Efficient Numerical Methods for Fluid-Structure Interaction

Efficient Numerical Methods for Fluid-Structure Interaction

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de Technische Universiteit Delft, op gezag van de Rector Magnificus prof. dr. ir. J.T. Fokkema, voorzitter van het College voor Promoties, in het openbaar te verdedigen op dinsdag 14 juni 2005 om 15:30 uur

 door

Christian MICHLER

Diplom-Ingenieur des Chemieingenieurwesens, Technische Universität Clausthal, Duitsland

geboren te Celle, Duitsland

Dit proefschrift is goedgekeurd door de promotor:

Prof. dr. ir. R. de Borst

Samenstelling promotiecommissie:

	TT •••
Rector Magnificus,	Voorzitter
Prof. dr. ir. R. de Borst,	Technische Universiteit Delft, promotor
Prof. dr. ir. B. Koren,	Technische Universiteit Delft
Prof. dr. ir. D.J. Rixen,	Technische Universiteit Delft
Prof. dr. ir. H. Deconinck,	Von Karman Institute for Fluid Dynamics, België
Prof. DrIng. E. Ramm,	Universität Stuttgart, Duitsland
Prof. DrIng. W.A. Wall,	Technische Universität München, Duitsland
Dr. ir. E.H. van Brummelen,	Technische Universiteit Delft

Dr. ir. E.H. van Brummelen heeft als begeleider in belangrijke mate aan de totstandkoming van het proefschrift bijgedragen.

Copyright © 2005 by C. Michler

All rights reserved. No part of this material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any other information storage and retrieval system, without written permission from the copyright owner.

Printed in The Netherlands by PrintPartners Ipskamp, Enschede

ISBN 90-9019533-5

Keywords: Fluid-structure interaction, partitioning, subiteration, GMRES, Newton-Krylov methods, efficiency, energy conservation, space/time finite-element method

Contents

Preface		$\mathbf{i}\mathbf{x}$		
1	Introduction			1
	1.1	Motivation		1
	1.2	Outline		4
2	Mathematical description of fluid-structure interaction		9	
	2.1	Introductio	n	9
	2.2	Kinematic	description of a continuum	10
		2.2.1	Eulerian description	10
		2.2.2	Lagrangian description	12
		2.2.3	Arbitrary Lagrangian-Eulerian formulation	13
		2.2.4	Space/time formulation	14
	2.3	Governing	equations for fluid flows	15
		2.3.1	Classical formulation	15
		2.3.2	Variational space/time formulation	17
	2.4 Governing equations for structures		equations for structures	18
		2.4.1	Classical formulation	18
		2.4.2	Variational space/time formulation	21
	2.5	Interface co	onditions	21
		2.5.1	Kinematic interface conditions	21
		2.5.2	Dynamic interface condition	22
	2.6	Aggregated	variational problem	22
	2.7	Discretizati	on of the variational problem	23
		2.7.1	Space/time finite-element discretization	23
		2.7.2	Implications of the discretization	24
	2.8	Concluding	remarks	25
3	Conservation under discretization			27
	3.1	Introductio	n	27
	3.2	Problem statement		28
		3.2.1	The piston problem $\ldots \ldots \ldots \ldots \ldots \ldots \ldots$	29

		3.2.2	Variational formulation	31
	3.3	System	behaviour	32
		3.3.1	Conservation at the fluid-structure interface	32
		3.3.2	Linearized-system behaviour	33
	3.4	Conserv	ation properties under discretization	34
		3.4.1	Discrete approximations	35
		3.4.2	Conservation in the discretized subsystems	35
		3.4.3	Conservation for compatible discretizations	36
		3.4.4	Conservation for incompatible discretizations	37
		3.4.5	Coupling of trapezoidal with time-discontinuous Ga-	
			lerkin time integration	38
	3.5	Numeric	cal experiments	40
		3.5.1	Experimental setup	40
		3.5.2	Conservative versus non-conservative discretization	43
		3.5.3	Coupling of trapezoidal with time-discontinuous Ga-	
			lerkin time integration	45
	3.6	Conclud	ing remarks	47
4	Effic	iency of c	conventional solution methods	51
	4.1	Introduc	ction	51
	4.2	Monolit	hic versus partitioned solution methods	52
		4.2.1	Monolithic methods	53
		4.2.2	Strongly-coupled partitioned methods	55
		4.2.3	Loosely-coupled partitioned methods	56
		4.2.4	Prediction techniques	57
		4.2.5	Strongly-coupled versus loosely-coupled methods	58
		4.2.6	Improving the efficiency of strongly-coupled methods	59
	4.3 Numerical experiments			60
	1.0	4.3.1	Experimental setup	60
		432	Investigation of stability properties	61
		433	Investigation of accuracy properties	63
		4.3.4	Comparison of efficiency	64
		435	Improving the efficiency of strongly-coupled methods	65
	4.4	Conclud	ing remarks	66
5	Inte	rface-GM	$\operatorname{RES}(\mathbf{R})$ acceleration of subiteration (69
0	5.1	Introduc	tion	69
	5.2	Interface	e-GMRES(R) acceleration	71
	0.2	5 2 1	Subiteration: Interface fixed-point iteration	71
		5.2.2	Interface-GMRES acceleration of subiteration	71
		523	Reuse of Krylov vectors: Interface-GMRESR	75
	5.3	Numeric	al experiments	76
	0.0	531	Convergence of Interface-GMRES(R)	77
		539	Rause of Krylov vectors in subsequent time stops	89
		0.0.4	recuse of region vectors in subsequent time steps .	04

	5.4	Concluding remarks	86
6	Error	-amplification analysis of Interface-GMRES(R)	89
	6.1	Introduction	89
	6.2	Algebraic problem statement	90
		6.2.1 The system of fluid-structure-interaction equations	91
		6.2.2 Error-amplification of inexact Newton methods	92
	6.3	Analysis of subiteration	93
		$6.3.1 The subiteration method \ldots \ldots \ldots \ldots \ldots \ldots$	93
		6.3.2 Error-amplification analysis	94
		6.3.3 Nonnormality of subiteration	96
		6.3.4 Subiteration preconditioning	97
	6.4	Analysis of Interface-GMRES(R)	99
		6.4.1 Subiteration-preconditioned GMRES	100
		6.4.2 Reuse of Krylov vectors	104
		6.4.3 Error-amplification analysis	105
		6.4.4 Analysis of the reuse of Krylov vectors	107
		6.4.5 GMRES convergence bounds	107
	6.5	Numerical experiments	108
		6.5.1 Experimental setup	108
		6.5.2 Convergence of subiteration	110
		6.5.3 GMRES convergence bounds	112
		6.5.4 Analysis of the reuse of Krylov vectors	113
	6.6	Concluding remarks	115
7	Asses	sment of Interface-GMRES(R) on the panel problem	119
	7.1	Introduction	119
	7.2	Problem statement	120
		7.2.1 The panel problem	120
		7.2.2 System behaviour	122
	7.3	Numerical experiments	123
		7.3.1 Experimental setup	124
		7.3.2 Assessment of Interface-GMRES(R) convergence .	126
		7.3.3 Loosely-coupled versus strongly-coupled methods .	127
		7.3.4 Effect of the initial conditions	131
	7.4	Concluding remarks	132
8	Conc	lusions	135
9	Reco	mmendations	139
Bib	liograp	hy	141
٨	Con	provided and the state internetion	140
A	A.1	Assessment of energy conservation	149 149
	· -		- 10

	A.2	Assessment of momentum conservation $\ldots \ldots \ldots \ldots$	150
Sub	ject In	dex	153
Sun	nmary		155
San	nenvatt	ing	157
Cur	riculun	n Vitae	161

Preface

This thesis summarizes the research that I carried out at the Faculty of Aerospace Engineering at Delft University of Technology.

