic blood flow to human phonation.

9 Concluding remarks and future perspectives

Within this work, we set out to develop and improve numerical methods for incompressible flow and FSI problems targeting haemodynamics, therefore putting a focus on blood flow through soft biological tissue, modelled via generalised Newtonian fluids and incompressible hyperelastic, fibre-reinforced solids. For patient-specific biomedical settings, robust, efficient and reliable methods are of crucial importance, as digital twinning, parameter studies on virtual cohorts and computer-simulated surgery for clinical support are promising fields for future application of potent software frameworks. The computational (bio-)mechanics community has seen tremendous progress over the past decade and continues to pursue the goal of improving health and wellbeing of mankind. The present work contributes to these efforts by advancing the field of tailored solution methods for generalised Newtonian fluid flow. We extend and improve techniques readily available for incompressible flow of Newtonian fluids towards the cardiovascular setting, allowing for variable viscosity to capture the complex rheology.

In this regard, a novel PPE-based stabilisation technique is introduced, allowing for equal-order interpolation within a coupled velocity-pressure formulation involving generalised Newtonian fluids. This basic scheme is then extended towards instationary flow, where solver aspects such as decoupling and linearising the velocity-pressure and viscosity systems and preconditioner design play a critical role especially when applied to large-scale real-world applications. Furthermore, we present two alternative methods for adaptively choosing the time step size via (i) a stabilised predictor-corrector approach or (ii) a fractional step θ -scheme with an embedded lower-order scheme. Several examples showcase the increased accuracy and robustness with respect to the stabilisation parameter, whose choice is critical, but in practical applications often left to the user.

Following this trail of thought, we then proceed in completely decoupling the velocity and pressure fields, deriving an additional Dirichlet condition for the pressure and hence rendering the PPE invertible based on a given velocity field. Such a technique frequently encountered in inverse haemodynamics—computing a pressure field from measured velocities—allows decoupling the momentum balance and PPE steps through higher-order extrapolations in (pseudo-)time. With this, only standard problems need to be solved, since suitably rewriting the PPE allows employing C^0 -continuous function spaces only. Based on this idea, an iteration-free split-step scheme is developed, which enables the use of equal orders for all involved fields, each of which are updated only once per time step.

Aiming for fluid–structure interaction, we summarise standard displacement-based formulations for (nearly) incompressible, hyperelastic and fibre-reinforced continua as a

placeholder for more complex tissue models to be coupled to the PPE-based split-step flow solver. The main reasons for limiting ourselves here to a coupling to the split-step scheme are (i) that semi-implicit FSI coupling is enabled and (ii) a comparison of the flow solvers showed that the scheme decoupling velocity and pressure is faster in our implementation, which heavily relies on block-based preconditioning via algebraic multigrid methods. Adopting a partitioned setup, one of the main contributions of this work is then the combination of semi-implicit coupling, Robin interface conditions and acceleration schemes such as Aitken's relaxation or the Interface Quasi-Newton Inverse Least-Squares method. This mix of numerical techniques leads in various challenging computational experiments to tremendous speed-ups, which is crucial for solving even a patient-specific case of aortic dissection using only 32 cores in a reasonable time.

Concerning such patient-specific simulations in the cardiovascular context, another important contribution of this work consists of carefully weaving in well-established modelling techniques capturing aspects vital in biomedical applications. These methods account for, e.g., inflow over non-circular inlets, the neglected downstream vascular network via lumped parameter models or the prestress present in the tissue at the time of image acquisition. We further propose a novel algorithm to automatically construct suitable material orientations in complex geometries involving thin structure layers wetted from both sides as encountered in aortic dissection. The PPE-based split-step scheme is then extended by classical stabilisations for reentrant flow over Neumann boundaries and convection-dominant flow. Within this work, each of these included models and methods can be interpreted as placeholders for more complex and potentially more accurate or more robust schemes, but were simply chosen for them being well-established and their applicability and reliability in the targeted applications.

At this point, we are confident in recommending the introduced methods as several numerical experiments in academic setups, benchmark cases and more applied or even patient-specific scenarios demonstrate the improvements over the state of the art in incompressible flow and FSI in haemodynamics. We sincerely believe the present work can function as a starting point, inspiration or stepping stone for future developments and improvements. Finally concluding this work, we want to hint at various possibilities for such extensions or point towards areas, where further research is required.

The most pressing issue related to this work is the missing sound mathematical analysis of the fundamental PPE-based schemes, which we do not provide. Here, we limit ourselves to rigorous testing and numerical experiments to provide convergence rates in space and time. Knowing that a sound mathematical basis is strongly required and hence a central open problem, we can for now only refer to numerous works in literature providing numerical evidence of the desired stability and convergence properties, even though proofs thereof are not yet available.

Higher-order methods in space and time can be employed to increase accuracy, were simply increasing the polynomial order of the Lagrangian finite elements, backward differentiation and extrapolation formulae combined with higher-order accurate generalised- α timestepping schemes for the structure are directly applicable. The proposed methods straight-forwardly extend to higher-order, but we herein focused on lower equal-order interpolation as the pragmatic choice. This is rooted in the expected input meshes being

rather fine but lower-order accurate due to challenging mesh construction in clinical applications, limiting the maximum admissible time step size via standard CFL conditions. In this regard, the FSI scheme's drawback of coupling via the (potentially first-order accurate) structure acceleration can be mitigated, leading to higher-order temporal accuracy of the pressure field as well. In the pure flow solvers too, the velocity and pressure approximations yield similar temporal convergence rates, where the pressure proved to be delicate once more, motivating further investigations.

In addition to that, we employ standard continuous Galerkin schemes to all of the encountered systems, but many of the involved subproblems and individual steps might be tackled with more appropriate numerical methods such as, e.g., discontinuous Galerkin formulations applied to the fluid momentum step, or following promising concepts such as matrix-free geometric (Newton-)multigrid. As pointed out already several times within this work, the proposed algorithms greatly benefit from them leading to a sequence of standard problems such as a Poisson equation to recover the pressure, a mass-matrix problem to project the viscosity, advection-diffusion-reaction equations governing the fluid velocity (components) or standard Lagrangian formulations of (hyper-)elasticity. Hence, advanced solution or preconditioning techniques may be readily applied to further facilitate a fast solution process.

The selection of a suitable Robin parameter in the Robin–Neumann FSI schemes—here chosen dependent on the time step size, but uniform in space—is a critical point demanding further attention. Comparing the present semi-implicit method to the family of currently rapidly progressing fully explicit coupling approaches, possible improvements might be achieved, whereas analogies of the Robin interface condition to methods exploiting artificial compressibility on the interface might also prove very useful. This aspect can yield vast improvements to the scheme as the Robin parameter is shown to be *the* central parameter to achieve a stable coupling scheme when considering physiological parameters, adaptive timestepping and lumped parameter models to determine adequate pressure levels and flow splits.

Naturally, the individual models incorporated in the overall framework might be extended to capture even more relevant physics. This is a never-ending endeavour, which we started in the direction of (i) thrombus formation [8] being central to various cardiovascular diseases (aneurysms, aortic dissection and more) and medical device design (blood pumps or stents) and (ii) multiphase flow [7], which requires adaptations to the PPE lying at the very heart of the presented methods. Concerning practical and patient-specific applications, higher-order mesh generation is also a matter of ongoing research as is the incorporation of more suitable subproblem solvers and models, most pressingly the inclusion of more stable lumped parameter models [272] and improved prestressing strategies such as [355].

In a closing remark we want to emphasise our sincere belief that computational science and engineering can assist in advancing medical care and device design and might be used, e.g., in training medical personnel. Further, computer tools can help discover and test new treatment options via digital surgery or parameter studies on virtual cohorts, in all scenarios reducing the risk involved while at the same time significantly reducing costs. Connecting back to the starting point of this thesis, we hope that the present work contributes its share to the development of computational (bio-)mechanics.

Bibliography

- [1] D.R.Q. Pacheco, R. Schussnig, O. Steinbach, and T.-P. Fries. A global residual-based stabilization for equal-order finite element approximations of incompressible flows. *Int. J. Numer. Methods Eng.*, 122(8):2075–2094, 2021.
- [2] R. Schussnig, D.R.Q. Pacheco, and T.-P. Fries. Robust stabilised finite element solvers for generalised Newtonian fluid flows. *J. Comput. Phys.*, 442:110436, 2021.
- [3] D.R.Q. Pacheco, R. Schussnig, and T.-P. Fries. An efficient split-step framework for non-Newtonian incompressible flow problems with consistent pressure boundary conditions. *Comput. Methods Appl. Mech. Eng.*, 382:113888, 2021.
- [4] D.R.Q. Pacheco. *Stable and stabilised finite element methods for incompressible flows of generalised Newtonian fluids*. PhD thesis, Technische Universität Graz, 2021.
- [5] R. Schussnig, D.R.Q. Pacheco, and T.-P. Fries. Efficient split-step schemes for fluid–structure interaction involving incompressible generalised Newtonian flows. *Comput. Struct.*, 260:106718, 2022.
- [6] R. Schussnig, D.R.Q. Pacheco, M. Kaltenbacher, and T.-P. Fries. Semi-implicit fluid–structure interaction in biomedical applications. *Comput. Methods Appl. Mech. Eng.*, 400:115489, 2022.
- [7] D.R.Q. Pacheco and R. Schussnig. Consistent pressure Poisson splitting methods for incompressible multi-phase flows: eliminating numerical boundary layers and inf-sup compatibility restrictions. *Comput. Mech.*, 2022. doi: 10.1007/s00466-022-02190-x.
- [8] R. Schussnig, S. Dreyman, A. Jafarinaia, T. Hochrainer, and T.-P. Fries. A semi-implicit method for thrombus formation in haemodynamic fluid–structure interaction. In *Proceedings of the 8th European Congress on Computational Methods in Applied Sciences and Engineering*, Oslo, 2022. submitted.
- [9] R. Schussnig and T.-P. Fries. A concept for aortic dissection with fluid–structure–crack interaction. *Proc. Appl. Math. Mech.*, 19(1):e201900100, 2019.
- [10] R. Schussnig, M. Rolf-Pissarczyk, G. Holzapfel, and T.-P. Fries. Fluid–structure interaction simulations of aortic dissection. *Proc. Appl. Math. Mech.*, 20(1):e202000125, 2021.
- [11] R. Schussnig, K. Bäumlner, and T.-P. Fries. Multi-layered tissue models in patient-specific simulations of aortic dissection. *Proc. Appl. Math. Mech.*, 21(1):e202100090, 2021.

- [12] R. Schussnig and T.-P. Fries. Coupled multiphysics modeling of aortic dissection. In *Proceedings of the 14th WCCM & ECCOMAS Congress 2020*, Virtual Congress, 2020.
- [13] G.P. Galdi, A.M. Robertson, R. Rannacher, and S. Turek. *Hemodynamical Flows*, volume 37 of *Oberwolfach Seminars*. Birkhäuser, Basel, 2008.
- [14] P. Nithiarasu. A fully explicit characteristic based split (CBS) scheme for viscoelastic flow calculations. *Int. J. Numer. Methods Eng.*, 60(5):949–978, 2004.
- [15] I.G. Donev and B.D. Reddy. Time-dependent finite element simulations of a shear-thinning viscoelastic fluid with application to blood flow. *Int. J. Numer. Meth. Fluids*, 75(9):668–686, 2014.
- [16] G. Böhme and J. Broszeit. Numerical flow simulation for Bingham plastics in a single-screw extruder. *Theor. Comput. Fluid Dyn.*, 9(1):65–74, 1997.
- [17] L. John, P. Pustějovská, and O. Steinbach. On the influence of the wall shear stress vector form on hemodynamic indicators. *Comput. Vis. Sci.*, 18(4-5):113–122, 2017.
- [18] S. Ranftl, T.S. Müller, U. Windberger, W. von der Linden, and G. Brenn. A Bayesian Approach to Blood Rheological Uncertainties in Aortic Hemodynamics. *Int. J. Numer. Method. Biomed. Eng.*, 2022.
- [19] D.R.Q. Pacheco and O. Steinbach. A continuous finite element framework for the pressure Poisson equation allowing non-Newtonian and compressible flow behavior. *Int. J. Numer. Meth. Fluids*, 93:1435–1445, 2021.
- [20] H. Johnston and J.-G. Liu. Accurate, stable and efficient Navier–Stokes solvers based on explicit treatment of the pressure term. *J. Comput. Phys.*, 199(1):221–259, 2004.
- [21] J. Liu. Open and traction boundary conditions for the incompressible Navier–Stokes equations. *J. Comput. Phys.*, 228(19):7250–7267, 2009.
- [22] R. Rannacher. On the numerical solution of the incompressible Navier-Stokes equations. *J. Appl. Math. Mech.*, 73(9):203–216, 1993.
- [23] P.M. Gresho and R.L. Sani. On pressure boundary conditions for the incompressible Navier-Stokes equations. *Int. J. Numer. Meth. Fluids*, 7(10):1111–1145, 1987.
- [24] J.G. Heywood, R. Rannacher, and S. Turek. Artificial boundaries and flux and pressure conditions for the incompressible Navier-Stokes equations. *Int. J. Numer. Meth. Fluids*, 22(5):325–352, 1996.
- [25] T. Richter. *Fluid-Structure Interactions: Models, Analysis and Finite Elements*. Springer International Publishing, Basel, 2017.
- [26] D.R.Q. Pacheco, T.S. Müller, O. Steinbach, and G. Brenn. On outflow boundary conditions in finite element simulations of non-Newtonian internal flows. *Int. J. Comput. Vis. Sci. Eng.*, 1(6), 2021.
- [27] V. John. *Finite Element Methods for Incompressible Flow Problems*. Springer International Publishing, Basel, 2016.

- [28] D. Boffi, F. Brezzi, and M. Fortin. *Mixed Finite Element Methods and Applications*, volume 44. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [29] P. Hood and C. Taylor. Navier-Stokes equations using mixed interpolation. In J.T. Oden, O.C. Zienkiewicz, R.H. Gallagher, and C. Taylor, editors, *Finite Element Methods in Flow Problems*, pages 57–66, Huntsville, 1974. UAH Press.
- [30] J.L. Guermond, P. Mineev, and Jie Shen. An overview of projection methods for incompressible flows. *Comput. Methods Appl. Mech. Eng.*, 195(44-47):6011–6045, 2006.
- [31] F. Brezzi and J. Pitkäranta. On the stabilization of finite element approximations of the Stokes equations. In W. Hackbusch, editor, *Efficient Solutions of Elliptic Systems. Notes on Numerical Fluid Mechanics*, volume 10, pages 11–19, Wiesbaden, 1984.
- [32] T.J.R. Hughes, L.P. Franca, and M. Balestra. A new finite element formulation for computational fluid dynamics: V. Circumventing the Babuška-Brezzi condition: a stable Petrov-Galerkin formulation of the Stokes problem accommodating equal-order interpolations. *Comput. Methods Appl. Mech. Eng.*, 59(1):85–99, 1986.
- [33] T.J.R. Hughes and L.P. Franca. A new finite element formulation for computational fluid dynamics: VII. The Stokes problem with various well-posed boundary conditions: Symmetric formulations that converge for all velocity/pressure spaces. *Comput. Methods Appl. Mech. Eng.*, 65(1):85–96, 1987.
- [34] F. Brezzi and J. Douglas Jr. Stabilized mixed methods for the Stokes problem. *Numer. Math.*, 53(1-2):225–235, 1988.
- [35] J. Douglas Jr. and J. Wang. An Absolutely Stabilized Finite Element Method for the Stokes Problem. *Math. Comput.*, 52(186):495, 1989.
- [36] F. Shakib. *Finite element analysis of the compressible Euler and Navier-Stokes equations*. PhD thesis, Stanford University, 1988.
- [37] T.E. Tezduyar. Stabilized Finite Element Formulations for Incompressible Flow Computations. In J.W. Hutchinson and T.Y. Wu, editors, *Stabilized Finite Element Formulations for Incompressible Flow Computations*, volume 28 of *Advances in Applied Mechanics*, pages 1–44. Elsevier, 1991.
- [38] B. Hübner and D. Dinkler. A simultaneous solution procedure for strong interactions of generalized Newtonian fluids and viscoelastic solids at large strains. *Int. J. Numer. Methods Eng.*, 64(7):920–939, 2005.
- [39] J. Donea and A. Huerta. *Finite Element Methods for Flow Problems*. John Wiley & Sons, Chichester, 2003.
- [40] V. John and J. Novo. Analysis of the Pressure Stabilized Petrov–Galerkin Method for the Evolutionary Stokes Equations Avoiding Time Step Restrictions. *SIAM J. Numer. Anal.*, 53(2):1005–1031, 2015.
- [41] P. Nithiarasu and C.-B. Liu. An artificial compressibility based characteristic based split (CBS) scheme for steady and unsteady turbulent incompressible flows. *Comput. Methods Appl. Mech. Eng.*, 195(23-24):2961–2982, 2006.