Many people have contributed to the realization of this thesis. I would like to thank my promotor René de Borst for his unconditional support, trust and encouragement during this work. I am also very much indebted to my supervisor Harald van Brummelen for his careful guidance and great support. I really enjoyed our pleasant collaboration and the many fruitful discussions.

Moreover, I wish to express my sincere gratitude to Steven Hulshoff for his advice, and for his support in code development. I also would like to acknowledge numerous helpful discussions and support from Edwin Munts, Erik Jan Lingen and Martijn Stroeven regarding implementation and coding. Furthermore, I am grateful to Sergio Turteltaub and Hester Bijl for many valuable discussions.

Our secretary Carla Roovers deserves special credits for her kind support and invaluable help. I also would like to thank Harold Thung for his system-administrative support.

I enjoyed the pleasant working atmosphere and company of my former office mates Doobo Chung, Alexei Kononov, Alberto Bordallo, Otto Heeres and Marcela Cid. I also would like to acknowledge the many interesting conversations with Michiel Hagenbeek which significantly helped me in learning Dutch. Moreover, I am grateful to Kris van der Zee for numerous stimulating discussions. The supervision of Rob in 't Groen and the interaction with Clemens Verhoosel and Ido Akkerman provided complementary insights for my research which I thankfully acknowledge.

I really appreciated the inspiring working environment and the good atmosphere in the Engineering Mechanics group, and for this I thank all present and former colleagues, in particular Joris Remmers, Akke Suiker, Ellen Kuhl and Akihiro Matsuda.

I am also grateful to the many friends I have made here during my stay in Delft, especially to my former housemate Weihong He and to Yadira Cordero, Anna Dall'Acqua, Carmen Lai, Pavel Paclik, Jun Chen and Alex and Huihua van der Klij.

I also would like to thank my parents and my brother for their strong support during this work. Finally, I wish to deeply thank Ana for her understanding, encouragement and continuous support.

Christian Michler Delft, April 2005

Chapter 1 Introduction

1.1 Motivation

The interaction between a fluid and a structure occurs in a wide variety of physical systems and engineering applications. In aerospace engineering, for instance, the interaction between the airflow and the aircraft structure can cause structural oscillations of increasing amplitude, known as flutter. Buffet of control surfaces and fins due to (self-induced) turbulence is another relevant example. These aeroelastic phenomena can ultimately result in failure of the structure; see Fig. 1.1 on the following page for an illustration. In civil engineering, wind-induced vibrations of high-rise buildings and bridges can cause the collapse of the construction. The most prominent example is the failure of the Tacoma Narrows Bridge in 1940; see Fig. 1.2. These examples drastically highlight the impact of fluid-structure-interaction phenomena on the safety and reliability of constructions. Other relevant application areas in which fluid-structure interaction plays a crucial role are, for instance, inflatable structures such as airbags and parachutes, and artificial heart valves; see, e.g., Refs. [63, 67].

To determine the effects of fluid-structure interaction for a given system, engineering design often involves extensive experimental testing. However, experiments may be costly, time consuming, especially for large design-parameter spaces, and in some cases even infeasible, e.g., in case of hazards. As numerical models and techniques have matured over the last decades to deliver more accurate predictions, and with the advent of increasing computing power for affordable prices, numerical simulation has become more and more established in the design process to support or even replace experimental testing. Consequently, there is a growing demand for the accurate and efficient numerical solution of fluid-structure interaction problems arising in the various engineering disciplines.

To solve a fluid-structure-interaction problem on a computer requires first



Figure 1.1: F-16 ventral fin damage due to buffet; reproduced from [11] with permission.



Figure 1.2: The collapse of the Tacoma Narrows Bridge due to wind-induced vibrations.

of all a mathematical model, which is generally expressed in terms of partial differential equations. To make these equations amenable to numerical treatment, discretization techniques such as the finite-element or the finite-volume method are applied. Such discretization methods translate the continuum equations into a system of discrete algebraic equations. From these algebraic equations, discrete approximations to the solution can then be extracted by means of a numerical solution algorithm; see Fig. 1.3 for an illustration. Depending on the application area, this procedure is commonly referred to as *Computational Fluid Dynamics* (CFD) or *Computational Structure Dynamics* (CSD). The computation of fluid-structure interactions involves both disciplines, CFD and CSD.

The computational challenges in the numerical solution of fluid-structure interaction problems are closely linked to the generic features of fluid-structure-interaction phenomena. Fluid-structure interaction constitutes a *multi-physics problem* on account of the interaction between subsystems with very different physics. These subsystems are interconnected through kinematic and dynamic



Figure 1.3: Steps in the numerical solution of a fluid-structure-interaction problem (and reference to the corresponding chapters in this thesis).

conditions at their common *interface*. The position of the fluid-structure interface is not known a priori, but it forms part of the solution. Therefore, fluid-structure interaction can also be classified as a *free-boundary problem*. The inherent interconnection requires, in principle, the simultaneous solution of fluid and structure. However, such simultaneous treatment results in the loss of software modularity. In the scientific community, different solution methods are controversially advocated to tackle this problem; see the special journal issue [51] for an overview. However, the computation of fluid-structure interactions remains challenging.

Many of the difficulties in the computation of fluid-structure-interaction problems can be traced to the inherently different length and time scales of the fluid and structure subsystems. These disparate scales in fluid and structure translate into different resolution requirements in the discrete numerical model. This typically results in non-matching meshes and different orders of approximation at the fluid-structure interface. Under such an *incompatible discretization* of the fluid and structure it is difficult to maintain the conservation properties of the continuum system at the fluid-structure interface. However, the correct representation of the energy transfer at the interface is crucial for the correct prediction of physical instabilities such as flutter.

Moreover, the inherently different length and time scales of the fluid and structure subsystems typically render the discretized fluid-structure system severely ill-conditioned. This considerably slows down the convergence of iterative methods for the simultaneous solution of the fluid and structure.

The above elaboration supports the truism that the computation of fluidstructure interaction is profoundly more difficult than the computation of the fluid and the structure separately. Most state-of-the-art techniques obscure if not violate energy conservation at the interface. Moreover, the customary solution techniques often suffer from a lack of robustness and are computationally inefficient. These flaws of conventional numerical techniques provide the motivation for the research presented in this thesis, which is concerned with conservative discretization methods and robust and efficient solution techniques for fluid-structure interaction.

1.2 Outline

The contents of this thesis are organized as follows.

Chapter 2 presents the mathematical formulation of fluid-structure interaction. We elaborate the different reference frames that can be used for the kinematic description of a continuum fluid and structure. We then introduce the governing equations for the fluid and the structure, and the kinematic and dynamic interface conditions which describe the interaction of the fluid and the structure at their common interface. To identify the generic features of fluid-structure-interaction problems, we rephrase the governing equations into a variational formulation in space/time. As this formulation does not make any stipulations on the particular model for the fluid, the structure or the interface conditions, it comprises *any* fluid-structure-interaction problem. Finally, we elaborate the discretization of the aggregated variational statement by the space/time finite-element method, and we discuss the particularities for fluid-structure-interaction problems.

In Chapter 3, we investigate the conservation properties of the discrete numerical model at the fluid-structure interface. To this end, we establish the conservation properties of the continuum system and of its discretization by the finiteelement method. The conservation properties of the continuum system can be lost under discretization. It appears that energy conservation at the interface is only trivially maintained under restrictive compatibility conditions on the approximation spaces of fluid and structure, i.e., matching meshes and identical orders of interpolation at the interface. These conditions are prohibitive for practical use. We then consider an approach based on coincidence and orthogonality conditions which enables conservation also for incompatible discretizations. To assess the implications that particular pairings of fluid and structure discretizations can have on the conservation properties at the interface, we consider an alternative pairing of discretization methods, viz., the trapezoidal method for the structure with a time-discontinuous Galerkin method for the fluid. We show that the trapezoidal method is generally not conservative for forced structures and, hence, it does not admit a conservative fluid-structure coupling. In our numerical experiments, we demonstrate that a method which maintains conservation at the interface yields a much more accurate solution than a non-conservative method, at the same computational expense. Conversely, violating the conservation properties can induce numerical instability. The presented results clearly warrant a preference for conservative discretizations.

In Chapter 4, we assess the efficiency of conventional solution methods for fluid-structure-interaction problems. Numerical solution methods for fluid-structure interaction typically employ partitioning, i.e., fluid and structure equations are separately integrated in time subject to complementary partitions of the interface conditions. Thus, a simultaneous solution of the fluid and structure equations is circumvented. This process is also referred to as *subiteration*, as it operates within a time step. If multiple subiterations are carried out per time step, the solution method is referred to as a *strongly-coupled* partitioned method. If, on the other hand, only a single subiteration is carried out per time step, the solution method is called a *loosely-coupled* partitioned method. This implies that the computational cost per time step is higher for strongly-coupled methods. On the other hand, only strongly-coupled methods can resolve the aggregated fluid-structure equations, whereas loosely-coupled methods induce a numerical evaluation error, which impedes conservation and restricts the admissible time-step size for reasons of stability and accuracy. An immediate question is then whether for a given accuracy the higher computational cost of strongly-coupled methods is compensated for by the larger time steps they can afford in comparison with loosely-coupled methods. In our numerical experiments, we compare loosely-coupled and stronglycoupled methods in terms of stability, accuracy and efficiency, where we conceive efficiency as the ratio of accuracy to computational cost. We show that the numerical evaluation error incurred by loosely-coupled methods can compromise the stability, accuracy and efficiency of the method. Strongly-coupled methods are superior to loosely-coupled methods provided that the underlying discretization maintains the conservation properties. Indeed, our results indicate that the higher computational cost of strongly-coupled methods is only justified by a greater accuracy if the underlying discretization is conservative. Our results therefore refute a common belief that for the same accuracy strongly-coupled methods are generally more expensive than loosely-coupled methods with a reduced time-step size.

In Chapter 5, we devise a novel solution method for fluid-structure interaction. Although subiteration is an apt solver for many problems, it suffers from three essential drawbacks. Firstly, subiteration is only conditionally stable. Secondly, transient divergence can precede asymptotic convergence due to nonnormality. Thirdly, subiteration is in general inefficient because of its inability to reuse information. Indeed, subiteration operates in a sequential time-integration process and, thus, solves a sequence of similar problems. However, it cannot exploit this property by reusing information from previously solved similar problems, for instance, for preconditioning purposes. To overcome the aforementioned disadvantages, we propose to combine subiteration with GMRES acceleration. We establish the theoretical background and the algorithmic aspects of the combined subiteration/GMRES method on the basis of the generic formulation of fluid-structureinteraction problems from Chapter 2, which implies that the proposed method is in principle generic and, thus, applicable to any fluid-structure system. The combined method is based on the observation that subiteration can be conceived as a fixedpoint iteration for the interface position. Therefore, the GMRES acceleration can be confined to the interface position. We refer to this process as Interface-GMRES. Because the Krylov vectors need only contain the discrete representation associated with the interface position, the GMRES acceleration requires only negligible computational resources. Another asset is that the acceleration of subiteration allows for optional reuse of Krylov vectors in subsequent invocations of GMRES, which we refer to as GMRESR. Such reuse can considerably enhance the efficiency of the method. The implementation of the Interface-GMRES(R) method in codes which use subiteration as a solver is straightforward, because the method retains the modularity of the underlying subiteration method. We illustrate the effectiveness of the proposed method through numerical experiments on the prototypical piston model problem. Detailed convergence studies and a comparison to standard subiteration show that the Interface-GMRES(R) method is much more robust, and that it converges even if subiteration itself diverges. Our results also demonstrate that the accelerated method is much more efficient than subiteration, and that the reuse of Krylov vectors can yield considerable computational savings.

Chapter 6 analyses the linear-algebra aspects of the Interface-GMRES(R)method on the basis of properties of the error-amplification matrix of the aggregated system. This complements the exposition from Chapter 5. By virtue of the linear-algebra setting, it is possible to derive precise expressions for the erroramplification properties of subiteration separately, and of subiteration combined with GMRES, with and without the reuse option. We show that subiteration condenses errors into a low-dimensional subspace which can be associated with the interface degrees-of-freedom. The rank of the error-amplification matrix associated with subiteration is at most equal to the dimension N of the approximation space of the interface variables. This implies that a Krylov method terminates in at most N steps, independent of the choice of the acceleration space, e.g., aggregated variables, structure variables, or interface variables. However, the acceleration on the interface variables is the most efficient, because the computational cost and the storage required by the Krylov acceleration itself increase with the dimension of the acceleration space. The linear-algebra setting enables a clear explanation of the relation between the *local* GMRES acceleration (i.e., on the interface degreesof-freedom), and the *qlobal* error-amplification properties (i.e., for the aggregated system). Moreover, the nonnormality of subiteration, and its implications for the combined subiteration/GMRES method, can be traced immediately to properties of the error-amplification matrix. Such nonnormality can degrade the sharpness of GMRES convergence bounds. Next, we show that the subiteration iterates span a Krylov space corresponding to a preconditioned aggregated system. The properties of the subiteration-preconditioned GMRES method are then considered and, in particular, the relation between GMRES convergence and nonnormality of the subiteration preconditioner is addressed. Numerical experiments on a model problem illustrate the developed theory. We analyse the convergence of the respective methods in terms of spectral radius, matrix norms and sharp convergence upper bounds. Moreover, we investigate the effect of the reuse of Krylov vectors on the error-amplification properties of the subiteration-preconditioned GMRES method to assess the limitations on the reuse option.

In Chapter 7, we assess the Interface-GMRES(R) method on an alternative model problem, viz., the panel problem, to demonstrate the versatility of the solution method. The panel problem is distinctly different from the piston problem considered in Chapter 5. In contrast to the one-dimensional piston problem, for the two-dimensional panel problem the interface degrees-of-freedom pertain to both space and time. This distinction is relevant for further testing of the Interface-GMRES(R) method, as the method operates on the interface degrees-offreedom. Another relevant feature that distinguishes the panel problem from the piston problem relates to the aspect of parameter-dependent stability behaviour. Many fluid-structure-interaction problems can display instabilities such as flutter and divergence for certain parameter settings, whereas other parameter settings yield stable behaviour. The piston problem does not have this property, as it is (marginally) stable for all parameter settings. The panel problem, on the other hand, can exhibit physical instability such as flutter and divergence. An investigation of the convergence behaviour of the Interface-GMRES(R) method for different stability regimes is therefore relevant. In our numerical experiments on the panel problem, we examine the convergence behaviour of Interface-GMRES(R), assess its robustness and efficiency, and compare its performance to standard subiteration. We investigate the effect of changes in the solution behaviour due to flutter on the convergence of the Interface-GMRES(R) method and on the effectiveness of reuse of the Krylov space. Moreover, we study the influence of the initial conditions on the system behaviour and on the convergence of Interface-GMRES(R). Finally, we consider loosely-coupled and strongly-coupled partitioned solution methods for a stable system and a system undergoing instability in the form of flutter.

Chapters 8 and 9 contain concluding remarks and suggestions for future research, respectively.