- [42] T.J.R. Hughes, W.K. Liu, and A. Brooks. Finite element analysis of incompressible viscous flows by the penalty function formulation. *J. Comput. Phys.*, 30(1):1–60, 1979.
- [43] L.P. Franca and A. Russo. Approximation of the Stokes problem by residual-free macro bubbles. *East-West J. Numer. Math.*, 4:265–278, 1996.
- [44] C.R. Dohrmann and P.B. Bochev. A stabilized finite element method for the Stokes problem based on polynomial pressure projections. *Int. J. Numer. Meth. Fluids*, 46(2):183–201, 2004.
- [45] J.-J. Droux and T.J.R. Hughes. A boundary integral modification of the Galerkin least squares formulation for the Stokes problem. *Comput. Methods Appl. Mech. Eng.*, 113(1-2):173–182, 1994.
- [46] K.E. Jansen, S.S. Collis, C. Whiting, and F. Shaki. A better consistency for low-order stabilized finite element methods. *Comput. Methods Appl. Mech. Eng.*, 174(1-2):153–170, 1999.
- [47] E. Burman and P. Hansbo. Edge stabilization for the generalized Stokes problem: A continuous interior penalty method. *Comput. Methods Appl. Mech. Eng.*, 195(19-22):2393–2410, 2006.
- [48] P. Bochev and M. Gunzburger. An absolutely stable pressure-Poisson stabilized finite element method for the Stokes equations. *SIAM J. Numer. Anal.*, 42(3):1189–1207, 2004.
- [49] R. Codina and J. Blasco. A finite element formulation for the Stokes problem allowing equal velocity-pressure interpolation. *Comput. Methods Appl. Mech. Eng.*, 143(3-4):373–391, 1997.
- [50] R. Codina, J. Blasco, G.C. Buscaglia, and A. Huerta. Implementation of a stabilized finite element formulation for the incompressible Navier-Stokes equations based on a pressure gradient projection. *Int. J. Numer. Meth. Fluids*, 37(4):419–444, 2001.
- [51] R. Becker and M. Braack. A finite element pressure gradient stabilization for the Stokes equations based on local projections. *Calcolo*, 38(4):173–199, 2001.
- [52] M. Braack and G. Lube. Finite elements with local projection stabilization for incompressible flow problems. *J. Comput. Math.*, 27:116–147, 2009.
- [53] L. Rubart and G. Böhme. Numerical simulation of shear-thinning flow problems in mixing vessels. *Theor. Comput. Fluid Dyn.*, 3(2):95–115, 1991.
- [54] H.R. Tamaddon-Jahromi, D. Ding, M.F. Webster, and P. Townsend. A Taylor–Galerkin finite element method for non-Newtonian flows. *Int. J. Numer. Methods Eng.*, 34(3):741–757, 1992.
- [55] J.P.W. Baaijens, A.A. van Steenhoven, and J.D. Janssen. Numerical analysis of steady generalized Newtonian blood flow in a 2D model of the carotid artery bifurcation. *Biorheology*, 60(1):63–74, 1993.
- [56] C. Tu and M. Deville. Pulsatile flow of non-Newtonian fluids through arterial stenoses. *J. Biomech.*, 29(7):899–908, 1996.

- [57] N. Arada, M. Pires, and A. Sequeira. Viscosity effects on flows of generalized Newtonian fluids through curved pipes. *Comput. Math. Appl.*, 53(3-4):625–646, 2007.
- [58] W.L. Barth, L.V. Branets, and G.F. Carey. Non-Newtonian flow in branched pipes and artery models. *Int. J. Numer. Meth. Fluids*, 57(5):531–553, 2008.
- [59] M.G.H.M. Baltussen, Y.J. Choi, M.A. Hulsen, and P.D. Anderson. Weakly-imposed Dirichlet boundary conditions for non-Newtonian fluid flow. *J. Non-Newton. Fluid Mech.*, 166(17-18):993–1003, 2011.
- [60] H. Lee. Numerical approximation of quasi-Newtonian flows by ALE-FEM. *Numer. Methods Partial Differ. Equ.*, 28(5):1667–1695, 2012.
- [61] H. Damanik, J. Hron, A. Ouazzi, and S. Turek. Monolithic Newton-multigrid solution techniques for incompressible nonlinear flow models. *Int. J. Numer. Meth. Fluids*, 71(2):208–222, 2013.
- [62] X. Zheng, G. Chen, and X. Xie. A divergence-free weak Galerkin method for quasi-Newtonian Stokes flows. *Sci. China Math.*, 60(8):1515–1528, 2017.
- [63] G. Böhme and L. Rubart. Non-Newtonian flow analysis by finite elements. *Fluid Dyn. Res.*, 5(3):147–158, 1989.
- [64] G.F. Carey, K.C. Wang, and W.D. Joubert. Performance of iterative methods for Newtonian and generalized Newtonian flows. *Int. J. Numer. Meth. Fluids*, 9(2):127–150, 1989.
- [65] M. Franta, J. Málek, and K.R. Rajagopal. On steady flows of fluids with pressure- and shear-dependent viscosities. *Proc. Math. Phys. Eng. Sci.*, 461(2055):651–670, 2005.
- [66] H. Sobhani, M. Razavi-Nouri, and M.H.R. Ghoreishy. Investigation of combination of finite element formulation and element type on the accuracy of 3D modeling of polymeric fluid flow in an extrusion die. *J. Appl. Polym. Sci.*, 120(3):1607–1615, 2011.
- [67] S. Knauf, S. Frei, T. Richter, and R. Rannacher. Towards a complete numerical description of lubricant film dynamics in ball bearings. *Comput. Mech.*, 53(2):239–255, 2014.
- [68] A. Masud and J. Kwack. A stabilized mixed finite element method for the incompressible shear-rate dependent non-Newtonian fluids: Variational Multiscale framework and consistent linearization. *Comput. Methods Appl. Mech. Eng.*, 200(5-8):577–596, 2011.
- [69] V.L. Marrero, J.A. Tichy, O. Sahni, and K.E. Jansen. Numerical study of purely viscous non-Newtonian flow in an abdominal aortic aneurysm. *J. Biomech. Eng.*, 136(10), 2014.
- [70] L. Gesenhues, J.J. Camata, A.M.A. Côrtes, F.A. Rochinha, and A.L.G.A. Coutinho. Finite element simulation of complex dense granular flows using a well-posed regularization of the $\mu(I)$ -rheology. *Comput. Fluids*, 188:102–113, 2019.

- [71] F. Zinani and S. Frey. Finite element approximations for quasi-Newtonian flows employing a multi-field GLS method. *Comput. Mech.*, 48(2):139–152, 2011.
- [72] E. Castillo and R. Codina. Stabilized stress–velocity–pressure finite element formulations of the Navier–Stokes problem for fluids with non-linear viscosity. *Comput. Methods Appl. Mech. Eng.*, 279:554–578, 2014.
- [73] A.N. Brooks and T.J.R. Hughes. Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations. *Comput. Methods Appl. Mech. Eng.*, 32(1):199–259, 1982.
- [74] L.P. Franca and S.L. Frey. Stabilized finite element methods: II. The incompressible Navier-Stokes equations. *Comput. Methods Appl. Mech. Eng.*, 99(2-3):209–233, 1992.
- [75] R. Hartmann. Higher-order and adaptive discontinuous Galerkin methods with shock-capturing applied to transonic turbulent delta wing flow. *Int. J. Numer. Meth. Fluids*, 72(8):883–894, 2013.
- [76] L.P. Franca and T.J.R. Hughes. Two classes of mixed finite element methods. *Comput. Methods Appl. Mech. Eng.*, 69(1):89–129, 1988.
- [77] M. Olshanskii, G. Lube, T. Heister, and J. Löwe. Grad-div stabilization and subgrid pressure models for the incompressible Navier-Stokes equations. *Comput. Methods Appl. Mech. Eng.*, 198(49-52):3975–3988, 2009.
- [78] L.P. Franca, T.J.R. Hughes, and R. Stenberg. Stabilized Finite Element Methods. In M.D. Gunzburger and R.A. Nicolaides, editors, *Incompressible Computational Fluid Dynamics: Trends and Advances*, pages 87–108. Cambridge University Press, 1993.
- [79] L.P. Franca, G. Hauke, and A. Masud. Revisiting stabilized finite element methods for the advective–diffusive equation. *Comput. Methods Appl. Mech. Eng.*, 195(13-16):1560–1572, 2006.
- [80] V. John, P. Knobloch, and U. Wilbrandt. Finite Element Pressure Stabilizations for Incompressible Flow Problems. In T. Bodnár, P. Galdi, and Š. Nečasová, editors, *Fluids Under Pressure*, pages 483–573. Springer International Publishing, Cham, 2020.
- [81] V. John, P. Knobloch, and J. Novo. Finite elements for scalar convection-dominated equations and incompressible flow problems: a never ending story? *Comput. Vis. Sci.*, 19(5-6):47–63, 2018.
- [82] J. Deteix and D. Yakoubi. Shear rate projection schemes for non-Newtonian fluids. *Comput. Methods Appl. Mech. Eng.*, 354:620–636, 2019.
- [83] J. Rang. Pressure corrected implicit θ -schemes for the incompressible Navier-Stokes equations. *Appl. Math. Comput.*, 201(1-2):747–761, 2008.
- [84] D. Meidner and T. Richter. A posteriori error estimation for the fractional step theta discretization of the incompressible Navier-Stokes equations. *Comput. Methods Appl. Mech. Eng.*, 288:45–59, 2015.

- [85] L. Failer and T. Wick. Adaptive time-step control for nonlinear fluid–structure interaction. *J. Comput. Phys.*, 366:448–477, 2018.
- [86] V. John. On the efficiency of linearization schemes and coupled multigrid methods in the simulation of a 3D flow around a cylinder. *Int. J. Numer. Meth. Fluids*, 50(7):845–862, 2006.
- [87] B.M. Irons and R.C. Tuck. A version of the Aitken accelerator for computer iteration. *Int. J. Numer. Methods Eng.*, 1(3):275–277, 1969.
- [88] U. Küttler and W.A. Wall. Fixed-point fluid-structure interaction solvers with dynamic relaxation. *Comput. Mech.*, 43(1):61–72, 2008.
- [89] V.M. Verzhbitskiĭ and I.F. Yumanova. On the quadratic convergence of the Aitken Δ^2 process. *Comput. Math. Math. Phys.*, 51(10):1659–1663, 2011.
- [90] E. Hairer, S.P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I - Nonstiff problems*, volume 8 of *Springer Series in Computational Mathematics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2nd edition, 1993.
- [91] J.C. Simo and F. Armero. Unconditional stability and long-term behavior of transient algorithms for the incompressible Navier-Stokes and Euler equations. *Comput. Methods Appl. Mech. Eng.*, 111(1-2):111–154, 1994.
- [92] S. Turek. A comparative study of time-stepping techniques for the incompressible Navier-Stokes equations: From fully implicit non-linear schemes to semi-implicit projection methods. *Int. J. Numer. Meth. Fluids*, 22:987–1011, 1996.
- [93] R. Ingram. A new linearly extrapolated Crank-Nicolson time-stepping scheme for the Navier-Stokes equations. *Math. Comput.*, 82(284):1953–1973, 2013.
- [94] H. Elman, M. Mihajlović, and D. Silvester. Fast iterative solvers for buoyancy driven flow problems. *J. Comput. Phys.*, 230(10):3900–3914, 2011.
- [95] V.E. Henson and U.M. Yang. BoomerAMG: A parallel algebraic multigrid solver and preconditioner. *Appl. Numer. Math.*, 41(1):155–177, 2002.
- [96] G. Alzetta, D. Arndt, W. Bangerth, V. Boddu, B. Brands, D. Davydov, R. Gassmoeller, T. Heister, L. Heltai, K. Kormann, M. Kronbichler, M. Maier, J.-P. Pelteret, B. Turcksin, and D. Wells. The `deal.II` Library, Version 9.0. *J. Numer. Math.*, 26(4):173–183, 2018.
- [97] D. Arndt, W. Bangerth, B. Blais, T.C. Clevenger, M. Fehling, A.V. Grayver, T. Heister, L. Heltai, M. Kronbichler, M. Maier, P. Munch, J.-P. Pelteret, R. Rastak, I. Thomas, B. Turcksin, Z. Wang, and D. Wells. The `deal.II` Library, Version 9.2. *J. Numer. Math.*, 28(3):131–146, 2020.
- [98] D. Arndt, W. Bangerth, B. Blais, M. Fehling, R. Gassmüller, T. Heister, L. Heltai, U. Köcher, M. Kronbichler, M. Maier, P. Munch, J.-P. Pelteret, S. Proell, K. Simon, B. Turcksin, D. Wells, and J. Zhang. The `deal.II` library, Version 9.3. *J. Numer. Math.*, 29(3):171–186, 2021.
- [99] D. Arndt, W. Bangerth, M. Feder, M. Fehling, R. Gassmüller, T. Heister, L. Heltai, M. Kronbichler, M. Maier, P. Munch, J.-P. Pelteret, S. Sticker, B. Turcksin, and D. Wells. The `deal.II` library, version 9.4. *J. Numer. Math.*, 30(3):231–246, 2022.

- [100] D. Braess. *Finite elements: Theory, fast solvers, and applications in solid mechanics*. Cambridge University Press, Cambridge, 3rd edition, 2007.
- [101] A. Smith and D. Silvester. Implicit algorithms and their linearization for the transient incompressible Navier-Stokes equations. *IMA J. Numer. Anal.*, 17(4): 527–545, 1997.
- [102] E. Cuthill and J. McKee. Reducing the Bandwidth of Sparse Symmetric Matrices. In *Proceedings of the 1969 24th National Conference*, ACM '69, pages 157–172, New York, NY, 1969. Association for Computing Machinery.
- [103] I.P. King. An automatic reordering scheme for simultaneous equations derived from network systems. *Int. J. Numer. Methods Eng.*, 2:523–533, 1970.
- [104] S. Turek. *Efficient Solvers for Incompressible Flow Problems - An Algorithmic and Computational Approach*, volume 6 of *Lecture Notes in Computational Science and Engineering*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.
- [105] M. Benzi. Preconditioning techniques for large linear systems: A survey. *J. Comput. Phys.*, 182:418–477, 2002.
- [106] M. Benzi, G.H. Golub, and J. Liesen. Numerical solution of saddle point problems. *Acta Numer.*, 14:1–137, 2005.
- [107] H.C. Elman, D.J. Silvester, and A.J. Wathen. *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*. Oxford University Press, Oxford, 2014.
- [108] U. Langer and H. Yang. Algebraic Multigrid Based Preconditioners for Fluid-Structure Interaction and Its Related Sub-problems. In I. Lirkov, S.D. Margenov, and J. Waśniewski, editors, *Large-Scale Scientific Computing*, pages 91–98, Cham, 2015. Springer International Publishing.
- [109] O. Axelsson. *Iterative Solution Methods*. Cambridge University Press, Cambridge, 1994.
- [110] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Panama, 2003.
- [111] M.A. Heroux and J.M. Willenbring. A new overview of the Trilinos project. *Sci. Program.*, 20(2):83–88, 2012.
- [112] O. Axelsson and R. Blaheta. Preconditioning of matrices partitioned in 2×2 block form: eigenvalue estimates and Schwarz DD for mixed FEM. *Numer. Linear Algebra Appl.*, 17(5):787–810, 2010.
- [113] O. Axelsson and M. Neytcheva. A general approach to analyse preconditioners for two-by-two block matrices. *Numer. Linear Algebra Appl.*, 20(5):723–742, 2013.
- [114] Z.-Z. Bai. Structured preconditioners for nonsingular matrices of block two-by-two structures. *Math. Comput.*, 75(254):791–815, 2005.
- [115] Z.-Z. Bai and M.K. Ng. On Inexact Preconditioners for Nonsymmetric Matrices. *SIAM J. Sci. Comput.*, 26(5):1710–1724, 2005.