Parts of the Chapters 2 to 6 have been published before. Their content is based on the following journal publications:

- C. MICHLER, E.H. VAN BRUMMELEN, S.J. HULSHOFF, AND R. DE BORST, The relevance of conservation for stability and accuracy of numerical methods for fluid-structure interaction, Comput. Methods Appl. Mech. Engrg. 192 (2003), nos. 37–38, 4195–4215,
- C. MICHLER, S.J. HULSHOFF, E.H. VAN BRUMMELEN, AND R. DE BORST, A monolithic approach to fluid-structure interaction, Computers and Fluids **33** (2004), nos. 5–6, 839–848,
- C. MICHLER, E.H. VAN BRUMMELEN, AND R. DE BORST, An interface Newton-Krylov solver for fluid-structure interaction, Int. J. Num. Meth. Fluids 47 (2005), nos. 10–11, 1189–1195, Special issue containing the Proceedings of the Eighth ICFD Conference on Numerical Methods for Fluid Dynamics (Oxford, U.K., March 29 - April 1, 2004),
- C. MICHLER, E.H. VAN BRUMMELEN, AND R. DE BORST, Error-amplification analysis of subiteration-preconditioned GMRES for fluid-structure interaction, Comput. Methods Appl. Mech. Engrg. - (2005), -, (Accepted for publication),
- E.H. VAN BRUMMELEN, C. MICHLER, AND R. DE BORST, Interface-GM-RES(R) acceleration of subiteration for fluid-structure-interaction problems, (Submitted for publication), (2005). Also available as: Technical Report DACS-05-001, Delft University of Technology, The Netherlands. Available at : http://www.em.lr.tudelft.nl/downloads/DACS-05-001.pdf

and on the following conference proceedings:

- C. MICHLER, S.J. HULSHOFF, E.H. VAN BRUMMELEN, H. BIJL, AND R. DE BORST, *Space-time discretizations for fluid-structure interaction*, Proceedings of the Fifth World Congress on Computational Mechanics (WCCM V) (H.A. Mang, F.G. Rammerstorfer, and J.Eberhardsteiner, eds.), Vienna University of Technology, Austria, July 7-12, 2002,
- C. MICHLER, E.H. VAN BRUMMELEN, S.J. HULSHOFF, AND R. DE BORST, Efficiency of monolithic and partitioned methods for fluid-structure interaction, Proceedings of the International Forum on Aeroelasticity and Structural Dynamics (IFASD) (Amsterdam, The Netherlands, June 4-6, 2003), (on CD ROM),
- C. MICHLER, AND E.H. VAN BRUMMELEN, *GMRES acceleration of subiteration for fluid-structure interactions*, Moving Boundaries 2003: Proceedings of the Seventh International Conference on Computational Modelling of Free and Moving Boundary Problems (Santa Fe, New Mexico, USA, November 4-6 2003) (A.A. Mammoli and C.A. Brebbia, eds.), WIT Press, Wessex, 2003, 109–118.

Chapter 2 Mathematical description of fluid-structure interaction

2.1 Introduction

In this chapter, we present the mathematical model for fluid-structure interaction. This model comprises the governing equations for the fluid and the structure, augmented by interface conditions. These interface conditions describe the interaction of the fluid and the structure at their common interface, and they can be separated into kinematic, dynamic and thermal interface conditions. As for most applications the thermal interaction can be neglected, we shall not consider it in our work.

The descriptions of a continuum fluid and a continuum structure customarily employ different reference frames, viz., the Eulerian and the Lagrangian reference frame, respectively. When fluid and structure interact, the induced movement of the fluid domain conflicts with the Eulerian fluid description. To describe a fluid on a moving domain, two methodologies can be distinguished, viz., the arbitrary Lagrangian-Eulerian formulation and the space/time formulation. Moreover, the difference in fluid and structure reference frames necessitates a transformation of interface quantities from one reference frame into the other.

To identify the generic features of fluid-structure-interaction problems, we rephrase the governing equations into a variational statement in space/time in conformity with Ref. [71]. This formulation is generic, i.e., it does not make any stipulations on the particular model for the fluid or the structure. Therefore, it comprises *any* fluid-structure-interaction problem. As generic features of fluid-structure-interaction problems we identify the interconnection of a fluid and a structure subsystem by kinematic and dynamic interface conditions and, moreover, the free-boundary character of the interface.

Since in this thesis we are concerned with the numerical solution of fluidstructure-interaction problems, we consider the discretization of the aggregated variational statement and, in particular, the relation between the generic features of fluid-structure interaction and the discretization. We discuss the advantages of the space/time finite-element method over conventional methods that discretize space and time separately, i.e., by different discretization methods. In contrast to conventional methods, the space/time finite-element method can easily handle moving meshes, and it inherently satisfies the so-called *Discrete Geometric Conservation Law* (DGCL). The DGCL expresses the requirement for a numerical scheme on a moving mesh that the change in a control volume is equal to the area swept by the moving boundaries. Compliance with the DGCL is relevant for the stability and accuracy of a numerical method on a moving mesh; see, e.g., Refs. [17, 35] and also the review article [16].

The contents of this chapter are organized as follows: Section 2.2 briefly recalls the concepts of Eulerian, Lagrangian and arbitrary Lagrangian-Eulerian description of a continuum and, moreover, compares the latter to the space/time description. In Sections 2.3–2.5, we present the governing equations for the fluid, the structure and the interface conditions, both in strong form and in variational form. In Section 2.6, we present the aggregated variational statement, and in Section 2.7, we discuss its discretization by the space/time finite-element method. Section 2.8 contains concluding remarks.

2.2 Kinematic description of a continuum

To set the reference frames for the ensuing description of the fluid, the structure and the interface conditions, we briefly recall the classical Eulerian and Lagrangian description of a continuum; a detailed exposition can be found in any standard textbook on continuum mechanics, for instance, in [60, 64]. We elaborate that, if the fluid domain moves, for instance due to fluid-structure interaction, the movement of the domain renders the Eulerian fluid description involved. We consider two approaches for conveniently describing a fluid on a moving domain, viz., the customary arbitrary Lagrangian-Eulerian formulation and the space/time formulation.

2.2.1 Eulerian description

In the Eulerian description, the reference frame is fixed with respect to a specific location in space. Any property of the continuum is then described as a function of a so-called *spatial coordinate*, i.e., a specific location in space through which different material points pass in subsequent instances. Accordingly, the Eulerian description is commonly also referred to as the *spatial description*. Because the spatial description dissociates the reference frame from the material point, the continuum moves relative to the reference frame, which gives rise to convective terms in the formulation. By virtue of the fixed reference frame, the Eulerian description is suitable for large deformations and distortions of the continuum such as arise in, for instance, fluid flow. The fluid equations are therefore customarily formulated

in an Eulerian reference frame. However, if the domain of the continuum moves or deforms, a fixed reference frame becomes inconvenient, because it requires the explicit tracking of the domain boundary.

To illustrate this concept, we shall consider the simple example of the advection of a continuum. Let x and t denote spatial and temporal coordinates, respectively, $\phi(x,t)$ a scalar material property and a(x,t) the advection velocity. The homogeneous advection of ϕ is then described by

$$\frac{\partial \phi}{\partial t} + a \frac{\partial \phi}{\partial x} = 0, \qquad c_0 t \le x \le 1 + c_1 t, \quad 0 \le t \le T.$$
(2.1a)

Specifically, in (2.1a), we consider the case in which the domain is moving and deforming in time, with $c_0 \in \mathbb{R}_+ \setminus \{0\}, c_1 \in \mathbb{R}_- \setminus \{0\}$ denoting the constant velocity of the left and right boundary, respectively; see Fig. 2.1 for an illustration. We shall assume that Eq. (2.1a) is complemented with suitable initial and boundary conditions. For convenience and later use, we postulate that the advection velocity *a* conforms to the Burger's equation

$$\frac{\partial a}{\partial t} + a \frac{\partial a}{\partial x} = 0, \qquad c_0 t \le x \le 1 + c_1 t, \quad 0 \le t \le T$$
 (2.1b)

and satisfies on the domain boundaries

$$a(c_0t,t) = c_0, \quad a(1+c_1t,t) = c_1, \qquad 0 \le t \le T,$$
(2.1c)

which translates into the impermeability of the boundaries. The length of the time interval, T, is chosen such that shocks in the Burger's equation (2.1b) do not occur.

The presented example serves to illustrate certain distinguishing features of the Eulerian description which pertain to the presence of convective terms in the formulation and, moreover, to the need for an explicit tracking of the moving domain boundary; see Eq. (2.1a) and also Fig. 2.1.



Figure 2.1: Illustration of the Eulerian description: spatial domain of the continuum (grey), Eulerian coordinate isolines denoted by (--).

2.2.2 Lagrangian description

In the Lagrangian description, the reference frame is connected to the material point, and any property of the continuum is described as a function of a so-called *material coordinate*. Accordingly, the Lagrangian description is commonly also referred to as the *material description*.

We reconsider the example of the advection equation. To express Eq. (2.1a) in Lagrangian form, we need to establish the relation between the spatial coordinate x and the material coordinate denoted by y. By definition of the material description, the material coordinate y is constant along the path of a material point. For the considered example, this particle path is determined by the advection velocity a. Let us therefore assume that spatial and material coordinates are interrelated by x = y + at. Stipulating that time is invariant under a change of reference frame, we use the same symbol to denote time in either reference frame. The transformation from the material to the spatial domain is then given by

$$(y,t) \longmapsto (x,t) = (\varphi(y,t),t).$$
 (2.2a)

To transform the advection equation (2.1a) from the spatial domain into the material domain, we require the derivative $\partial(y,t)/\partial(x,t)$, which for the considered example can be expressed as

$$\frac{\partial(y,t)}{\partial(x,t)} = \begin{pmatrix} 1 - t\frac{\partial a}{\partial x} & -a - t\frac{\partial a}{\partial t} \\ 0 & 1 \end{pmatrix}.$$
 (2.2b)

The advection equation (2.1a) then translates into

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} - t\frac{\partial\phi}{\partial y}(\frac{\partial a}{\partial t} + a\frac{\partial a}{\partial x}) = 0, \qquad 0 \le y \le 1, \quad 0 \le t \le T, \qquad (2.3a)$$

where we have invoked Eq. (2.1c). Subject to the postulation on *a* stated by Eq. (2.1b), Eq. (2.3a) can be simplified to

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = 0, \qquad 0 \le y \le 1, \quad 0 \le t \le T.$$
(2.3b)

In Eq. (2.3), $d/dt := \partial/\partial t|_y = \partial/\partial t|_x + a\partial/\partial x$ denotes the material time derivative, i.e., the time derivative at a fixed coordinate y. Note that there are no convective terms in Eq. (2.3b), because the reference frame is connected to the material point. Moreover, by virtue of the transformation (2.2), Eq. (2.3b) is expressed on a fixed reference domain (y, t); see also Fig. 2.2 for an illustration. This renders the Lagrangian description attractive for problems with moving boundaries and interfaces. However, to prevent ill-conditioning of the transformation φ , it is essential that the continuum undergoes only relatively small deformations. For most problems involving structures, this is indeed the case. Therefore, the Lagrangian description has become customary in structural mechanics. If, on the other hand, the continuum undergoes very large deformations or distortions, the transformation becomes ill-conditioned. This downside of the Lagrangian description is also indicated in Fig. 2.2 by the focusing of the coordinate isolines. It is easily conceivable that for multiple space dimensions this can lead to entanglement and, hence, to singularity of the mapping φ .



Figure 2.2: Illustration of the Lagrangian description: mapping from spatial to material domain, Lagrangian coordinate isolines denoted by (--).

2.2.3 Arbitrary Lagrangian-Eulerian formulation

To overcome the respective restrictions of the Eulerian and the Lagrangian description, the arbitrary Lagrangian-Eulerian (ALE) formulation has been developed in, among others, Refs. [13, 49] in the finite-element and finite-difference context, respectively. In the ALE description, the reference frame can be moved arbitrarily. Any material property is then described with respect to a so-called *referential coordinate*. To express the advection equation (2.1a) in ALE form, we introduce the referential coordinate, ξ , and the velocity of the reference frame, $\beta(x,t)$. Stipulating the invariance of the temporal coordinate, we interrelate spatial and referential coordinates by $x = \xi + \beta t$. The transformation from the referential to the spatial domain is then given by

$$(\xi, t) \longmapsto (x, t) = (\psi(\xi, t), t).$$
(2.4a)

To transform the advection equation (2.1a) from the spatial domain into the referential domain, we require the derivative

$$\frac{\partial(\xi,t)}{\partial(x,t)} = \begin{pmatrix} 1 - t\frac{\partial\beta}{\partial x} & -\beta - t\frac{\partial\beta}{\partial t} \\ 0 & 1 \end{pmatrix}.$$
 (2.4b)

The advection equation (2.1a) then translates into

$$\frac{\partial \phi}{\partial t}\Big|_{\xi} + (a-\beta)\frac{\partial \phi}{\partial \xi} - t\frac{\partial \phi}{\partial \xi}(\frac{\partial \beta}{\partial t} + a\frac{\partial \beta}{\partial x}) = 0,$$

$$(c_0 - \beta)t \le \xi \le 1 + (c_1 - \beta)t, \quad 0 \le t \le T, \quad (2.5)$$

where the time derivative is taken at a fixed referential coordinate. Two special instances of the ALE formulation can be distinguished: Upon setting $\beta = 0$ in Eqs. (2.4) and (2.5), the derivative $\partial(\xi, t)/\partial(x, t)$ assumes the form of the identity, and Eq. (2.5) reverts to Eq. (2.1a). Thus, the Eulerian description is recovered. If, on the other hand, we set $\beta = a$, Eqs. (2.4) and (2.5) revert to Eqs. (2.2) and (2.3), respectively, and we recover the Lagrangian description. To combine the merits of both Eulerian and Lagrangian description, a suitable choice for β is to set β identical to a on the moving domain boundary, and to set β in the domain interior such that an even distribution of coordinate isolines is obtained. This allows for a convenient description of the moving domain on a fixed reference domain and, at the same time, prevents the entanglement of coordinate isolines; see Fig. 2.3. Therefore, the ALE description has become standard for problems on moving domains in which the continuum undergoes large deformations and distortions, such as fluid flows with a moving boundary or free-surface flows; see, e.g., Ref. [8] and also the review article [14].



Figure 2.3: Illustration of the arbitrary Lagrangian-Eulerian description: mapping from spatial to referential domain, ALE coordinate isolines denoted by (--).

2.2.4 Space/time formulation

In the space/time formulation, any property of the material is expressed as a function of a space/time coordinate; see, e.g., Refs. [37, 62, 65]. To formulate the space/time description of a problem, both the Eulerian and the Lagrangian

description can be adopted as a starting point. For exemplification, we consider the Eulerian form of the advection equation (2.1a) which is expressed in a space/time reference frame by

$$\frac{\partial \phi}{\partial t} + a \frac{\partial \phi}{\partial x} = 0 \quad \text{on} \quad \Omega \,, \tag{2.6}$$

where $\Omega := \{(x,t) \in \mathbb{R}^2 : c_0 t \leq x \leq 1 + c_1 t, 0 \leq t \leq T\}$ denotes the (Eulerian) space/time domain of the continuum. Eq. (2.6) is virtually identical to Eq. (2.1a). To elucidate the difference between the space/time and the standard Eulerian description, we cast Eq. (2.6) in variational form by multiplication with a suitable test function and integration over the space/time domain

$$\int_{\Omega} \lambda \left(\frac{\partial \phi}{\partial t} + a \frac{\partial \phi}{\partial x} \right) = 0, \qquad (2.7)$$

where λ denotes a suitable test function. In contrast, in the classical Eulerian, Lagrangian and ALE description, the variational statement is commonly expressed only over the spatial domain. The difference between the space/time and the classical descriptions is, admittedly, subtle. In fact, the descriptions can be shown to be identical under the appropriate mapping. More precisely, an equation expressed in the Eulerian space/time domain yields the classical ALE form when mapped on to the referential space/time domain; see Ref. [37].