- [116] M.F. Murphy, G.H. Golub, and A.J. Wathen. A Note on Preconditioning for Indefinite Linear Systems. *SIAM J. Sci. Comput.*, 21(6):1969–1972, 2000.
- [117] Y. Notay. A New Analysis of Block Preconditioners for Saddle Point Problems. *SIAM J. Matrix Anal. Appl.*, 35(1):143–173, 2014.
- [118] A. Segal. Preconditioners for Incompressible Navier-Stokes Solvers. *Numer. Math.*, 3(3):245–275, 2010.
- [119] A.C. de Niet and F.W. Wubs. Two preconditioners for saddle point problems in fluid flows. *Int. J. Numer. Meth. Fluids*, 54(4):355–377, 2007.
- [120] S.V. Patankar and D.B. Spalding. A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows. *Int. J. Heat Mass Transfer*, 15(10):1787–1806, 1972.
- [121] C. Vuik, A. Saghier, and G.P. Boerstoel. The Krylov accelerated SIMPLE(R) method for flow problems in industrial furnaces. *Int. J. Numer. Meth. Fluids*, 33(7):1027–1040, 2000.
- [122] M. ur Rehman, C. Vuik, and G. Segal. SIMPLE-type preconditioners for the Oseen problem. *Int. J. Numer. Meth. Fluids*, 61(4):432–452, 2009.
- [123] D. Silvester, H. Elman, D. Kay, and A. Wathen. Efficient preconditioning of the linearized Navier–Stokes equations for incompressible flow. *J. Comput. Appl. Math.*, 128(1-2):261–279, 2001.
- [124] D. Kay, D. Loghin, and A. Wathen. A preconditioner for the steady-state Navier–Stokes equations. *SIAM J. Sci. Comput.*, 24(1):237–256, 2003.
- [125] H. Elman, V.E. Howle, J. Shadid, R. Shuttleworth, and R. Tuminaro. Block preconditioners based on approximate commutators. *SIAM J. Sci. Comput.*, 27(5):1651–1668, 2006.
- [126] H. Elman, V.E. Howle, J. Shadid, D. Silvester, and R. Tuminaro. Least squares preconditioners for stabilized discretizations of the Navier–Stokes equations. *SIAM J. Sci. Comput.*, 30(1):290–311, 2008.
- [127] M. Benzi and M.A. Olshanskii. An augmented Lagrangian-based approach to the Oseen problem. *SIAM J. Sci. Comput.*, 28(6):2095–2113, 2006.
- [128] M. Benzi, M.A. Olshanskii, and Z. Wang. Modified augmented Lagrangian preconditioners for the incompressible Navier-Stokes equations. *Int. J. Numer. Meth. Fluids*, 66(4):486–508, 2011.
- [129] T. Heister and G. Rapin. Efficient augmented Lagrangian-type preconditioning for the Oseen problem using Grad-Div stabilization. *Int. J. Numer. Meth. Fluids*, 71(1):118–134, 2013.
- [130] H.C. Elman. Preconditioning strategies for models of incompressible flow. *J. Sci. Comput.*, 25(1):347–366, 2005.
- [131] M. ur Rehman, C. Vuik, and G. Segal. A comparison of preconditioners for incompressible Navier-Stokes solvers. *Int. J. Numer. Meth. Fluids*, 57(12):1731–1751, 2008.

- [132] J. Cahouet and J.-P. Chabard. Some fast 3D finite element solvers for the generalized Stokes problem. *Int. J. Numer. Meth. Fluids*, 8(8):869–895, 1988.
- [133] X. He, M. Neytcheva, and C. Vuik. On Preconditioning of Incompressible Non-Newtonian Flow Problems. *J. Comput. Math.*, 33(1):33–58, 2015.
- [134] M. Kronbichler, A. Diagne, and H. Holmgren. A fast massively parallel two-phase flow solver for microfluidic chip simulation. *Int. J. High Perform. Comput. Appl.*, 32(2):266–287, 2018.
- [135] Y. Saad and M.H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Statist.*, 7(3):856–869, 1986.
- [136] H.C. Elman and R.S. Tuminaro. Boundary conditions in approximate commutator preconditioners for the Navier-Stokes equations. *Electron. Trans. Numer. Anal.*, 35:257–280, 2009.
- [137] O.C. Zienkiewicz and Y.M. Xie. A simple error estimator and adaptive time stepping procedure for dynamic analysis. *Earthq. Eng. Struct. Dyn.*, 20(9):871–887, 1991.
- [138] P.M. Gresho, D.F. Griffiths, and D.J. Silvester. Adaptive time-stepping for incompressible flow part I: scalar advection-diffusion. *SIAM J. Sci. Comput.*, 30(4): 2018–2054, 2008.
- [139] D.A. Kay, P.M. Gresho, D.F. Griffiths, and D.J. Silvester. Adaptive time-stepping for incompressible flow part II: Navier–Stokes equations. *SIAM J. Sci. Comput.*, 32(1):111–128, 2010.
- [140] V. John and J. Rang. Adaptive time step control for the incompressible Navier–Stokes equations. *Comput. Methods Appl. Mech. Eng.*, 199(9):514–524, 2010.
- [141] J.G. Heywood and R. Rannacher. Finite-element approximation of the nonstationary Navier–Stokes problem. Part IV: Error analysis for second-order time discretization. *SIAM J. Numer. Anal.*, 27(2):353–384, 1990.
- [142] W. Dettmer and D. Perić. An analysis of the time integration algorithms for the finite element solutions of incompressible Navier–Stokes equations based on a stabilised formulation. *Comput. Methods Appl. Mech. Eng.*, 192(9-10):1177–1226, 2003.
- [143] M. Luskin, R. Rannacher, and W. Wendland. On the smoothing property of the Crank-Nicolson scheme. *Appl. Anal.*, 14(2):117–135, 1982.
- [144] R. Rannacher. Finite element solution of diffusion problems with irregular data. *Numer. Math.*, 43(2):309–327, 1984.
- [145] B. Lindberg. On smoothing and extrapolation for the trapezoidal rule. *BIT Numer. Math.*, 11(1):29–52, 1971.
- [146] O. Østerby. Five ways of reducing the Crank–Nicolson oscillations. *BIT Numer. Math.*, 43(4):811–822, 2003.
- [147] W.A. Wall. *Fluid-Struktur-Interaktion mit stabilisierten Finiten Elementen*. PhD thesis, Universität Stuttgart, 1999.

- [148] J.C. Butcher. Implicit Runge-Kutta processes. *Math. Comput.*, 18(85):50–64, 1964.
- [149] P. Solin and L. Korous. Adaptive higher-order finite element methods for transient PDE problems based on embedded higher-order implicit Runge-Kutta methods. *J. Comput. Phys.*, 231(4):1635–1649, 2012.
- [150] K. Gustafsson, M. Lundh, and G. Söderlind. A PI stepsize control for the numerical solution of ordinary differential equations. *BIT Numer. Math.*, 28(2):270–287, 1988.
- [151] M. Mayr, W.A. Wall, and M.W. Gee. Adaptive time stepping for fluid-structure interaction solvers. *Finite Elem. Anal. Des.*, 141:55–69, 2018.
- [152] P.R. Amestoy, I.S. Duff, J.Y. L'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.*, 23(1):15–41, 2001.
- [153] L.I.G. Kovasznay. Laminar flow behind a two-dimensional grid. *Math. Proc. Camb. Philos. Soc.*, 44(1):58–62, 1948.
- [154] J.-G. Liu, J. Liu, and R.L. Pego. Stable and accurate pressure approximation for unsteady incompressible viscous flow. *J. Comput. Phys.*, 229(9):3428–3453, 2010.
- [155] J. de Frutos, V. John, and J. Novo. Projection methods for incompressible flow problems with WENO finite difference schemes. *J. Comput. Phys.*, 309:368–386, 2016.
- [156] E. Erturk, T.C. Corke, and C. Gökçöl. Numerical solutions of 2-D steady incompressible driven cavity flow at high Reynolds numbers. *Int. J. Numer. Meth. Fluids*, 48(7):747–774, 2005.
- [157] U. Ghia, K.N. Ghia, and C.T. Shin. High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method. *J. Comput. Phys.*, 48(3):387–411, 1982.
- [158] V. John, G. Matthies, and J. Rang. A comparison of time-discretization/linearization approaches for the incompressible Navier-Stokes equations. *Comput. Methods Appl. Mech. Eng.*, 195(44-47):5995–6010, 2006.
- [159] L. Jay. Convergence of a class of Runge-Kutta methods for differential-algebraic systems of index 2. *BIT Numer. Math.*, 33(1):137–150, 1993.
- [160] E. Hairer and G. Wanner. *Solving Ordinary Differential Equations II - Stiff and Differential-Algebraic Problems*, volume 14 of *Springer Series in Computational Mathematics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2nd edition, 1996.
- [161] O. Steinbach. A note on initial higher-order convergence results for boundary element methods with approximated boundary conditions. *Numer. Methods Partial Differ. Equ.*, 16(6):581–588, 2000.
- [162] D.R.Q. Pacheco and O. Steinbach. On the initial higher-order pressure convergence in equal-order finite element discretizations of the Stokes system. *Comput. Math. with Appl.*, 109:140–145, 2022.

- [163] Y.I. Cho and K.R. Kensey. Effects of the non-Newtonian viscosity of blood on flows in a diseased arterial vessel. Part 1: Steady flows. *Biorheology*, 28(3-4):241–262, 1991.
- [164] M. Thiriet. *Biology and Mechanics of Blood Flows - Part II: Mechanics and Medical Aspects*. CRM Series in Mathematical Physics. Springer New York, New York, 2008.
- [165] A.M. Gambaruto, J. Janela, A. Moura, and A. Sequeira. Sensitivity of hemodynamics in a patient specific cerebral aneurysm to vascular geometry and blood rheology. *Math. Biosci. Eng.*, 8(2):409–23, 2011.
- [166] G.E. Karniadakis and S.J. Sherwin. *Spectral/hp element methods for computational fluid dynamics*. Oxford University Press, Oxford, 2013.
- [167] J.L. Guermond, P. Mineev, and J. Shen. Error analysis of pressure-correction schemes for the time-dependent Stokes equations with open boundary conditions. *SIAM J. Numer. Anal.*, 43(1):239–258, 2005.
- [168] J.L. Guermond and J. Shen. Velocity-Correction Projection Methods for Incompressible Flows. *SIAM J. Numer. Anal.*, 41(1):112–134, 2003.
- [169] G.E. Karniadakis, M. Israeli, and S.A. Orszag. High-order splitting methods for the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 97(2):414–443, 1991.
- [170] S.A. Orszag, M. Israeli, and M.O. Deville. Boundary conditions for incompressible flows. *J. Sci. Comput.*, 1(1):75–111, 1986.
- [171] A.J. Chorin. Numerical solution of the Navier-Stokes equations. *Math. Comput.*, 22(104):745–762, 1968.
- [172] C.W Hirt and J.L Cook. Calculating three-dimensional flows around structures and over rough terrain. *J. Comput. Phys.*, 10(2):324–340, 1972.
- [173] J. Guermond and J. Shen. On the error estimates for the rotational pressure-correction projection methods. *Math. Comput.*, 73(248):1719–1737, 2003.
- [174] P.D. Mineev and P.M. Gresho. A remark on pressure correction schemes for transient viscous incompressible flow. *Commun. Numer. Methods Eng.*, 14(4):335–346, 1998.
- [175] L.J.P. Timmermans, P.D. Mineev, and F.N. van de Vosse. An approximate projection scheme for incompressible flow using spectral elements. *Int. J. Numer. Meth. Fluids*, 22(7):673–688, 1996.
- [176] A. Poux, S. Glockner, and M. Azaïez. Improvements on open and traction boundary conditions for Navier-Stokes time-splitting methods. *J. Comput. Phys.*, 230(10):4011–4027, 2011.
- [177] N. Fehn, W.A. Wall, and M. Kronbichler. On the stability of projection methods for the incompressible Navier-Stokes equations based on high-order discontinuous Galerkin discretizations. *J. Comput. Phys.*, 351:392–421, 2017.
- [178] J.-G. Liu, J. Liu, and R.L. Pego. Error estimates for finite-element Navier-Stokes solvers without standard Inf-Sup conditions. *Chin. Ann. Math. Ser. B*, 30(6):743–768, 2009.

- [179] J. Jia and J. Liu. Stable and spectrally accurate schemes for the Navier–Stokes equations. *SIAM J. Sci. Comput.*, 33(5):2421–2439, 2011.
- [180] D. Shirokoff and R.R. Rosales. An efficient method for the incompressible Navier–Stokes equations on irregular domains with no-slip boundary conditions, high order up to the boundary. *J. Comput. Phys.*, 230(23):8619–8646, 2011.
- [181] Z. Sheng, M. Thiriet, and F. Hecht. A high-order scheme for the incompressible Navier-Stokes equations with open boundary condition. *Int. J. Numer. Meth. Fluids*, 73(1):58–73, 2013.
- [182] H. Johnston, C. Wang, and J.-G. Liu. A local pressure boundary condition spectral collocation scheme for the three-dimensional Navier–Stokes equations. *J. Sci. Comput.*, 60(3):612–626, 2014.
- [183] L. Li. A split-step finite-element method for incompressible Navier-Stokes equations with high-order accuracy up-to the boundary. *J. Comput. Phys.*, 408:109274, 2020.
- [184] J. Deteix and D. Yakoubi. Improving the pressure accuracy in a projection scheme for incompressible fluids with variable viscosity. *Appl. Math. Lett.*, 79:111–117, 2018.
- [185] L. Plasman, J. Deteix, and D. Yakoubi. A projection scheme for Navier-Stokes with variable viscosity and natural boundary condition. *Int. J. Numer. Meth. Fluids*, pages 1–21, 2020.
- [186] A. Limache, S. Idelsohn, R. Rossi, and E. Oñate. The violation of objectivity in Laplace formulations of the Navier–Stokes equations. *Int. J. Numer. Meth. Fluids*, 54(6-8):639–664, 2007.
- [187] M. Heil. An efficient solver for the fully coupled solution of large-displacement fluid–structure interaction problems. *Comput. Methods Appl. Mech. Eng.*, 193(1-2):1–23, 2004.
- [188] T.J.R. Hughes, W.K. Liu, and T.K. Zimmermann. Lagrangian-Eulerian finite element formulation for incompressible viscous flows. *Comput. Methods Appl. Mech. Eng.*, 29(3):329–349, 1981.
- [189] P. Le Tallec and J. Mouro. Fluid structure interaction with large structural displacements. *Comput. Methods Appl. Mech. Eng.*, 190(24-25):3039–3067, 2001.
- [190] A. Leuprecht, K. Perktold, M. Prosi, T. Berk, W. Trubel, and H. Schima. Numerical study of hemodynamics and wall mechanics in distal end-to-side anastomoses of bypass grafts. *J. Biomech.*, 35(2):225–236, 2002.
- [191] D. Forti, M. Bukač, A. Quaini, S. Čanić, and S. Deparis. A monolithic approach to fluid–composite structure interaction. *J. Sci. Comput.*, 72(1):396–421, 2017.
- [192] A. Quaini and A. Quarteroni. A semi-implicit approach for fluid-structure interaction based on an algebraic fractional step method. *Math. Models Methods Appl. Sci.*, 17(06):957–983, 2007.
- [193] A. Quarteroni, M. Tuveri, and A. Veneziani. Computational vascular fluid dynamics: problems, models and methods. *Comput. Vis. Sci.*, 2(4):163–197, 2000.

- [194] J. Donea, S. Giuliani, and J.P. Halleux. An arbitrary Lagrangian-Eulerian finite element method for transient dynamic fluid-structure interactions. *Comput. Methods Appl. Mech. Eng.*, 33(1-3):689–723, 1982.
- [195] L.J. Fauci and R. Dillon. Biofluidmechanics of reproduction. *Annu. Rev. Fluid Mech.*, 38(1):371–394, 2006.
- [196] A.L. Fogelson. Platelet-wall interactions in continuum models of platelet thrombosis: formulation and numerical solution. *Math. Med. Biol.*, 21(4):293–334, 2004.
- [197] B.E. Griffith. Immersed boundary model of aortic heart valve dynamics with physiological driving and loading conditions. *Int. J. Numer. Method. Biomed. Eng.*, 28(3):317–345, 2012.
- [198] B.E. Griffith, X. Luo, D.M. McQueen, and C.S. Peskin. Simulating the fluid dynamics of natural and prosthetic heart valves using the immersed boundary method. *Int. J. Appl. Mech.*, 01(01):137–177, 2009.
- [199] J.D. Bransen, A. Viré, S.R. Turteltaub, and G.J.W. van Bussel. A comparative analysis of Lagrange multiplier and penalty approaches for modelling fluid-structure interaction. *Eng. Comput.*, 38(4):1677–1705, 2021.
- [200] C. Hesch, A.J. Gil, A. Arranz Carreño, J. Bonet, and P. Betsch. A mortar approach for fluid–structure interaction problems: Immersed strategies for deformable and rigid bodies. *Comput. Methods Appl. Mech. Eng.*, 278:853–882, 2014.
- [201] F.P.T. Baaijens. A fictitious domain/mortar element method for fluid-structure interaction. *Int. J. Numer. Methods Fluids*, 35(7):743–761, 2001.
- [202] R. van Loon, P.D. Anderson, J. de Hart, and F.P.T. Baaijens. A combined fictitious domain/adaptive meshing method for fluid–structure interaction in heart valves. *Int. J. Numer. Methods Fluids*, 46(5):533–544, 2004.
- [203] D. Boffi and L. Gastaldi. A fictitious domain approach with Lagrange multiplier for fluid-structure interactions. *Numer. Math.*, 135(3):711–732, 2017.
- [204] Y. Wang, P.K. Jimack, and M.A. Walkley. A one-field monolithic fictitious domain method for fluid–structure interactions. *Comput. Methods Appl. Mech. Eng.*, 317: 1146–1168, 2017.
- [205] C.W. Hirt, A.A. Amsden, and J.L. Cook. An arbitrary Lagrangian-Eulerian computing method for all flow speeds. *J. Comput. Phys.*, 14(3):227–253, 1974.
- [206] J. Donea, P. Fasoli-Stella, and S. Giuliani. Lagrangian and Eulerian finite element techniques for transient fluid-structure interaction problems. *Therm. and Fluid/Struct. Dyn. Anal.*, B, 1977.
- [207] J. Donea, A. Huerta, J.-P. Ponthot, and A. Rodríguez-Ferran. Arbitrary Lagrangian-Eulerian Methods. In *Encyclopedia of Computational Mechanics*, pages 1–23. John Wiley & Sons, Chichester, UK, 2nd edition, 2017.
- [208] P.D. Thomas and C.K. Lombard. Geometric conservation law and its application to flow computations on moving grids. *AIAA J.*, 17, 1979.

- [209] H. Guillard and C. Farhat. On the significance of the geometric conservation law for flow computations on moving meshes. *Comput. Methods Appl. Mech. Eng.*, 190(11-12):1467–1482, 2000.
- [210] C. Farhat, P. Geuzaine, and C. Grandmont. The Discrete Geometric Conservation Law and the Nonlinear Stability of ALE Schemes for the Solution of Flow Problems on Moving Grids. *J. Comput. Phys.*, 174(2):669–694, 2001.
- [211] C. Farhat and P. Geuzaine. Design and analysis of robust ALE time-integrators for the solution of unsteady flow problems on moving grids. *Comput. Methods Appl. Mech. Eng.*, 193(39):4073–4095, 2004.
- [212] D. Boffi and L. Gastaldi. Stability and geometric conservation laws for ALE formulations. *Comput. Methods Appl. Mech. Eng.*, 193(42-44):4717–4739, 2004.
- [213] Étienne, A. Garon, and D. Pelletier. Perspective on the geometric conservation law and finite element methods for ALE simulations of incompressible flow. *J. Comput. Phys.*, 228(7):2313–2333, 2009.
- [214] C. Förster, W.A. Wall, and E. Ramm. On the geometric conservation law in transient flow calculations on deforming domains. *Int. J. Numer. Methods Fluids*, 50(12):1369–1379, 2006.
- [215] N. Fehn, J. Heinz, W.A. Wall, and M. Kronbichler. High-order arbitrary Lagrangian-Eulerian discontinuous Galerkin methods for the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 430:110040, 2021.
- [216] P. Geuzaine, C. Grandmont, and C. Farhat. Design and analysis of ALE schemes with provable second-order time-accuracy for inviscid and viscous flow simulations. *J. Comput. Phys.*, 191(1):206–227, 2003.
- [217] S. Badia, F. Nobile, and C. Vergara. Robin–Robin preconditioned Krylov methods for fluid–structure interaction problems. *Comput. Methods Appl. Mech. Eng.*, 198(33-36):2768–2784, 2009.
- [218] S. Badia, F. Nobile, and C. Vergara. Fluid–structure partitioned procedures based on Robin transmission conditions. *J. Comput. Phys.*, 227(14):7027–7051, 2008.
- [219] F. Nobile and C. Vergara. An Effective Fluid-Structure Interaction Formulation for Vascular Dynamics by Generalized Robin Conditions. *SIAM J. Sci. Comput.*, 30(2):731–763, 2008.
- [220] L. Gerardo-Giorda, F. Nobile, and C. Vergara. Analysis and Optimization of Robin–Robin Partitioned Procedures in Fluid-Structure Interaction Problems. *SIAM J. Numer. Anal.*, 48(6):2091–2116, 2010.
- [221] J. Degroote. On the similarity between Dirichlet–Neumann with interface artificial compressibility and Robin–Neumann schemes for the solution of fluid-structure interaction problems. *J. Comput. Phys.*, 230(17):6399–6403, 2011.
- [222] E. Burman and M.A. Fernández. Explicit strategies for incompressible fluid-structure interaction problems: Nitsche type mortaring versus Robin-Robin coupling. *Int. J. Numer. Methods Eng.*, 97(10):739–758, 2014.

- [223] G. Gigante and C. Vergara. On the Choice of Interface Parameters in Robin–Robin Loosely Coupled Schemes for Fluid–Structure Interaction. *Fluids*, 6(6), 2021.
- [224] G. Gigante and C. Vergara. On the stability of a loosely-coupled scheme based on a Robin interface condition for fluid-structure interaction. *Comput. Math. with Appl.*, 96:109–119, 2021.
- [225] M. Astorino, F. Chouly, and M.A. Fernández. Robin Based Semi-Implicit Coupling in Fluid-Structure Interaction: Stability Analysis and Numerics. *SIAM J. Sci. Comput.*, 31(6):4041–4065, 2010.
- [226] T.-P. Fries and H.G. Matthies. A stabilized and coupled meshfree/meshbased method for the incompressible Navier–Stokes equations – Part II: Coupling. *Comput. Methods Appl. Mech. Eng.*, 195(44–47):6191–6204, 2006.
- [227] A. Porpora, P. Zunino, C. Vergara, and M. Piccinelli. Numerical treatment of boundary conditions to replace lateral branches in hemodynamics. *Int. J. Numer. Methods Biomed. Eng.*, 28(12):1165–1183, 2012.
- [228] K. Takizawa, J. Christopher, T.E. Tezduyar, and S. Sathe. Space-time finite element computation of arterial fluid-structure interactions with patient-specific data. *Int. J. Numer. Methods Biomed. Eng.*, 26(1):101–116, 2010.
- [229] J. Fouchet-Incaux. Artificial boundaries and formulations for the incompressible Navier-Stokes equations. Applications to air and blood flows. *Bol. Soc. Esp. Mat. Apl. SeMA*, 40:1–40, 2014.
- [230] L. Baffico, C. Grandmont, and B. Maury. Multiscale modelling of the respiratory tract. *Math. Models Methods Appl. Sci.*, 20(1):59–93, 2010.
- [231] M. Braack and P.B. Mucha. Directional Do-Nothing Condition for the Navier-Stokes Equations. *J. Comput. Math.*, 32(5):507–521, 2014.
- [232] F. Nobile and L. Formaggia. A Stability Analysis for the Arbitrary Lagrangian Eulerian Formulation with Finite Elements. *East-West J. Numer. Math.*, 7:105–132, 1999.
- [233] D. Arndt, M. Braack, and G. Lube. Finite Elements for the Navier-Stokes Problem with Outflow Condition. In B. Karasözen, M. Manguoğlu, M. Tezer-Sezgin, S. Göktepe, and Ö. Uğur, editors, *Numerical Mathematics and Advanced Applications ENUMATH 2015*, pages 95–103, Cham, 2016. Springer International Publishing.
- [234] C. Conca, F. Murat, and O. Pironneau. The Stokes and Navier-Stokes equations with boundary conditions involving the pressure. *Japan. J. Math. (N.S.)*, 20(2): 279–318, 1994.
- [235] C. Bègue, C. Conca, and O. Murat, F. Pironneau. A nouveau sur les équations de Stokes et de Navier-Stokes avec des conditions aux limites sur la pression. *CR Acad. Sc. Paris*, 304(1):23–28, 1987.
- [236] Ch.-H. Bruneau and P. Fabrie. Effective downstream boundary conditions for incompressible Navier–Stokes equations. *Int. J. Numer. Methods Fluids*, 19(8): 693–705, 1994.

- [237] C.-H. Bruneau and P. Fabrie. New efficient boundary conditions for incompressible Navier-Stokes equations : a well-posedness result. *Esaim Math. Model. Numer. Anal.*, 30(7):815–840, 1996.
- [238] C. Bertoglio and A. Caiazzo. A tangential regularization method for backflow stabilization in hemodynamics. *J. Comput. Phys.*, 261:162–171, 2014.
- [239] C. Bertoglio and A. Caiazzo. A Stokes-residual backflow stabilization method applied to physiological flows. *J. Comput. Phys.*, 313:260–278, 2016.
- [240] L. Formaggia, J.F. Gerbeau, F. Nobile, and A. Quarteroni. Numerical Treatment of Defective Boundary Conditions for the Navier–Stokes Equations. *SIAM J. Numer. Anal.*, 40(1):376–401, 2002.
- [241] G. Arbia, I.E. Vignon-Clementel, T.-Y. Hsia, and J.-F. Gerbeau. Modified Navier–Stokes equations for the outflow boundary conditions in hemodynamics. *Eur. J. Mech. B Fluids*, 60:175–188, 2016.
- [242] S.H. Attaran, H. Niroomand-oscuii, and F. Ghalichi. A novel, simple 3D/2D outflow boundary model for blood flow simulations in compliant arteries. *Comput. Fluids*, 174:229–240, 2018.
- [243] H.J. Kim, C.A. Figueroa, T.J.R. Hughes, K.E. Jansen, and C.A. Taylor. Augmented Lagrangian method for constraining the shape of velocity profiles at outlet boundaries for three-dimensional finite element simulations of blood flow. *Comput. Methods Appl. Mech. Eng.*, 198(45):3551–3566, 2009.
- [244] M. Esmaily Moghadam, Y. Bazilevs, T.-Y. Hsia, I.E. Vignon-Clementel, and A.L. Marsden. A comparison of outlet boundary treatments for prevention of backflow divergence with relevance to blood flow simulations. *Comput. Mech.*, 48(3):277–291, 2011.
- [245] C. Bertoglio, A. Caiazzo, Y. Bazilevs, M. Braack, M. Esmaily Moghadam, V. Gravemeier, A.L. Marsden, P. Olivier, I.E. Vignon-Clementel, and W.A. Wall. Benchmark problems for numerical treatment of backflow at open boundaries. *Int. J. Numer. Methods Biomed. Eng.*, 34(2):e2918, 2018.
- [246] Y. Bazilevs, J.R. Gohean, T.J.R. Hughes, R.D. Moser, and Y. Zhang. Patient-specific isogeometric fluid–structure interaction analysis of thoracic aortic blood flow due to implantation of the Jarvik 2000 left ventricular assist device. *Comput. Methods Appl. Mech. Eng.*, 198(45-46):3534–3550, 2009.
- [247] I.E. Vignon-Clementel, C.A. Figueroa, K.E. Jansen, and C.A. Taylor. Outflow boundary conditions for 3D simulations of non-periodic blood flow and pressure fields in deformable arteries. *Comput. Methods Biomech. Biomed. Eng.*, 13(5): 625–640, 2010.
- [248] J. Liu, W. Yang, I.S. Lan, and A.L. Marsden. Fluid-structure interaction modeling of blood flow in the pulmonary arteries using the unified continuum and variational multiscale formulation. *Mech. Res. Commun.*, 107:103556, 2020.
- [249] J.M. Oakes, A.L. Marsden, C. Grandmont, S.C. Shadden, C. Darquenne, and I.E. Vignon-Clementel. Airflow and Particle Deposition Simulations in Health and

- Emphysema: From In Vivo to In Silico Animal Experiments. *Ann. Biomed. Eng.*, 42:899–914, 2015.
- [250] J.M. Oakes, A.L. Marsden, C. Grandmont, C. Darquenne, and I.E. Vignon-Clementel. Distribution of aerosolized particles in healthy and emphysematous rat lungs: Comparison between experimental and numerical studies. *J. Biomech.*, 48(6):1147–1157, 2015.
- [251] O. Frank. Die Grundform des arteriellen Pulses: Mathematische Analyse. Erste Abhandlung. *Z. Biol.*, 37:483–526, 1899.
- [252] T.J.R. Hughes and J. Lubliner. On the one-dimensional theory of blood flow in the larger vessels. *Math. Biosci.*, 18(1-2):161–170, 1973.
- [253] N. Stergiopoulos, D.F. Young, and T.R. Rogge. Computer simulation of arterial flow with applications to arterial and aortic stenoses. *J. Biomech.*, 25(12):1477–1488, 1992.
- [254] B.W. Schaaf and P.H. Abbrecht. Digital computer simulation of human systemic arterial pulse wave transmission: A nonlinear model. *J. Biomech.*, 5(4):345–364, 1972.
- [255] J.C. Stettler, P. Niederer, and M. Anliker. Theoretical analysis of arterial hemodynamics including the influence of bifurcations - Part I: Mathematical model and prediction of normal pulse patterns. *Ann. Biomed. Eng.*, 9(2):145–164, 1981.
- [256] L. Formaggia, J.F. Gerbeau, F. Nobile, and A. Quarteroni. On the coupling of 3D and 1D Navier–Stokes equations for flow problems in compliant vessels. *Comput. Methods Appl. Mech. Eng.*, 191(6-7):561–582, 2001.
- [257] I.E. Vignon-Clementel, C.A. Figueroa, K.E. Jansen, and C.A. Taylor. Outflow boundary conditions for three-dimensional finite element modeling of blood flow and pressure in arteries. *Comput. Methods Appl. Mech. Eng.*, 195(29):3776–3796, 2006.
- [258] P.J. Blanco, S.M. Watanabe, and R.A. Feijóo. Identification of vascular territory resistances in one-dimensional hemodynamics simulations. *J. Biomech.*, 45(12):2066–2073, 2012.
- [259] M. Oshima, R. Torii, S. Tokuda, S. Yamada, and A. Koizumi. Patient-Specific Modeling and Multi-Scale Blood Simulation for Computational Hemodynamic Study on the Human Cerebrovascular System. *Curr. Pharm. Biotechnol.*, 13(11):2153–2165, 2012.
- [260] I.E. Vignon-Clementel, A.L. Marsden, and J.A. Feinstein. A primer on computational simulation in congenital heart disease for the clinician. *Prog. Pediatr. Cardiol.*, 30(1):3–13, 2010.
- [261] F. Migliavacca, R. Balossino, G. Pennati, G. Dubini, T.-Y. Hsia, M.R. de Leval, and E.L. Bove. Multiscale modelling in biofluidynamics: Application to reconstructive paediatric cardiac surgery. *J. Biomech.*, 39(6):1010–1020, 2006.
- [262] H.J. Kim, I.E. Vignon-Clementel, C.A. Figueroa, K.E. Jansen, and C.A. Taylor. Developing computational methods for three-dimensional finite element simulations of coronary blood flow. *Finite Elem. Anal. Des.*, 46(6):514–525, 2010.