By virtue of expressing the variational statement over the space/time domain, a movement of the domain in time is inherently accounted for in the formulation. The space/time formulation is thus suitable for the Eulerian description of problems on moving domains such as a fluid flow with a moving boundary or a free surface; see, e.g., Ref. [66]. This renders the space/time description an attractive alternative to the classical ALE description. Therefore, we shall use the space/time formulation in the sequel for the specification of the governing equations of the fluid, the structure and the interface conditions, which allows us to treat the subsystems and their interaction in a unified framework.

2.3 Governing equations for fluid flows

2.3.1 Classical formulation

The motion of fluids (liquids and gases) is governed by conservation laws which state that mass, momentum and energy of a fluid are conserved during its motion. To describe a fluid flow, the state of the flow is expressed by so-called *state variables*. These state variables are characteristic fluid properties such as density, velocity, pressure, etc. The mathematical description of the conservation laws with the conserved quantities expressed in terms of the state variables yields a system of partial differential equations. A derivation of these equations can be found in many textbooks, for instance, in the one by Batchelor [3].

To introduce the governing equations for fluid flow on a domain which deforms in time, we adopt the space/time formulation of the fluid equations in Eulerian form. Denoting by $x \in \mathbb{R}^d$ (d = 1, 2, 3) and $t \in \mathbb{R}_+$ spatial and temporal coordinates, respectively, we consider a fluid in an open bounded space/time domain $\Omega_{\alpha} \subset \mathbb{R}^d \times \mathbb{R}_+$. Its boundary consists of the moving boundary, Γ_{α} , and the fixed boundary, $\partial \Omega_{\alpha} \setminus \Gamma_{\alpha}$, where α connotes a representation function of the moving boundary, which will be elaborated in the sequel.

Appropriate state variables for the description of a viscous, compressible fluid are the density $\rho(x,t)$, the velocity v(x,t), the pressure p(x,t), the absolute temperature $\theta(x,t)$ and the specific internal energy of the fluid e(x,t). The total energy is defined by $E := \rho(e + v^2/2)$. Conservation of mass, momentum and energy are then expressed by, respectively,

$$\frac{\partial}{\partial t}\rho + \nabla \cdot (\rho v) = 0, \qquad (x,t) \in \Omega_{\alpha}, \qquad (2.8a)$$

$$\frac{\partial}{\partial t}\rho v + \nabla \cdot (\rho v v + pI - \tau) - \rho g = 0, \qquad (x, t) \in \Omega_{\alpha}, \qquad (2.8b)$$

$$\frac{\partial}{\partial t}E + \nabla \cdot \left((E+p)v - v \cdot \tau - k\nabla\theta\right) - \rho v \cdot g = 0, \qquad (x,t) \in \Omega_{\alpha}, \qquad (2.8c)$$

with all quantities and operations evaluated in the Eulerian space/time configuration. Specifically, ∇ and ∇ · denote the gradient operator and the divergence operator, respectively; k is the thermal conductivity of the fluid, τ the viscous stress tensor, g the gravitational acceleration and I the identity. With k = 0 and $\tau = 0$, the equations (2.8) are called the *Euler equations*.

Closure of the system of equations (2.8) requires complementary relations, viz., two additional *thermodynamic relations* and a *constitutive relation*. The thermodynamic relations give a relation between the state variables; see, e.g., Ref. [2]. The constitutive relation relates the viscous stress tensor to the state variables. Whereas the conservation statements of mass, momentum and energy are independent of the considered fluid, the thermodynamic relations and the constitutive relation are fluid specific. For instance, in the case of a (calorically) perfect gas, such thermodynamic relations are given by

$$p = \rho R \theta, \qquad e = e(\theta),$$
 (2.8d)

where the first relation is also referred to as the equation of state, and R denotes the specific gas constant. The constitutive relation for the viscous stress tensor in the case of a Newtonian fluid is

$$\tau := \mu \left([\nabla v] + [\nabla v]^T - \frac{2}{3} [\nabla \cdot v] I \right), \tag{2.8e}$$

with μ the dynamic viscosity of the fluid; see, e.g., Ref. [3]. The momentum equations (2.8b) with the viscous stress tensor defined according to relation (2.8e) are referred to as the *Navier–Stokes equations*.

To complete the description of the initial-boundary-value problem for the fluid, Eqs. (2.8) must be supplemented with suitable initial and boundary conditions. In particular, at the moving boundary Γ_{α} , the boundary velocity reappears in the boundary condition. For viscous flows, the identification of the fluid velocity at the moving boundary with the boundary velocity forms an appropriate boundary condition, and corresponds to the so-called 'no-slip' condition. For inviscid flows, a specification of the boundary velocity in the normal direction suffices. In either case, this condition translates into the impermeability of the boundary Γ_{α} ; cf. Section 2.5. Moreover, at material boundaries, the specification of the temperature or the heat flux is required. However, we shall not consider thermal effects in this work.

2.3.2 Variational space/time formulation

To phrase the classical problem statement into a space/time variational form, we multiply the equations by appropriate test functions and integrate the resulting expression over the space/time domain; see, e.g., Refs. [37, 62, 65].

To formulate the variational statement for the fluid problem, we represent the moving boundary Γ_{α} by a representation function α which we accommodate in a space of admissible moving-boundary representations \mathbb{A} . Thus, to each $\alpha \in \mathbb{A}$ corresponds a specific $\Gamma(\alpha) := \Gamma_{\alpha}$ and, accordingly, an $\Omega(\alpha) := \Omega_{\alpha}$. Note that the specifics of the mapping $\alpha \mapsto \Gamma_{\alpha}$ depend on whether the fluid is viscous or inviscid. In the viscous case, the mapping bears the form $\alpha : \Xi \times \mathbb{R}_+ \mapsto \mathbb{R}^d (\Xi \subseteq \mathbb{R}^{d-1})$ and specifies the location of each point on the moving boundary explicitly by

$$\Gamma_{\alpha} = \{ (x,t) : x = \alpha(\chi,t), t \in \mathbb{R}_+, \chi \in \Xi \subseteq \mathbb{R}^{d-1} \}.$$

$$(2.9)$$

The corresponding velocity of the moving boundary, $\partial \alpha / \partial t$, enters the initialboundary-value problem for the fluid in the form of a boundary condition. Whereas viscous flows require a specification of the boundary velocity in all directional components, for inviscid flows only the normal component of the boundary velocity is required. Therefore, for inviscid flows, a weaker description than (2.9) suffices. In the inviscid case, we therefore redefine $\alpha : \Xi \times \mathbb{R}_+ \mapsto \mathbb{R}$. This mapping can describe, for instance, the displacement of a reference surface Γ_0 in the direction of its outward unit normal vector n(x,t) according to $\Gamma_\alpha := \{(x,t) + \alpha(x,t)n(x,t) : (x,t) \in \Gamma_0\}$. We remark that alternative descriptions of Γ_α in compliance with $\alpha : \Xi \times \mathbb{R}_+ \mapsto \mathbb{R}$ are possible.