- [263] M. Esmaily Moghadam, F. Migliavacca, I.E. Vignon-Clementel, T.-Y. Hsia, A.L. Marsden, and Modeling of Congenital Hearts Alliance (MOCHA) Investigators. Optimization of Shunt Placement for the Norwood Surgery Using Multi-Domain Modeling. *J. Biomech. Eng.*, 134(5), 2012.
- [264] A Baretta, C. Corsini, A.L. Marsden, I.E. Vignon-Clementel, T.-Y. Hsia, G. Dubini, F. Migliavacca, and G. Pennati. Respiratory effects on hemodynamics in patient-specific CFD models of the Fontan circulation under exercise conditions. *Eur. J. Mech. B Fluids*, 35:61–69, 2012.
- [265] P.J. Blanco and R.A. Feijóo. A dimensionally-heterogeneous closed-loop model for the cardiovascular system and its applications. *Med. Eng. Phys.*, 35(5):652–667, 2013.
- [266] K. Bäumlner, V. Vedula, A.M. Sailer, J. Seo, P. Chiu, G. Mistelbauer, F.P. Chan, M.P. Fischbein, A.L. Marsden, and D. Fleischmann. Fluid–structure interaction simulations of patient-specific aortic dissection. *Biomech. Model. Mechanobiol.*, 19(5):1607–1628, 2020.
- [267] C. Bertoglio, A. Caiazzo, and M.A. Fernández. Fractional-Step Schemes for the Coupling of Distributed and Lumped Models in Hemodynamics. *SIAM J. Sci. Comput.*, 35(3):B551–B575, 2013.
- [268] A.M. de Villiers, A.T. McBride, B.D. Reddy, T. Franz, and B.S. Spottiswoode. A validated patient-specific FSI model for vascular access in haemodialysis. *Biomech. Model. Mechanobiol.*, 17(2):479–497, 2018.
- [269] A. Prasad, L.K. To, M.L. Gorrepati, C.K. Zarins, and C.A. Figueroa. Computational Analysis of Stresses Acting on Intermodular Junctions in Thoracic Aortic Endografts. *J. Endovasc. Ther.*, 18(4):559–568, 2011.
- [270] W. Yang, I.E. Vignon-Clementel, G. Troianowski, V.M. Reddy, J.A. Feinstein, and A.L. Marsden. Hepatic blood flow distribution and performance in conventional and novel Y-graft Fontan geometries: A case series computational fluid dynamics study. *J. Thorac. Cardiovasc. Surg.*, 143(5):1086–1097, 2012.
- [271] M. Esmaily Moghadam, I.E. Vignon-Clementel, R. Figliola, and A.L. Marsden. A modular numerical method for implicit 0D/3D coupling in cardiovascular finite element simulations. *J. Comput. Phys.*, 244:63–79, 2013.
- [272] L. Grinberg and G.E. Karniadakis. Outflow Boundary Conditions for Arterial Networks with Multiple Outlets. *Ann. Biomed. Eng.*, 36:1496–1514, 2008.
- [273] L. Formaggia, A. Quarteroni, and A. Veneziani. *Cardiovascular Mathematics: Modeling and simulation of the circulatory system*, volume 1 of *Modeling, Simulation and Applications*. Springer New York, Milan, 2009.
- [274] T.J.R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice-Hall, Englewood Cliffs, NJ, 1987.
- [275] I. Babuška and M. Suri. Locking effects in the finite element approximation of elasticity problems. *Numer. Math.*, 62:439–463, 1992.

- [276] G.A. Holzapfel. *Nonlinear Solid Mechanics - A Continuum Approach for Engineering*. John Wiley & Sons, Chichester, 2000.
- [277] P. Wriggers. *Nonlinear Finite Element Methods*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [278] O.C. Zienkiewicz, R.L. Taylor, and D. Fox. *The Finite Element Method for Solid and Structural Mechanics*. Butterworth-Heinemann, Oxford, 7th edition, 2014.
- [279] O. Gültekin, H. Dal, and G.A. Holzapfel. On the quasi-incompressible finite element analysis of anisotropic hyperelastic materials. *Comput. Mech.*, 63:443–453, 2019.
- [280] T. Sussman and K.-J. Bathe. A finite element formulation for nonlinear incompressible elastic and inelastic analysis. *Comput. Struct.*, 26:357–409, 1987.
- [281] R. Stenberg and M. Suri. Mixed hp finite element methods for problems in elasticity and Stokes flow. *Numer. Math.*, 72:367–389, 1996.
- [282] M.E. Canga and E.B. Becker. An iterative technique for the finite element analysis of near-incompressible materials. *Comput. Methods Appl. Mech. Eng.*, 170:79–101, 1999.
- [283] O. Klaas, A. Maniatty, and M.S. Shephard. A stabilized mixed finite element method for finite elasticity. *Comput. Methods Appl. Mech. Eng.*, 180:65–79, 1999.
- [284] A.M. Maniatty, Y. Liu, O. Klaas, and M.S. Shephard. Higher order stabilized finite element method for hyperelastic finite deformation. *Comput. Methods Appl. Mech. Eng.*, 191:1491–1503, 2002.
- [285] A. Masud and K. Xia. A Stabilized Mixed Finite Element Method for Nearly Incompressible Elasticity. *J. Appl. Mech.*, 72:711–720, 2005.
- [286] R. Mahnen, I. Caylak, and G. Laschet. Two mixed finite element formulations with area bubble functions for tetrahedral elements. *Comput. Methods Appl. Mech. Eng.*, 197:1147–1165, 2008.
- [287] I. Caylak and R. Mahnen. Stabilization of mixed tetrahedral elements at large deformations. *Int. J. Numer. Methods Eng.*, 90:218–242, 2012.
- [288] A.J. Gil, C.H. Lee, J. Bonet, and M. Aguirre. A stabilised Petrov–Galerkin formulation for linear tetrahedral elements in compressible, nearly incompressible and truly incompressible fast dynamics. *Comput. Methods Appl. Mech. Eng.*, 276:659–690, 2014.
- [289] M. Cervera, M. Chiumenti, L. Benedetti, and R. Codina. Mixed stabilized finite element methods in nonlinear solid mechanics. Part III: Compressible and incompressible plasticity. *Comput. Methods Appl. Mech. Eng.*, 285:752–775, 2015.
- [290] S. Rossi, N. Abboud, and G. Scovazzi. Implicit finite incompressible elastodynamics with linear finite elements: A stabilized method in rate form. *Comput. Methods Appl. Mech. Eng.*, 311:208–249, 2016.
- [291] E. Karabelas, G. Haase, G. Plank, and C.M. Augustin. Versatile stabilized finite element formulations for nearly and fully incompressible solid mechanics. *Comput. Mech.*, 65:193–215, 2020.

- [292] J.C. Simo and M.S. Rifai. A class of mixed assumed strain methods and the method of incompatible modes. *Int. J. Numer. Methods Eng.*, 29:1595–1638, 1990.
- [293] J.C. Simo and F. Armero. Geometrically non-linear enhanced strain mixed methods and the method of incompatible modes. *Int. J. Numer. Methods Eng.*, 33:1413–1449, 1992.
- [294] J.C. Simo, F. Armero, and R.L. Taylor. Improved versions of assumed enhanced strain tri-linear elements for 3D finite deformation problems. *Comput. Methods Appl. Mech. Eng.*, 110:359–386, 1993.
- [295] R.L. Taylor. A mixed-enhanced formulation tetrahedral finite elements. *Int. J. Numer. Methods Eng.*, 47:205–227, 2000.
- [296] J.-S. Chen and C. Pan. A Pressure Projection Method for Nearly Incompressible Rubber Hyperelasticity, Part I: Theory. *J. Appl. Mech.*, 63:862–868, 1996.
- [297] J.-S. Chen, C.-T. Wu, and C. Pan. A Pressure Projection Method for Nearly Incompressible Rubber Hyperelasticity, Part II: Applications. *J. Appl. Mech.*, 63:869–876, 1996.
- [298] J.C. Simo, R.L. Taylor, and K.S. Pister. Variational and projection methods for the volume constraint in finite deformation elasto-plasticity. *Comput. Methods Appl. Mech. Eng.*, 51(1):177–208, 1985.
- [299] O.C. Zienkiewicz, J. Rojek, R.L. Taylor, and M. Pastor. Triangles and tetrahedra in explicit dynamic codes for solids. *Int. J. Numer. Methods Eng.*, 43:565–583, 1998.
- [300] Y. Yu, H. Baek, M.L. Bittencourt, and G.E. Karniadakis. Mixed spectral/hp element formulation for nonlinear elasticity. *Comput. Methods Appl. Mech. Eng.*, 213-216:42–57, 2012.
- [301] M. Cervera, M. Chiumenti, Q. Valverde, and C. Agelet de Saracibar. Mixed linear/linear simplicial elements for incompressible elasticity and plasticity. *Comput. Methods Appl. Mech. Eng.*, 192:5249–5263, 2003.
- [302] N.M. Lafontaine, R. Rossi, M. Cervera, and M. Chiumenti. Explicit mixed strain-displacement finite element for dynamic geometrically non-linear solid mechanics. *Comput. Mech.*, 55:543–559, 2015.
- [303] M. Chiumenti, M. Cervera, and R. Codina. A mixed three-field FE formulation for stress accurate analysis including the incompressible limit. *Comp. Methods Appl. Mech. Eng.*, 283:1095–1116, 2015.
- [304] J. Schröder, P. Wriggers, and D. Balzani. A new mixed finite element based on different approximations of the minors of deformation tensors. *Comput. Methods Appl. Mech. Eng.*, 200:3583–3600, 2011.
- [305] J. Bonet, A.J. Gil, and R. Ortigosa. A computational framework for polyconvex large strain elasticity. *Comput. Methods Appl. Mech. Eng.*, 283:1061–1094, 2015.
- [306] J. Schröder, N. Viebahn, D. Balzani, and P. Wriggers. A novel mixed finite element for finite anisotropic elasticity; the SKA-element Simplified Kinematics for Anisotropy. *Comput. Methods Appl. Mech. Eng.*, 310:475–494, 2016.

- [307] H. Kabaria, A.J. Lew, and B. Cockburn. A hybridizable discontinuous Galerkin formulation for non-linear elasticity. *Comput. Methods Appl. Mech. Eng.*, 283: 303–329, 2015.
- [308] A. Ten Eyck and A. Lew. Discontinuous Galerkin methods for non-linear elasticity. *Int. J. Numer. Methods Eng.*, 67:1204–1243, 2006.
- [309] D. Di Pietro and S. Lemaire. An extension of the Crouzeix–Raviart space to general meshes with application to quasi-incompressible linear elasticity and Stokes flow. *Math. Comput.*, 84:1–31, 2014.
- [310] R.S. Falk. Nonconforming finite element methods for the equations of linear elasticity. *Math. Comput.*, 57:529–550, 1991.
- [311] D. Davydov, J.-P. Pelteret, D. Arndt, M. Kronbichler, and P. Steinmann. A matrix-free approach for finite-strain hyperelastic problems using geometric multigrid. *Int. J. Numer. Methods Eng.*, 121:2874–2895, 2020.
- [312] M. Suri. Analytical and computational assessment of locking in the hp finite element method. *Comput. Methods Appl. Mech. Eng.*, 133:347–371, 1996.
- [313] J.C. Simo and T.J.R. Hughes. *Computational Inelasticity*, volume 7 of *Interdisciplinary Applied Mathematics*. Springer New York, New York, 1998.
- [314] A. Düster, S. Hartmann, and E. Rank. p-FEM applied to finite isotropic hyperelastic bodies. *Comput. Methods Appl. Mech. Eng.*, 192:5147–5166, 2003.
- [315] Z. Yosibash, S. Hartmann, U. Heisserer, A. Düster, E. Rank, and M. Szanto. Axisymmetric pressure boundary loading for finite deformation analysis using p-FEM. *Comput. Methods Appl. Mech. Eng.*, 196, 2007.
- [316] C.M. Augustin, A. Neic, M. Liebmam, A.J. Prassl, S.A. Niederer, G. Haase, and G. Plank. Anatomically accurate high resolution modeling of human whole heart electromechanics: A strongly scalable algebraic multigrid solver method for non-linear deformation. *J. Comput. Phys.*, 305:622–646, 2016.
- [317] E. Karabelas, M.A.F. Gsell, C.M. Augustin, L. Marx, A. Neic, A.J. Prassl, L. Goubergrits, T. Kuehne, and G. Plank. Towards a Computational Framework for Modeling the Impact of Aortic Coarctations Upon Left Ventricular Load. *Frontiers in Physiology*, 9, 2018.
- [318] M.C. Hsu and Y. Bazilevs. Blood vessel tissue prestress modeling for vascular fluid–structure interaction simulation. *Finite Elem. Anal. Des.*, 47(6):593–599, 2011.
- [319] J. Bonet and R.D. Wood. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge University Press, Cambridge, 2nd edition, 2008.
- [320] P.G. Ciarlet. *Mathematical Elasticity Volume I: Three-Dimensional Elasticity*, volume 20 of *Studies in Mathematics and its Applications*. North-Holland, Amsterdam, 1988.
- [321] M.E. Gurtin. *An Introduction to Continuum Mechanics*, volume 158 of *Mathematics in Science and Engineering*. Academic Press, New York, 1981.

- [322] C. Truesdell and W. Noll. *The Non-Linear Field Theories of Mechanics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2nd edition, 1992.
- [323] P.J. Flory. Thermodynamic relations for high elastic materials. *Trans. Faraday Soc.*, 57:829–838, 1961.
- [324] L.R.G. Treloar. The elasticity of a network of long-chain molecules. I. *Trans. Faraday Soc.*, 39:36–41, 1943.
- [325] M. Mooney. A Theory of Large Elastic Deformation. *J. Appl. Phys.*, 11:582–592, 1940.
- [326] R.S. Rivlin. Large elastic deformations of isotropic materials IV. Further developments of the general theory. *Philos. Trans. Royal Soc. A*, 241:379–397, 1948.
- [327] J.C. Simo and C. Miehe. Associative coupled thermoplasticity at finite strains: Formulation, numerical analysis and implementation. *Comput. Methods Appl. Mech. Eng.*, 98:41–104, 1992.
- [328] S. Doll and K. Schweizerhof. On the Development of Volumetric Strain Energy Functions. *J. Appl. Mech.*, 67:17–21, 2000.
- [329] S. Hartmann and P. Neff. Polyconvexity of generalized polynomial-type hyperelastic strain energy functions for near-incompressibility. *Int. J. Solids Struct.*, 40:2767–2791, 2003.
- [330] G.A. Holzapfel and R.W. Ogden. Biomechanical relevance of the microstructure in artery walls with a focus on passive and active components. *Am. J. Physiol. Heart Circ. Physiol.*, 315(3):H540–H549, 2018.
- [331] G.A. Holzapfel, R.W. Ogden, and S. Sherifova. On fibre dispersion modelling of soft biological tissues: a review. *Proc. R. Soc. A: Math. Phys. Eng. Sci.*, 475:20180736, 2019.
- [332] S. Sherifova and G.A. Holzapfel. Biomechanics of aortic wall failure with a focus on dissection and aneurysm: A review. *Acta Biomater.*, 99:1–17, 2019.
- [333] G.A. Holzapfel and R.W. Ogden. Constitutive modelling of arteries. *Proc. R. Soc. A*, 466(2118):1551–1597, 2010.
- [334] G.A. Holzapfel, T.C. Gasser, and R.W. Ogden. A new constitutive framework for arterial wall mechanics and a comparative study of material models. *J. Elast.*, 61:1–48, 2000.
- [335] T.C. Gasser, R.W. Ogden, and G.A. Holzapfel. Hyperelastic modelling of arterial layers with distributed collagen fibre orientations. *J. R. Soc. Interface*, 3(6):15–35, 2006.
- [336] G.A. Holzapfel, J.A. Niestrawska, R.W. Ogden, A.J. Reinisch, and A.J. Schriefl. Modelling non-symmetric collagen fibre dispersion in arterial walls. *J. R. Soc. Interface*, 12(106):20150188, 2015.
- [337] K. Li, R.W. Ogden, and G.A. Holzapfel. A discrete fibre dispersion method for excluding fibres under compression in the modelling of fibrous tissues. *J. R. Soc. Interface*, 15:20170766, 2018.

- [338] M. Rolf-Pissarczyk, K. Li, D. Fleischmann, and G.A. Holzapfel. A discrete approach for modeling degraded elastic fibers in aortic dissection. *Comput. Methods Appl. Mech. Eng.*, 373:113511, 2021.
- [339] G.A. Holzapfel and R.W. Ogden. On the tension–compression switch in soft fibrous solids. *Eur. J. Mech. A Solids*, 49:561–569, 2015.
- [340] C. Sansour. On the physical assumptions underlying the volumetric-isochoric split and the case of anisotropy. *Eur. J. Mech. A Solids*, 27(1):28–39, 2008.
- [341] J. Helfenstein, M. Jabareen, E. Mazza, and S. Govindjee. On non-physical response in models for fiber-reinforced hyperelastic materials. *Int. J. Solids Struct.*, 47(16):2056–2061, 2010.
- [342] J. Chung and G.M. Hulbert. A time integration algorithm for structural dynamics with improved numerical dissipation: The generalized- α method. *J. Appl. Mech. Trans ASME*, 60(2):371–375, 1993.
- [343] S. Erlicher, L. Bonaventura, and O.S. Bursi. The analysis of the Generalized- α method for non-linear dynamic problems. *Comput. Mech.*, 28(2):83–104, 2002.
- [344] N.M. Newmark. A method of computation for structural dynamics. *J. Eng. Mech.*, 85(EM3):67–94, 1959.
- [345] H.M. Hilber, T.J.R. Hughes, and R.L. Taylor. Improved numerical dissipation for time integration algorithms in structural dynamics. *Earthq. Eng. Struct. Dyn.*, 5(3):283–292, 1977.
- [346] W.L. Wood, M. Bossak, and O.C. Zienkiewicz. An alpha modification of Newmark’s method. *Int. J. Numer. Methods Eng.*, 15(10):1562–1566, 1980.
- [347] P. Behnoudfar, Q. Deng, and V.M. Calo. High-order generalized-alpha method. *Appl. Eng. Sci.*, 4:100021, 2020.
- [348] P. Behnoudfar, Q. Deng, and V.M. Calo. Higher-order generalized- α methods for hyperbolic problems. *Comput. Methods Appl. Mech. Eng.*, 378:113725, 2021.
- [349] P. Behnoudfar, G. Loli, A. Reali, G. Sangalli, and V.M. Calo. Explicit high-order generalized- α methods for isogeometric analysis of structural dynamics. *Comput. Methods Appl. Mech. Eng.*, 389:114344, 2022.
- [350] J. Rang. Adaptive timestep control for the generalised- α method. In J.P. Moitinho de Almeida, P. Díez, C. Tiago, and N. Parés, editors, *Adaptive Modeling and Simulation*, pages 559–570, Barcelona, 2013. International Center for Numerical Methods in Engineering (CIMNE).
- [351] F. Nobile, M. Pozzoli, and C. Vergara. Time accurate partitioned algorithms for the solution of fluid–structure interaction problems in haemodynamics. *Comput. Fluids*, 86:470–482, 2013.
- [352] F. Nobile, M. Pozzoli, and C. Vergara. Inexact accurate partitioned algorithms for fluid–structure interaction problems with finite elasticity in haemodynamics. *J. Comput. Phys.*, 273:598–617, 2014.

- [353] U. Küttler, M. Gee, C. Förster, A. Comerford, and W.A. Wall. Coupling strategies for biomedical fluid-structure interaction problems. *Int. J. Numer. Methods Biomed. Eng.*, 26:305–321, 2010.
- [354] J.D. Bayer, R.C. Blake, G. Plank, and N.A. Trayanova. A Novel Rule-Based Algorithm for Assigning Myocardial Fiber Orientation to Computational Heart Models. *Ann. Biomed. Eng.*, 40(10):2243–2254, 2012.
- [355] L. Marx, J.A. Niestrawska, M.A.F. Gsell, F. Caforio, G. Plank, and C.M. Augustin. Robust and efficient fixed-point algorithm for the inverse elastostatic problem to identify myocardial passive material parameters and the unloaded reference configuration. *J. Comput. Phys.*, 463:111266, 2022.
- [356] R. Erbel, F. Alfonso, C. Boileau, O. Dirsch, B. Eber, A. Haverich, H. Rakowski, J. Struyven, K. Radegran, U. Sechtem, J. Taylor, C. Zollikofer, W. W. Klein, B. Mulder, and L. A. Providencia. Diagnosis and management of aortic dissection. *Eur. Heart J.*, 22(18):1642–81, 2001.
- [357] P.D. Patel and R.R. Arora. Pathophysiology, diagnosis, and management of aortic dissection. *Ther. Adv. Cardiovasc. Dis.*, 2(6):439–68, 2008.
- [358] Y. Liu, C. Dang, M. Garcia, H. Gregersen, and G.S. Kassab. Surrounding tissues affect the passive mechanics of the vessel wall: theory and experiment. *Am. J. Physiol. Heart Circ. Physiol.*, 293:H3290–H3300, 2007.
- [359] P. Moireau, N. Xiao, M. Astorino, C.A. Figueroa, D. Chapelle, C.A. Taylor, and J.-F. Gerbeau. External tissue support and fluid-structure simulation in blood flows. *Biomech. Model. Mechanobiol.*, 11(1-2):1–18, 2012.
- [360] T.C. Shidhore, H.L. Cebull, M.C. Madden, I.C. Christov, V.L. Rayz, and C.J. Goergen. Estimating external tissue support parameters with fluid-structure interaction models from 4D ultrasound of murine thoracic aortae. *Eng. Comput.*, 2022. doi: 10.1007/s00366-022-01735-1.
- [361] Y. Bazilevs, M.-C. Hsu, Y. Zhang, W. Wang, T. Kvamsdal, S. Hentschel, and J.G. Isaksen. Computational vascular fluid-structure interaction: methodology and application to cerebral aneurysms. *Biomech. Model. Mechanobiol.*, 9(4):481–498, 2010.
- [362] F. Nobile and C. Vergara. Partitioned Algorithms for Fluid-Structure Interaction Problems in Haemodynamics. *Milan J. Math.*, 80(2):443–467, 2012.
- [363] P. Reymond, P. Crosetto, S. Deparis, A. Quarteroni, and N. Stergiopulos. Physiological simulation of blood flow in the aorta: Comparison of hemodynamic indices as predicted by 3-D FSI, 3-D rigid wall and 1-D models. *Med. Eng. Phys.*, 35(6):784–791, 2013.
- [364] P. Crosetto, P. Reymond, S. Deparis, D. Kontaxakis, N. Stergiopulos, and A. Quarteroni. Fluid-structure interaction simulation of aortic blood flow. *Comput. Fluids*, 43(1):46–57, 2011.
- [365] P. Causin, J.F. Gerbeau, and F. Nobile. Added-mass effect in the design of partitioned algorithms for fluid-structure problems. *Comput. Methods Appl. Mech. Eng.*, 194(42-44):4506–4527, 2005.

- [366] C. Förster, W.A. Wall, and E. Ramm. Artificial added mass instabilities in sequential staggered coupling of nonlinear structures and incompressible viscous flows. *Comput. Methods Appl. Mech. Eng.*, 196(7):1278–1293, 2007.
- [367] M. Lesoinne and C. Farhat. Higher-Order Subiteration-Free Staggered Algorithm for Nonlinear Transient Aeroelastic Problems. *AIAA J.*, 36(9):1754–1757, 1998.
- [368] C. Kassiotis, A. Ibrahimbegovic, R. Niekamp, and H.G. Matthies. Nonlinear fluid–structure interaction problem. Part I: implicit partitioned algorithm, nonlinear stability proof and validation examples. *Comput. Mech.*, 47(3):305–323, 2011.
- [369] R.M. Kirby, Z. Yosibash, and G.E. Karniadakis. Towards stable coupling methods for high-order discretization of fluid–structure interaction: Algorithms and observations. *J. Comput. Phys.*, 223(2):489–518, 2007.
- [370] M.W. Gee, C. Reeps, H.H. Eckstein, and W.A. Wall. Prestressing in finite deformation abdominal aortic aneurysm simulation. *J. Biomech.*, 42(11):1732–1739, 2009.
- [371] M.W. Gee, C. Förster, and W.A. Wall. A computational strategy for prestressing patient-specific biomechanical problems under finite deformation. *Int. J. Numer. Methods Biomed. Eng.*, 26(1):52–72, 2010.
- [372] T.E. Tezduyar, S. Sathe, M. Schwaab, and B.S. Conklin. Arterial fluid mechanics modeling with the stabilized space–time fluid–structure interaction technique. *Int. J. Numer. Methods Fluids*, 57(5):601–629, 2008.
- [373] T. Machairas, A. Kontogiannis, A. Karakalas, A. Solomou, V. Riziotis, and D. Saravanos. Robust fluid-structure interaction analysis of an adaptive airfoil using shape memory alloy actuators. *Smart Mater. Struct.*, 27(10):105035, 2018.
- [374] Y. Bazilevs, M.-C. Hsu, I. Akkerman, S. Wright, K. Takizawa, B. Henicke, T. Spielman, and T.E. Tezduyar. 3D simulation of wind turbine rotors at full scale. Part I: Geometry modeling and aerodynamics. *Int. J. Numer. Methods Fluids*, 65(1-3): 207–235, 2011.
- [375] T.A. Helgedagsrud, Y. Bazilevs, K.M. Mathisen, and O.A. Øiseth. ALE-VMS methods for wind-resistant design of long-span bridges. *J. Wind Eng. Ind. Aerodyn.*, 191:143–153, 2019.
- [376] D. Shiels, A. Leonard, and A. Roshko. Flow-induced vibration of a circular cylinder at limiting structural parameters. *J. Fluids Struct.*, 15(1):3–21, 2001.
- [377] P.W. Bearman. Circular cylinder wakes and vortex-induced vibrations. *J. Fluids Struct.*, 27(5-6):648–658, 2011.
- [378] Y.J. Chu, P.B. Ganesan, and M.A. Ali. Fluid–structure interaction simulation on flight performance of a dragonfly wing under different pterostigma weights. *J. Mech.*, 37:216–229, 2021.
- [379] Y. Bazilevs, V.M. Calo, Y. Zhang, and T.J.R. Hughes. Isogeometric Fluid–structure Interaction Analysis with Applications to Arterial Blood Flow. *Comput. Mech.*, 38(4-5):310–322, 2006.

- [380] P. Crosetto, S. Deparis, G. Fourestey, and A. Quarteroni. Parallel Algorithms for Fluid-Structure Interaction Problems in Haemodynamics. *SIAM J. Sci. Comput.*, 33(4):1598–1622, 2011.
- [381] R. Torii, M. Oshima, T. Kobayashi, K. Takagi, and T.E. Tezduyar. Computer modeling of cardiovascular fluid–structure interactions with the deforming-spatial-domain/stabilized space–time formulation. *Comput. Methods Appl. Mech. Eng.*, 195(13-16):1885–1895, 2006.
- [382] S.L. Thomson, L. Mongeau, and S.H. Frankel. Aerodynamic transfer of energy to the vocal folds. *J. Acoust. Soc. Am.*, 118(3):1689–1700, 2005.
- [383] H. Luo, R. Mittal, and S.A. Bielamowicz. Analysis of flow-structure interaction in the larynx during phonation using an immersed-boundary method. *J. Acoust. Soc.*, 126(2):816–824, 2009.
- [384] W.A. Wall and T. Rabczuk. Fluid–structure interaction in lower airways of CT-based lung geometries. *Int. J. Numer. Methods Fluids*, 57(5):653–675, 2008.
- [385] Y. Qiao, Y. Zeng, Y. Ding, J. Fan, K. Luo, and T. Zhu. Numerical simulation of two-phase non-Newtonian blood flow with fluid-structure interaction in aortic dissection. *Comput. Methods Biomech. Biomed. Eng.*, 22(6):620–630, 2019.
- [386] M. Alimohammadi, J. Sherwood, M. Karimpour, O. Agu, S. Balabani, and V. Díaz-Zuccarini. Aortic dissection simulation models for clinical support: fluid-structure interaction vs. rigid wall models. *Biomed. Eng. Online*, 14(34), 2015.
- [387] A. Qiao, W. Yin, and B. Chu. Numerical simulation of fluid–structure interaction in bypassed DeBakey III aortic dissection. *Comput. Methods Biomech. Biomed. Eng.*, 18(11):1173–1180, 2015.
- [388] J. Zimmermann, K. Bäumlner, M. Loecher, T.E. Cork, F.O. Kolawole, K. Gifford, A.L. Marsden, D. Fleischmann, and D.B. Ennis. Quantitative Hemodynamics in Aortic Dissection: Comparing in Vitro MRI with FSI Simulation in a Compliant Model. In D.B. Ennis, L.E. Perotti, and V.Y. Wang, editors, *Functional Imaging and Modeling of the Heart*, pages 575–586, Cham, 2021. Springer International Publishing.
- [389] B.J.B.M. Wolters, M.C.M. Rutten, G.W.H. Schurink, U. Kose, J. de Hart, and F.N. van de Vosse. A patient-specific computational model of fluid–structure interaction in abdominal aortic aneurysms. *Med. Eng. Phys.*, 27(10):871–883, 2005.
- [390] M. Xenos, S.H. Rambhia, Y. Alemu, S. Einav, N. Labropoulos, A. Tassiopoulos, J.J. Ricotta, and D. Bluenstein. Patient-Based Abdominal Aortic Aneurysm Rupture Risk Prediction with Fluid Structure Interaction Modeling. *Ann. Biomed. Eng.*, 38:3323–3337, 2010.
- [391] J.H.C. Fonken, E.J. Maas, A.H.M. Nievergeld, M.R.H.M. van Sambeek, F.N. van de Vosse, and R.G.P. Lopata. Ultrasound-Based Fluid-Structure Interaction Modeling of Abdominal Aortic Aneurysms Incorporating Pre-stress. *Front. Physiol.*, 12, 2021.

- [392] I.S. Lan, J. Liu, W. Yang, and A.L. Marsden. Numerical investigation of abdominal aortic aneurysm hemodynamics using the reduced unified continuum formulation for vascular fluid-structure interaction. *Forces. Mech.*, 7:100089, 2022.
- [393] S. Brandstaeter, S.L. Fuchs, J. Biehler, R.C. Aydin, W.A. Wall, and C.J. Cyron. Global Sensitivity Analysis of a Homogenized Constrained Mixture Model of Arterial Growth and Remodeling. *J. Elast.*, 145:191–221, 2021.
- [394] J.S. Tran, D.E. Schiavazzi, A.M. Kahn, and A.L. Marsden. Uncertainty quantification of simulated biomechanical stimuli in coronary artery bypass grafts. *Comput. Methods Appl. Mech. Eng.*, 345:402–428, 2019.
- [395] D.E. Schiavazzi, G. Arbia, C. Baker, A.M. Hlavacek, T.Y. Hsia, A.L. Marsden, I.E. Vignon-Clementel, and The Modeling of Congenital Hearts Alliance (MOCHA) Investigators. Uncertainty quantification in virtual surgery hemodynamics predictions for single ventricle palliation. *Int. J. Numer. Methods Biomed. Eng.*, 32(3):e02737, 2016.
- [396] J. Biehler, M.W. Gee, and W.A. Wall. Towards efficient uncertainty quantification in complex and large-scale biomechanical problems based on a Bayesian multi-fidelity scheme. *Biomech. Model. Mechanobiol.*, 14(3):489–513, 2015.
- [397] J. Biehler, S. Kehl, M.W. Gee, F. Schmies, J. Pelisek, A. Maier, C. Reeps, H.-H. Eckstein, and W.A. Wall. Probabilistic noninvasive prediction of wall properties of abdominal aortic aneurysms using Bayesian regression. *Biomech. Model. Mechanobiol.*, 16:45–61, 2017.
- [398] B. Wirthl, S. Brandstaeter, J. Nitzler, B.A. Schrefler, and W.A. Wall. Global sensitivity analysis based on Gaussian-process metamodelling for complex biomechanical problems. *arXiv*, 2022. doi: 10.48550/ARXIV.2202.01503.
- [399] A. Gerstenberger and W.A. Wall. An eXtended Finite Element Method/Lagrange multiplier based approach for fluid–structure interaction. *Comput. Methods Appl. Mech. Eng.*, 197:1699–1714, 2008.
- [400] U.M. Mayer, A. Gerstenberger, and W.A. Wall. Interface handling for three-dimensional higher-order XFEM-computations in fluid-structure interaction. *Int. J. Numer. Methods Eng.*, 79:846–869, 2009.
- [401] U.M. Mayer, A. Popp, A. Gerstenberger, and W.A. Wall. 3D fluid–structure-contact interaction based on a combined XFEM FSI and dual mortar contact approach. *Comput. Mech.*, 46:53–67, 2010.
- [402] L. Formaggia, F. Gatti, and S. Zonca. An XFEM/DG approach for fluid-structure interaction problems with contact. *Appl. Math.*, 66:183–211, 2021.
- [403] F. Alauzet, B. Fabrèges, M.A. Fernández, and M. Landajuela. Nitsche-XFEM for the coupling of an incompressible fluid with immersed thin-walled structures. *Comput. Methods Appl. Mech. Eng.*, 301:300–335, 2016.
- [404] M. Mayr, M.H. Noll, and M.W. Gee. A hybrid interface preconditioner for monolithic fluid–structure interaction solvers. *Adv. Model. Simul. Eng. Sci.*, 7(1), 2020.