To each admissible boundary representation Γ_{α} corresponds an initial-boundary-value problem for the fluid. We phrase this initial-boundary-value problem into the concise abstract variational statement: Find

$$\mathbf{u} \in \mathbb{U}_{\alpha}$$
: $\mathsf{F}_{\alpha}(\mathbf{v}, \mathbf{u}) = \mathsf{f}_{\alpha}(\mathbf{v}) \quad \forall \mathbf{v} \in \mathbb{V}_{\alpha},$ (2.10)

with the semi-linear functional $\mathsf{F}_{\alpha} : \mathbb{V}_{\alpha} \times \mathbb{U}_{\alpha} \mapsto \mathbb{R}$ associated with the space/time differential operator and the linear functional $\mathsf{f}_{\alpha} : \mathbb{V}_{\alpha} \mapsto \mathbb{R}$ associated with the prescribed data.

The initial and boundary conditions can be enforced strongly by incorporating them in the trial function space \mathbb{U}_{α} and restricting the test space \mathbb{V}_{α} accordingly. Alternatively, the auxiliary conditions can be imposed weakly and incorporated into the functionals by replacing the boundary integrals by their definitions conforming to the initial and boundary conditions. In either case, the function spaces depend on the function α through the domain on which they are defined, i.e., $\mathbb{H}_{\alpha} := \mathbb{H}(\Omega_{\alpha})$ ($\mathbb{H} = \mathbb{U}, \mathbb{V}$). Note that in fluid-structure interaction the function α actually constitutes an unknown: On account of the free-boundary character of the interface, α forms part of the solution. The dependence of the function spaces on α prevents us from casting the fluid problem and, thus, also the fluid-structure-interaction problem into a canonical variational form. To recover the canonical form, Ref. [71] proposes methodologies to remove the dependence on α from the function spaces. If the auxiliary conditions are enforced weakly, this dependence can be disposed of by embedding the union of all admissible domains in a so-called security set in combination with an extension mapping; see also Ref. [57]. If the auxiliary conditions are enforced strongly, then the dependence of the function spaces \mathbb{U}_{α} , \mathbb{V}_{α} on Ω_{α} can be removed by means of an α -dependent homeomorphic transformation, i.e., a one-to-one mapping between a moving and a fixed reference domain; see Ref. [71] for details. With the dependence on α removed from the function spaces, the fluid variational problem can be reformulated as

$$\mathbf{u} \in \mathbb{U}$$
: $\mathsf{F}(\mathbf{v}, \mathbf{u}, \alpha) = \mathsf{f}(\mathbf{v}) \quad \forall \mathbf{v} \in \mathbb{V}.$ (2.11)

Note that both the Euler and the Navier-Stokes equations on a domain with a moving boundary can be condensed into the variational form (2.11).

2.4 Governing equations for structures

2.4.1 Classical formulation

The motion of a structure is governed by the balance of momentum, and it is customarily described in a Lagrangian reference frame; cf. Section 2.2.2. The governing equations comprise the momentum balance laws, constitutive equations and kinematic relations. To introduce a model of the structure that is sufficiently general for most fluid-structure-interaction problems, we require a model that admits large structural deformations and, thus, a geometrically nonlinear description of the kinematics. However, we shall assume that the material behaviour is linear-elastic and, therefore, limit our considerations to stretches close to one. A structural model that complies with these requirements and assumptions is given by geometrically nonlinear elastodynamics in combination with the linear Saint-Venant Kirchhoff material law; see, e.g., Ref. [50].

To introduce the governing equations for the structure, let $x, y \in \mathbb{R}^d$ (d = 1, 2, 3) and $t \in \mathbb{R}_+$ denote spatial, material and temporal coordinates, respectively. We consider the structure on an open bounded material space/time domain, Ω_y . The displacement field z(y,t) is then expressed as the difference in coordinates of a material point between the spatial configuration and the material configuration

$$z(y,t) := x(y,t) - y = \varphi(y,t) - y, \qquad (2.12)$$

where the mapping φ relates the material configuration to the spatial configuration; see Eq. (2.2a). The displacement field z complies with the *balance of linear* momentum, which states the equilibrium of surface, volume and inertia forces

$$\rho \frac{\mathrm{d}^2 z}{\mathrm{d}t^2} = \nabla \cdot P + \rho b, \qquad (y,t) \in \Omega_y, \tag{2.13}$$

where ρ is the density, ρb designates the volume forces, and P is the first Piola-Kirchhoff stress tensor. In (2.13), all quantities and operations are evaluated in the material configuration Ω_y . For a convenient specification of the constitutive equation, we introduce the second Piola-Kirchhoff stress tensor S that is related to the first Piola-Kirchhoff stress tensor by

$$S = F^{-1}P. (2.14)$$

In (2.14), F denotes the material deformation gradient which maps a line element dy in the material domain onto a line element dx in the spatial domain:

$$dx = Fdy \quad \Leftrightarrow \quad F := \nabla x = I + \nabla z \,, \tag{2.15}$$

where the last equality is obtained upon invoking Eq. (2.12).

Assuming a homogeneous isotropic linear-elastic medium of *Saint-Venant Kirchhoff* type, the constitutive equation

$$S = \lambda(\mathrm{tr}E)I + 2\mu E \tag{2.16}$$

with tr denoting the trace operator, relates the second Piola-Kirchhoff stress tensor to the energetically conjugate Green-Lagrange strain tensor E by means of the Lamé constants λ and μ . The Green-Lagrange strain tensor is defined as

$$E := \frac{1}{2} (F^T F - I) \,. \tag{2.17}$$

Noting that the deformation gradient can alternatively be expressed as F = RU, where R and U denote the rotation and the stretch tensor, respectively, we can rewrite Eq. (2.17) as

$$E := \frac{1}{2} (U^T R^T R U - I) = \frac{1}{2} (U^2 - I), \qquad (2.18)$$

where the second identity follows from the fact that R is orthogonal and U is symmetric. For the assumption of a linear-elastic medium to be valid, the Green-Lagrange strain E has to be sufficiently small. In view of Eq. (2.18), this pertains to the requirement that the stretch tensor U has to be sufficiently close to the identity. However, arbitrary large rotations and, moreover, translations of the structure are permitted. A system of partial differential equations for the displacement field zcan be obtained by combining Eqs. (2.12)–(2.17).

To complete the description of the initial-boundary-value problem for the structure, the system must be supplemented with appropriate initial and boundary conditions. For the specification of these conditions, we subdivide the boundary of the space/time domain into complementary subsets $\partial \Omega_y = \Upsilon_0 \cup \Upsilon_T \cup \Gamma_d \cup \Gamma_n$, where Υ_0 and Υ_T denote the temporal boundary at the bottom and the top of the space/time domain, respectively, and Γ_d and Γ_n denote Dirichlet and Neumann subsets of the spatial boundary. The initial conditions prescribe the displacement and the velocity on Υ_0 , i.e., at time is equal zero

$$z(y,0) = z^{0}(y), \quad \frac{\mathrm{d}z}{\mathrm{d}t}(y,0) = \dot{z}^{0}(y), \quad \text{on} \quad \Upsilon_{0}$$
 (2.19)

with $z^0(y)$ and $\dot{z}^0(y)$ given. The boundary conditions prescribe surface displacements on Dirichlet boundaries

$$z(y,t) = \hat{z}(y,t), \quad \text{on} \quad \Gamma_d \tag{2.20}$$

with $\hat{z}(y,t)$ given, and surface tractions on Neumann boundaries

$$P \cdot N = \hat{t}(y, t), \quad \text{on} \quad \Gamma_n$$

$$(2.21)$$

with $\hat{t}(y,t)$ given and N the outward unit normal vector of the structure. The surface tractions $\hat{t}(y,t)$ can either be given explicitly, or they can derive from the inner product of a given stress tensor \hat{P} on Γ_n with the outward unit normal vector of the structure, i.e., $\hat{t}(y,t) = \hat{P} \cdot N$.