- [405] U. Langer and H. Yang. Numerical simulation of fluid–structure interaction problems with hyperelastic models: A monolithic approach. *Math. Comput. Simul.*, 145: 186–208, 2018.
- [406] U. Langer and H. Yang. Robust and efficient monolithic fluid-structure-interaction solvers. *Int. J. Numer. Methods Eng.*, 108(4):303–325, 2016.
- [407] A. Massing, M. Larson, A. Logg, and M. Rognes. A Nitsche-based cut finite element method for a fluid-structure interaction problem. *Comm. App. Math. Comp. Sci.*, 10(2):97–120, 2015.
- [408] B. Schott, C. Ager, and W.A. Wall. Monolithic cut finite element–based approaches for fluid–structure interaction. *Int. J. Numer. Methods Eng.*, 119(8):757–796, 2019.
- [409] E. Burman, M.A. Fernández, and S. Frei. A Nitsche-based formulation for fluid-structure interactions with contact. *Esaim Math. Model. Numer. Anal.*, 54(2): 531–564, 2020.
- [410] T. Klöppel, A. Popp, U. Küttler, and W.A. Wall. Fluid–structure interaction for non-conforming interfaces based on a dual mortar formulation. *Comput. Methods Appl. Mech. Eng.*, 200(45-46):3111–3126, 2011.
- [411] Y. Kim and C.S. Peskin. A penalty immersed boundary method for a rigid body in fluid. *Phys. Fluids*, 28(3):033603, 2016.
- [412] A. Viré, J. Xiang, and C.C. Pain. An immersed-shell method for modelling fluid–structure interactions. *Philos. Trans. Royal Soc. A*, 373(2035):20140085, 2015.
- [413] A. Viré, J. Spinneken, M.D. Piggott, C.C. Pain, and S.C. Kramer. Application of the immersed-body method to simulate wave–structure interactions. *Eur. J. Mech. B Fluids*, 55:330–339, 2016.
- [414] J. Hron and S. Turek. A Monolithic FEM/Multigrid Solver for an ALE Formulation of Fluid-Structure Interaction with Applications in Biomechanics. In *Fluid-Structure Interaction*, pages 146–170. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- [415] T. Richter. A monolithic geometric multigrid solver for fluid-structure interactions in ALE formulation. *Int. J. Numer. Methods Eng.*, 104(5):372–390, 2015.
- [416] T. Wick. Fully Eulerian fluid–structure interaction for time-dependent problems. *Comput. Methods Appl. Mech. Eng.*, 255:14–26, 2013.
- [417] D. Jodlbauer, U. Langer, and T. Wick. Parallel block-preconditioned monolithic solvers for fluid-structure interaction problems. *Int. J. Numer. Methods Eng.*, 117(6):623–643, 2019.
- [418] M. Balmus, A. Massing, J. Hoffman, R. Razavi, and D.A. Nordsletten. A partition of unity approach to fluid mechanics and fluid–structure interaction. *Comput. Methods Appl. Mech. Eng.*, 362:112842, 2020.
- [419] A.T. Barker and X.-C. Cai. Scalable parallel methods for monolithic coupling in fluid–structure interaction with application to blood flow modeling. *J. Comput. Phys.*, 229(3):642–659, 2010.

- [420] Y. Wu and X.-C. Cai. A fully implicit domain decomposition based ALE framework for three-dimensional fluid–structure interaction with application in blood flow computation. *J. Comput. Phys.*, 258:524–537, 2014.
- [421] T.E. Tezduyar, S. Sathe, R. Keedy, and K. Stein. Space–time finite element techniques for computation of fluid–structure interactions. *Comput. Methods Appl. Mech. Eng.*, 195(17-18):2002–2027, 2006.
- [422] M.W. Gee, U. Küttler, and W.A. Wall. Truly monolithic algebraic multigrid for fluid–structure interaction. *Int. J. Numer. Methods Eng.*, 85(8):987–1016, 2011.
- [423] E. Aulisa, S. Bnà, and G. Bornia. A monolithic ALE Newton–Krylov solver with Multigrid–Richardson–Schwarz preconditioning for incompressible Fluid–Structure Interaction. *Comput. Fluids*, 174:213–228, 2018.
- [424] F. Kong and X.-C. Cai. Scalability study of an implicit solver for coupled fluid–structure interaction problems on unstructured meshes in 3D. *Int. J. High Perform. Comput. Appl.*, 32:207–219, 2018.
- [425] J. Degroote. Partitioned Simulation of Fluid–Structure Interaction. *Arch. Comput. Methods Eng.*, 20(3):185–238, 2013.
- [426] G. Hou, J. Wang, and A. Layton. Numerical Methods for Fluid–Structure Interaction - A Review. *Commun. Comput. Phys.*, 12(2):337–377, 2012.
- [427] N. Hosters, J. Helmig, A. Stavrev, M. Behr, and S. Elgeti. Fluid–structure interaction with NURBS-based coupling. *Comput. Methods Appl. Mech. Eng.*, 332: 520–539, 2018.
- [428] D. Hilger, N. Hosters, F. Key, S. Elgeti, and M. Behr. A Novel Approach to Fluid–Structure Interaction Simulations Involving Large Translation and Contact. In H. van Brummelen, C. Vuik, M. Möller, C. Verhoosel, B. Simeon, and B. Jüttler, editors, *Isogeometric Analysis and Applications 2018*, pages 39–56, Cham, 2021. Springer International Publishing.
- [429] E.H. van Brummelen. Added Mass Effects of Compressible and Incompressible Flows in Fluid–Structure Interaction. *J. Appl. Mech.*, 76, 2009.
- [430] J.-F. Gerbeau and M. Vidrascu. A Quasi-Newton Algorithm Based on a Reduced Model for Fluid–Structure Interaction Problems in Blood Flows. *Esaim Math. Model. Numer. Anal.*, 37(4):631–647, 2003.
- [431] J.-F. Gerbeau, M. Vidrascu, and P. Frey. Fluid–structure interaction in blood flows on geometries based on medical imaging. *Comput. Struct.*, 83:155–165, 2005.
- [432] R. Haelterman, A.E.J. Bogaers, K. Scheufele, B. Uekermann, and M. Mehl. Improving the performance of the partitioned QN-ILS procedure for fluid–structure interaction problems: Filtering. *Comput. Struct.*, 171:9–17, 2016.
- [433] C. Michler, E.H. van Brummelen, and R. de Borst. An interface Newton–Krylov solver for fluid–structure interaction. *Int. J. Numer. Methods Fluids*, 47(10-11): 1189–1195, 2005.
- [434] M.Á. Fernández and M. Moubachir. A Newton method using exact Jacobians for solving fluid–structure coupling. *Comput. Struct.*, 83(2-3):127–142, 2005.

- [435] J. Degroote, K.-J. Bathe, and J. Vierendeels. Performance of a new partitioned procedure versus a monolithic procedure in fluid–structure interaction. *Comput. Struct.*, 87(11-12):793–801, 2009.
- [436] T. Spenke, N. Hosters, and M. Behr. A multi-vector interface quasi-Newton method with linear complexity for partitioned fluid–structure interaction. *Comput. Methods Appl. Mech. Eng.*, 361:112810, 2020.
- [437] J. Degroote, A. Swillens, P. Bruggeman, R. Haelterman, P. Segers, and J. Vierendeels. Simulation of fluid–structure interaction with the interface artificial compressibility method. *Int. J. Numer. Method. Biomed. Eng.*, 26(3-4):276–289, 2010.
- [438] A.E.J. Bogaers, S. Kok, B.D. Reddy, and T. Franz. Extending the robustness and efficiency of artificial compressibility for partitioned fluid–structure interactions. *Comput. Methods Appl. Mech. Eng.*, 283:1278–1295, 2015.
- [439] A. La Spina, C. Förster, M. Kronbichler, and W.A. Wall. On the role of (weak) compressibility for fluid–structure interaction solvers. *Int. J. Numer. Methods Fluids*, 92:129–147, 2020.
- [440] A. La Spina, M. Kronbichler, M. Giacomini, W.A. Wall, and A. Huerta. A weakly compressible hybridizable discontinuous Galerkin formulation for fluid–structure interaction problems. *Comput. Methods Appl. Mech. Eng.*, 372:113392, 2020.
- [441] M.A. Fernández, J.-F. Gerbeau, and C. Grandmont. A projection semi-implicit scheme for the coupling of an elastic structure with an incompressible fluid. *Int. J. Numer. Methods Eng.*, 69(4):794–821, 2007.
- [442] S. Badia, A. Quaini, and A. Quarteroni. Splitting Methods Based on Algebraic Factorization for Fluid-Structure Interaction. *SIAM J. Sci. Comput.*, 30(4):1778–1805, 2008.
- [443] M. Breuer, G. De Nayer, M. Münsch, T. Gallinger, and R. Wüchner. Fluid–structure interaction using a partitioned semi-implicit predictor–corrector coupling scheme for the application of large-eddy simulation. *J. Fluids Struct.*, 29:107–130, 2012.
- [444] A. Lozovskiy, M.A. Olshanskii, V. Salamatova, and Y.V. Vassilevski. An unconditionally stable semi-implicit FSI finite element method. *Comput. Methods Appl. Mech. Eng.*, 297:437–454, 2015.
- [445] T. He. On a Partitioned Strong Coupling Algorithm for Modeling Fluid–Structure Interaction. *Int. J. Appl. Mech.*, 07(02):1550021, 2015.
- [446] M. Landajuela, M. Vidrascu, D. Chapelle, and M.A. Fernández. Coupling schemes for the FSI forward prediction challenge: Comparative study and validation. *Int. J. Numer. Method. Biomed. Eng.*, 33(4):e2813, 2017.
- [447] A. Naseri, O. Lehmkuhl, I. Gonzalez, E. Bartrons, C.D. Pérez-Segarra, and A. Oliva. A semi-implicit coupling technique for fluid–structure interaction problems with strong added-mass effect. *J. Fluids Struct.*, 80:94–112, 2018.
- [448] M.A. Fernández and F.M. Gerosa. An unfitted mesh semi-implicit coupling scheme for fluid–structure interaction with immersed solids. *Int. J. Numer. Methods Eng.*, 122(19):5384–5408, 2021.

- [449] G. Guidoboni, R. Glowinski, N. Cavallini, and S. Canic. Stable loosely-coupled-type algorithm for fluid–structure interaction in blood flow. *J. Comput. Phys.*, 228(18):6916–6937, 2009.
- [450] M. Lukáčová-Medvid’ová, G. Rusnáková, and A. Hundertmark-Zaušková. Kinematic splitting algorithm for fluid–structure interaction in hemodynamics. *Comput. Methods Appl. Mech. Eng.*, 265:83–106, 2013.
- [451] J.W. Banks, W.D. Henshaw, and D.W. Schwendeman. An analysis of a new stable partitioned algorithm for FSI problems. Part II: Incompressible flow and structural shells. *J. Comput. Phys.*, 268:399–416, 2014.
- [452] M.A. Fernández, M. Landajuela, and M. Vidrascu. Fully decoupled time-marching schemes for incompressible fluid/thin-walled structure interaction. *J. Comput. Phys.*, 297:156–181, 2015.
- [453] E. Burman and M.A. Fernández. Stabilization of explicit coupling in fluid–structure interaction involving fluid incompressibility. *Comput. Methods Appl. Mech. Eng.*, 198(5-8):766–784, 2009.
- [454] J.W. Banks, W.D. Henshaw, and D.W. Schwendeman. An analysis of a new stable partitioned algorithm for FSI problems. Part I: Incompressible flow and elastic solids. *J. Comput. Phys.*, 269:108–137, 2014.
- [455] D.A. Serino, J.W. Banks, W.D. Henshaw, and D.W. Schwendeman. A Stable Added-Mass Partitioned (AMP) Algorithm for Elastic Solids and Incompressible Flow: Model Problem Analysis. *SIAM J. Sci. Comput.*, 41(4):A2464–A2484, 2019.
- [456] D.A. Serino, J.W. Banks, W.D. Henshaw, and D.W. Schwendeman. A stable added-mass partitioned (AMP) algorithm for elastic solids and incompressible flow. *J. Comput. Phys.*, 399:108923, 2019.
- [457] M. Heil, A.L. Hazel, and J. Boyle. Solvers for large-displacement fluid-structure interaction problems: Segregated versus monolithic approaches. *Comput. Mech.*, 43(1):91–101, 2008.
- [458] S. Badia, A. Quaini, and A. Quarteroni. Modular vs. non-modular preconditioners for fluid-structure systems with large added-mass effect. *Comput. Methods Appl. Mech. Eng.*, 197(49-50):4216–4232, 2008.
- [459] J. Degroote, R. Haelterman, S. Annerel, P. Bruggeman, and J. Vierendeels. Performance of partitioned procedures in fluid–structure interaction. *Comput. Struct.*, 88(7):446–457, 2010.
- [460] J.P. Sheldon, S.T. Miller, and J.S. Pitt. Methodology for Comparing Coupling Algorithms for Fluid-Structure Interaction Problems. *World J. Mech.*, 4:54–70, 2014.
- [461] A.E.J. Bogaers, S. Kok, B.D. Reddy, and T. Franz. Quasi-Newton methods for implicit black-box FSI coupling. *Comput. Methods Appl. Mech. Eng.*, 279:113–132, 2014.
- [462] T. Spenke, M. Make, and N. Hosters. A Robin-Neumann Scheme with quasi-Newton Acceleration for Partitioned Fluid-Structure Interaction. *Int. J. Numer. Methods Eng.*, 2022. doi: 10.1002/nme.7151.

- [463] J. Janela, A. Moura, and A. Sequeira. A 3D non-Newtonian fluid–structure interaction model for blood flow in arteries. *J. Comput. Appl. Math.*, 234(9):2783–2791, 2010.
- [464] L. Zhu, X. Yu, N. Liu, Y. Cheng, and X. Lu. A deformable plate interacting with a non-Newtonian fluid in three dimensions. *Phys. Fluids*, 29(8):083101, 2017.
- [465] L. Zhu. An IB Method for Non-Newtonian-Fluid Flexible-Structure Interactions in Three-Dimensions. *Comput. Model. Eng. Sci.*, 119(1):125–143, 2019.
- [466] A. Amani, A. Naseri, C.D. Pérez-Segarra, and A. Oliva. A method for fluid-structure interaction problems with non-Newtonian fluid. In *Proceedings of the 6th European Conference on Computational Mechanics and 7th European Conference on Computational Fluid Dynamics*, Glasgow, 2018.
- [467] A. Jendoubi, J. Deteix, and A. Fortin. A simple mesh-update procedure for fluid–structure interaction problems. *Comput. Struct.*, 169:13–23, 2016.
- [468] K. Takizawa, T.E. Tezduyar, J. Boben, N. Kostov, C. Boswell, and A. Buscher. Fluid–structure interaction modeling of clusters of spacecraft parachutes with modified geometric porosity. *Comput. Mech.*, 52:1351–1364, 2013.
- [469] H. Suito, K. Takizawa, V.Q.H. Huynh, D. Sze, and T. Ueda. FSI analysis of the blood flow and geometrical characteristics in the thoracic aorta. *Comput. Mech.*, 54:1035–1045, 2014.
- [470] K. Takizawa, T.E. Tezduyar, and R. Avsar. A low-distortion mesh moving method based on fiber-reinforced hyperelasticity and optimized zero-stress state. *Comput. Mech.*, 65:1567–1591, 2020.
- [471] A. Shamanskiy and B. Simeon. Mesh moving techniques in fluid-structure interaction: robustness, accumulated distortion and computational efficiency. *Comput. Mech.*, 67(2):583–600, 2021.
- [472] N. Biocca, P.J. Blanco, D.E. Caballero, J.M. Gimenez, G.E. Carr, and S.A. Urquiza. A biologically-inspired mesh optimizer based on pseudo-material remodeling. *Comput. Mech.*, 69:505–525, 2022.
- [473] T. Wick. Fluid-structure interactions using different mesh motion techniques. *Comput. Struct.*, 89(13-14):1456–1467, 2011.
- [474] A.A. Johnson and T.E. Tezduyar. Mesh update strategies in parallel finite element computations of flow problems with moving boundaries and interfaces. *Comput. Methods Appl. Mech. Eng.*, 119(1):73–94, 1994.
- [475] K. Stein, T. Tezduyar, and R. Benney. Mesh Moving Techniques for Fluid-Structure Interactions With Large Displacements. *J. Appl. Mech.*, 70:58–63, 2003.
- [476] C. Grandmont, V. Guimet, and Y. Maday. Numerical analysis of some decoupling techniques for the approximation of the unsteady fluid structure interaction. *Math. Models Methods Appl. Sci.*, 11(08):1349–1377, 2001.
- [477] D. Mok and W.A. Wall. Partitioned Analysis Schemes for the Transient Interaction of Incompressible Flows and Nonlinear Flexible Structures. In W.A. Wall, K.-U. Bletzinger, and K. Schweitzerhof, editors, *Trends in Computational Structural Mechanics*. CIMNE, Cham, 2001.

- [478] F. Lindner, M. Mehl, K. Scheufele, and B. Uekermann. A comparison of various quasi-Newton schemes for partitioned fluid-structure interaction. In B. Schrefler, E. Oñate, and M. Papadrakakis, editors, *Proceedings of the 6th International Conference on Coupled Problems in Science and Engineering*, Venice, 2015.
- [479] L. Radtke, A. Larena-Avellaneda, E.S. Debus, and A. Düster. Convergence acceleration for partitioned simulations of the fluid-structure interaction in arteries. *Comput. Mech.*, 57:901–920, 2016.
- [480] J. Degroote and J. Vierendeels. Multi-solver algorithms for the partitioned simulation of fluid–structure interaction. *Comput. Methods Appl. Mech. Eng.*, 200:2195–2210, 2011.
- [481] K. Scheufele and M. Mehl. Robust Multisecant Quasi-Newton Variants for Parallel Fluid-Structure Simulations—and Other Multiphysics Applications. *SIAM J. Sci. Comput.*, 39:S404–S433, 2017.
- [482] N. Fehn. *Robust and Efficient Discontinuous Galerkin Methods for Incompressible Flows*. PhD thesis, Technische Universität München, 2021.
- [483] J. Vierendeels, L. Lanoye, J. Degroote, and P. Verdonck. Implicit coupling of partitioned fluid–structure interaction problems with reduced order models. *Comput. Struct.*, 85:970–976, 2007.
- [484] G.H. Golub and C.F. van Loan. *Matrix computations*. Johns Hopkins University Press, Baltimore, 4th edition, 2013.
- [485] T.C.S. Rendall and C.B. Allen. Unified fluid–structure interpolation and mesh motion using radial basis functions. *Int. J. Numer. Methods Eng.*, 74(10):1519–1559, 2008.
- [486] T.C.S. Rendall and C.B. Allen. Efficient mesh motion using radial basis functions with data reduction algorithms. *J. Comput. Phys.*, 228(17):6231–6249, 2009.
- [487] A. de Boer, M.S. van der Schoot, and H. Bijl. Mesh deformation based on radial basis function interpolation. *Comput. Struct.*, 85(11):784–795, 2007.
- [488] P.-O. Persson, J. Peraire, and J. Bonet. A High Order Discontinuous Galerkin Method for Fluid–Structure Interaction. In *18th AIAA Computational Fluid Dynamics Conference*, page 4327, 2007.
- [489] B. Froehle and P.-O. Persson. A high-order discontinuous Galerkin method for fluid–structure interaction with efficient implicit–explicit time stepping. *J. Comput. Phys.*, 272:455–470, 2014.
- [490] J.T. Batina. Unsteady Euler algorithm with unstructured dynamic mesh for complex-aircraft aerodynamic analysis. *AIAA J.*, 29(3):327–333, 1991.
- [491] F.J. Blom. Considerations on the spring analogy. *Int. J. Numer. Methods Fluids*, 32(6):647–668, 2000.
- [492] C. Farhat, C. Degand, B. Koobus, and M. Lesoinne. Torsional springs for two-dimensional dynamic unstructured fluid meshes. *Comput. Methods Appl. Mech. Eng.*, 163(1):231–245, 1998.

- [493] C. Degand and C. Farhat. A three-dimensional torsional spring analogy method for unstructured dynamic meshes. *Comput. Struct.*, 80(3):305–316, 2002.
- [494] C.L. Bottasso, D. Detomi, and R. Serra. The ball-vertex method: a new simple spring analogy method for unstructured dynamic meshes. *Comput. Methods Appl. Mech. Eng.*, 194(39):4244–4264, 2005. ISSN 0045-7825.
- [495] T.J. Lin, Z.Q. Guan, J.H. Chang, and S.H. Lo. Vertex-Ball Spring Smoothing: An efficient method for unstructured dynamic hybrid meshes. *Comput. Struct.*, 136: 24–33, 2014.
- [496] H. Kanchi and A. Masud. A 3D adaptive mesh moving scheme. *Int. J. Numer. Methods Fluids*, 54(6-8):923–944, 2007.
- [497] R. Löhner and C. Yang. Improved ALE mesh velocities for moving bodies. *Commun. Numer. Methods Eng.*, 12(10):599–608, 1996.
- [498] P.A. Sackinger, P.R. Schunk, and R.R. Rao. A Newton–Raphson Pseudo-Solid Domain Mapping Technique for Free and Moving Boundary Problems: A Finite Element Implementation. *J. Comput. Phys.*, 125(1):83–103, 1996.
- [499] K. Stein, T.E. Tezduyar, and R. Benney. Automatic mesh update with the solid-extension mesh moving technique. *Comput. Methods Appl. Mech. Eng.*, 193(21): 2019–2032, 2004.
- [500] B.T. Helenbrook. Mesh deformation using the biharmonic operator. *Int. J. Numer. Methods Eng.*, 56(7):1007–1021, 2003.
- [501] T.E. Tezduyar. Finite element methods for flow problems with moving boundaries and interfaces. *Arch. Comput. Methods Eng.*, 8:83–130, 2001.
- [502] S. Yigit, M. Schäfer, and M. Heck. Grid movement techniques and their influence on laminar fluid–structure interaction computations. *J. Fluids Struct.*, 24(6):819–832, 2008.
- [503] M. Mayr. *A Monolithic Solver for Fluid-Structure Interaction with Adaptive Time Stepping and a Hybrid Preconditioner*. PhD thesis, Technische Universität München, 2016.
- [504] L. Bruder. A non-linear ALE mesh moving strategy for large-displacement fluid-structure interaction problems. Bachelor thesis, Technische Universität München, 2014.
- [505] P. Lamata, I. Roy, B. Blazevic, A. Crozier, S. Land, S.A. Niederer, D.R. Hose, and N.P. Smith. Quality metrics for high order meshes: Analysis of the mechanical simulation of the heart beat. *IEEE Trans. Med. Imaging*, 32(1):130–138, 2013.
- [506] T.J. Baker. Mesh Movement and Metamorphosis. *Eng. Comput.*, 18:188–198, 2002.
- [507] S.J. Sherwin and J. Peiró. Mesh generation in curvilinear domains using high-order elements. *Int. J. Numer. Methods Eng.*, 53:207–223, 2002.
- [508] S. Kim, Y.I. Cho, A.H. Jeon, B. Hogenauer, and K.R. Kensey. A new method for blood viscosity measurement. *J. Non-Newton. Fluid Mech.*, 94(1):47–56, 2000.

- [509] H. Weisbecker, D.M. Pierce, P. Regitnig, and G.A. Holzapfel. Layer-specific damage experiments and modeling of human thoracic and abdominal aortas with non-atherosclerotic intimal thickening. *J. Mech. Behav. Biomed. Mater.*, 12:93–106, 2012.
- [510] Y. Dimakopoulos, A.C.B. Bogaerds, P.D. Anderson, M.A. Hulsen, and F.P.T. Baaijens. Direct numerical simulation of a 2D-stented aortic heart valve at physiological flow rates. *Comput. Methods Biomech. Biomed. Eng.*, 15(11):1157–1179, 2012.
- [511] C. Geuzaine and J.-F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *Int. J. Numer. Methods Eng.*, 79(11): 1309–1331, 2009.
- [512] R. Schussnig, D.R.Q. Pacheco, M. Kaltenbacher, and T.-P. Fries. Efficient and higher-order accurate split-step methods for generalised Newtonian fluid flow. In *Extended Abstracts of the 6th ECCOMAS Young Investigators Conference*, València, 2021.
- [513] R. Schmidt and K. Singh. Meshmixer: An Interface for Rapid Mesh Composition. In *ACM SIGGRAPH 2010 Talks, SIGGRAPH '10*, Los Angeles, California, 2010.
- [514] J.D. Humphrey and C.A. Taylor. Intracranial and Abdominal Aortic Aneurysms: Similarities, Differences, and Need for a New Class of Computational Models. *Annu. Rev. Biomed. Eng.*, 10:221–246, 2008.
- [515] T. Canchi, S.D. Kumar, E.Y.K. Ng, and S. Narayanan. A Review of Computational Methods to Predict the Risk of Rupture of Abdominal Aortic Aneurysms. *BioMed. Res. Int.*, 2015:1–12, 2015.
- [516] A.H. Malkawi, R.J. Hinchliffe, Y. Xu, P.J. Holt, I.M. Loftus, and M.M. Thompson. Patient-specific biomechanical profiling in abdominal aortic aneurysm development and rupture. *J. Vasc. Surg.*, 52:480–488, 2010.
- [517] S. Lin, X. Han, Y. Bi, S. Ju, and L. Gu. Fluid-Structure Interaction in Abdominal Aortic Aneurysm: Effect of Modeling Techniques. *Biomed Res. Int.*, 2017:1–10, 2017.
- [518] S. Turek, J. Hron, M. Mádlík, M. Razzaq, H. Wobker, and J. F. Acker. Numerical Simulation and Benchmarking of a Monolithic Multigrid Solver for Fluid-Structure Interaction Problems with Application to Hemodynamics. In H.-J. Bungartz, M. Mehl, and M. Schäfer, editors, *Fluid Structure Interaction II*, pages 193–220, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [519] Y. Bazilevs, M.-C. Hsu, Y. Zhang, W. Wang, X. Liang, T. Kvamsdal, R. Brekken, and J.G. Isaksen. A fully-coupled fluid-structure interaction simulation of cerebral aneurysms. *Comput. Mech.*, 46(1):3–16, 2010.
- [520] B.J. Doyle, T.M. McGloughlin, K. Miller, J.T. Powell, and P.E. Norman. Regions of High Wall Stress Can Predict the Future Location of Rupture of Abdominal Aortic Aneurysm. *CardioVasc. Interv. Radiol.*, 37:815–818, 2014.
- [521] A. Hemmler, B. Lutz, C. Reeps, G. Kalender, and M.W. Gee. A methodology for in silico endovascular repair of abdominal aortic aneurysms. *Biomech. Model. Mechanobiol.*, 17:1139–1164, 2018.

- [522] C.A. Meyer, E. Bertrand, O. Boiron, and V. Deplano. Stereoscopically Observed Deformations of a Compliant Abdominal Aortic Aneurysm Model. *J. Biomech. Eng.*, 133(11), 2011.
- [523] V. Deplano, Y. Knapp, E. Bertrand, and E. Gaillard. Flow behaviour in an asymmetric compliant experimental model for abdominal aortic aneurysm. *J. Biomech.*, 40(11):2406–2413, 2007.
- [524] C.J. Mills, I.T. Gabe, J.H. Gault, D.T. Mason, J. Ross, E. Braunwald, and J.P. Shillingford. Pressure-flow relationships and vascular impedance in man. *Cardio-vasc. Res.*, 4(4):405–417, 1970.
- [525] K. Takizawa, T.E. Tezduyar, and T. Sasaki. Estimation of Element-Based Zero-Stress State in Arterial FSI Computations with Isogeometric Wall Discretization. In P. Wriggers and T. Lenarz, editors, *Biomedical Technology: Modeling, Experiments and Simulation*, pages 101–122. Springer International Publishing, Cham, 2018.
- [526] Y. Jin, A. Pepe, J. Li, C. Gsaxner, F. Zhao, J. Kleesiek, A.F. Frangi, and J. Egger. AI-based Aortic Vessel Tree Segmentation for Cardiovascular Diseases Treatment: Status Quo. *arXiv*, 2021. doi: 10.48550/ARXIV.2108.02998.
- [527] C. Dionysio, D. Wild, A. Pepe, C. Gsaxner, J. Li, L. Alvarez, and J. Egger. A cloud-based centerline algorithm for Studierfenster. In T.M. Deserno and B.J. Park, editors, *Medical Imaging 2021: Imaging Informatics for Healthcare, Research, and Applications*, volume 11601, pages 201–206, Online only, 2021. International Society for Optics and Photonics, SPIE.
- [528] N. Heller, F. Isensee, K.H. Maier-Hein, X. Hou, C. Xie, F. Li, Y. Nan, G. Mu, Z. Lin, M. Han, G. Yao, Y. Gao, Y. Zhang, Y. Wang, F. Hou, J. Yang, G. Xiong, J. Tian, C. Zhong, J. Ma, J. Rickman, J. Dean, B. Stai, R. Tejpaul, M. Oestreich, P. Blake, H. Kaluzniak, S. Raza, J. Rosenberg, K. Moore, E. Walczak, Z. Rengel, Z. Edgerton, R. Vasdev, M. Peterson, S. McSweeney, S. Peterson, A. Kalapara, N. Sathianathen, N. Papanikolopoulos, and C. Weight. The state of the art in kidney and kidney tumor segmentation in contrast-enhanced CT imaging: Results of the KiTS19 challenge. *Med. Image Anal.*, 67:101821, 2021.
- [529] T. Wick. Flapping and contact FSI computations with the fluid-solid interface-tracking/interface-capturing technique and mesh adaptivity. *Comput. Mech.*, 53(1):29–43, 2014.
- [530] M. Bukač, S. Čanić, R. Glowinski, J. Tambača, and A. Quaini. Fluid–structure interaction in blood flow capturing non-zero longitudinal structure displacement. *J. Comput. Phys.*, 235:515–541, 2013.
- [531] H. Baek and G.E. Karniadakis. A convergence study of a new partitioned fluid–structure interaction algorithm based on fictitious mass and damping. *J. Comput. Phys.*, 231(2):629–652, 2012.
- [532] S. Kniesburges, S.L. Thomson, A. Barney, M. Triep, P. Sidlof, J. Horacek, C. Brucker, and S. Becker. In Vitro Experimental Investigation of Voice Production. *Curr. Bioinform.*, 6(3):305–322, 2011.

-
- [533] M. de Oliveira Rosa, J.C. Pereira, M. Grellet, and A. Alwan. A contribution to simulating a three-dimensional larynx model using the finite element method. *J. Acoust. Soc.*, 114(5):2893–2905, 2003.
- [534] C. Tao, J.J. Jiang, and Y. Zhang. Simulation of vocal fold impact pressures with a self-oscillating finite-element model. *J. Acoust. Soc.*, 119(6):3987–3994, 2006.
- [535] S. Falk, S. Kniesburges, S. Schoder, B. Jakubaß, P. Maurerlehner, M. Echternach, M. Kaltenbacher, and M. Döllinger. 3D-FV-FE Aeroacoustic Larynx Model for Investigation of Functional Based Voice Disorders. *Front. Physiol.*, 12, 2021.

Monographic Series TU Graz

Structural Analysis

Wolfgang Moser

Transient coupled BEM-BEM and BEM-FEM analyses

2006

ISBN 978-3-902465-54-2

Tatiana Ribeiro

Elastoplastic Boundary Element Method with adaptive cell generation

2007

ISBN 978-3-902465-55-9

Christian Dünser

Simulation of Sequential Tunnel Excavation with the Boundary Element Method

2007

ISBN 978-3-902465-65-8

André Pereira

Coupling BEM/FEM for fluid-soil-structure interaction

2008

ISBN 978-3-85125-010-7

Gudrun Stettner

Dynamic behavior of semi-rigid frames considering SSI

2009

ISBN 978-3-85125-064-0

Klaus Thöni

The Boundary Element Method for Non-Linear Problems

2009

ISBN 978-3-85125-078-7

Plínio Glauber Carvalho dos Prazeres

Simulation of NATM Tunnel Construction with the BEM

2010

ISBN 978-3-85125-088-6

Monographic Series TU Graz

Structural Analysis

Katharina Riederer

Modelling of Ground Support in Tunnelling using the BEM

2010

ISBN 978-3-85125-090-9

Jürgen Zechner

Fast elasto-plastic BEM with hierarchical matrices

2012

ISBN 978-3-85125-233-0

Benjamin Marussig

Seamless Integration of Design and Analysis through Boundary Integral Equations

2016

ISBN 978-3-85125-440-2

Daniel Schöllhammer

Unified modelling and higher-order accurate analysis of curved shells and membranes

2021

ISBN 978-3-85125-813-4

Richard Schussnig

Generalised Newtonian Fluids in Cardiovascular Fluid-Structure Interaction

2023

ISBN 978-3-85125-943-8