In advance of Section 2.5, we shall identify a subset of the Neumann boundary $\Theta \subseteq \Gamma_n$ as the interface to the fluid. The stress tensor on Θ is supplied by the fluid via the dynamic interface condition; cf. Section 2.5. We note that the dynamic condition provides the *Cauchy stress tensor* which is defined in the spatial domain, in conformity with the Eulerian description of the fluid. Since the structure is generally described in the material domain, the Cauchy stress tensor obtained from the dynamic condition needs to be transformed into a corresponding stress tensor defined in the material domain. The transformation between the Cauchy stress tensor π and the first Piola-Kirchhoff stress tensor P is given by

$$P = (\det F)\pi F^{-T} \qquad \Leftrightarrow \qquad \pi = \frac{1}{\det F}PF^{T}$$
 (2.22)

with $\det F$ the determinant of the material deformation gradient F.

Note that in many cases the description of a structure admits simplified models, for instance, in the case of beams and shells; cf., e.g., Refs. [5, Ch.3] and [4, Ch.5]. In this thesis, we shall restrict ourselves to a geometrically linear description of the kinematics and consider the Euler-Bernoulli beam equation, $Mz_{,tt} + Dz_{,yyyy} = q$, with M, D and q denoting the mass, the bending stiffness and the distributed loading of the beam, respectively.

2.4.2 Variational space/time formulation

We consider the initial-boundary-value problem for the structure subject to prescribed initial and boundary conditions. The latter comprise prescribed displacements and a given stress tensor σ on subsets of the structure boundary. The structure problem can be written concisely in variational form: Find

$$z \in \mathbb{Z}: \qquad \mathsf{S}_{\sigma}(\lambda, z) = \mathsf{s}(\lambda) \qquad \forall \lambda \in \mathbb{L}$$
(2.23)

with $\mathbb{Z} := \mathbb{Z}(\Omega_y)$ and $\mathbb{L} := \mathbb{L}(\Omega_y)$. Dirichlet boundary conditions, i.e., prescribed displacements, are incorporated in \mathbb{Z} , and \mathbb{L} is constrained accordingly, i.e., the test functions vanish on Dirichlet boundaries.

The stress tensor on the boundary Θ derives from the fluid solution and, thus, it is unknown a priori; cf. Section 2.5. To make this dependence apparent in the equation, we introduce the notation $S(\lambda, z, \pi) := S_{\sigma}(\lambda, z)$, where π represents the stress tensor on Θ . Note that we have incorporated π into the functional S instead of the functional s, which contains the prescribed data, because the transformation from the Cauchy stress tensor π to the first Piola-Kirchhoff stress tensor P depends on z; see Eq. (2.22). The dependence on the stress tensor at $\Gamma_n \setminus \Theta$ is tacitly incorporated into the functional S. Upon these modifications, Eq. (2.23) can be rewritten in the following form

$$z \in \mathbb{Z}$$
: $S(\lambda, z, \pi) = s(\lambda) \quad \forall \lambda \in \mathbb{L}$. (2.24)

We assume that the variational problem (2.24) admits a unique solution for all π in a space of admissible stress-tensor functions \mathbb{P} .

Note that the equations governing geometrically nonlinear elastodynamics as well as the Euler-Bernoulli beam equation can be condensed into the variational form (2.24). In the latter case, the functionals are specified as $S = \langle \lambda, Mz_{,tt} + Dz_{,yyyy} - q \rangle_{\mathbb{Z}}$ and s = 0.

2.5 Interface conditions

The fluid and the structure problem are connected at their common interface by kinematic and dynamic conditions, which we specify below.

2.5.1 Kinematic interface conditions

The kinematic interface conditions specify that the fluid boundary and the structure boundary coincide at the interface and, moreover, that the fluid velocity at the interface equals the velocity of the interface. The latter condition actually constitutes a boundary condition for the initial-boundary-value problem of the fluid and, as such, it is incorporated in the variational statement (2.11), see Section 2.3. The former condition interrelates the displacement of the structure at the interface $z|_{\Theta}$ and the representation of the fluid boundary α through the identity $\Gamma_{\alpha} = z|_{\Theta}(\Theta)$. We phrase this identity in the variational statement: Find

$$\alpha \in \mathbb{A}: \qquad \mathsf{K}(\eta, \alpha, z) = \mathsf{k}(\eta) \qquad \forall \eta \in \mathbb{H}.$$
(2.25)

We stipulate that the functional K can be separated into $\mathsf{K}(\eta, \alpha, z) = \mathsf{K}_0(\eta, \alpha) + \mathsf{K}_1(\eta, z)$, which holds, if K is linear. For instance, if α is specified as $\alpha = \alpha^0 + z|_{\Theta}$ with $\alpha^0, z|_{\Theta}$ residing in \mathbb{A} , then the functionals take on the following form : $\mathsf{K}(\eta, \alpha, z) = \langle \eta, \alpha - z|_{\Theta} \rangle_{\mathbb{A}} = \langle \eta, \alpha \rangle_{\mathbb{A}} - \langle \eta, z|_{\Theta} \rangle_{\mathbb{A}}$ and $\mathsf{k}(\eta) = \langle \eta, \alpha^0 \rangle_{\mathbb{A}}$. Moreover, \mathbb{A} and \mathbb{H} can then be identified. In this work, we shall generally assume K and k to be of this form.

2.5.2 Dynamic interface condition

The dynamic interface condition identifies the surface tractions exerted on the interface by the fluid and the structure. Since these surface tractions derive from the inner product of the stress tensor with the outward unit normal vector, and this normal vector depends on the structure solution, we formulate the dynamic interface condition in terms of the stress tensor rather than in terms of the surface traction.

To specify the dynamic interface condition, let $T \in \mathbb{T}(\Omega)$ denote the fluid stress tensor which constitutes an operator according to $T : \mathbb{U}(\Omega) \mapsto \mathbb{T}(\Omega)$, induced by a constitutive relation. For instance, for an inviscid compressible fluid, the stress tensor reduces to the pressure according to $T : \mathbf{u} \mapsto p(\mathbf{u})I$ with p a given equation of state and I the identity in $\mathbb{R}^{d \times d}$. Another relevant example is the stress tensor for an incompressible Newtonian fluid $T : (v, p) \mapsto pI - \operatorname{Re}^{-1}([\nabla v] + [\nabla v]^T)$, where $(v, p) =: \mathbf{u}$ connotes a velocity/pressure pair and Re represents the Reynolds number. The dynamic condition identifies the stress tensor π at the structure boundary $\alpha^0 + z|_{\Theta}$ and the stress tensor in the fluid $T(\mathbf{u})$ at the boundary Γ_{α} through $\pi = T(\mathbf{u})|_{\Gamma_{\alpha}}$. We phrase this identity in the variational statement

$$\pi \in \mathbb{P}: \qquad \mathsf{D}(\phi, \mathbf{u}, \alpha, \pi) = 0 \qquad \forall \phi \in \mathbb{F}, \tag{2.26}$$

e.g., if $T(\mathbf{u})|_{\Gamma_{\alpha}} \in \mathbb{P}$, then $\mathsf{D}(\phi, \mathbf{u}, \alpha, \pi) = \langle \phi, \pi - T(\mathbf{u})|_{\Gamma_{\alpha}} \rangle_{\mathbb{P}}$.

Note that, in conformity with the Eulerian description of the fluid, the identification $\pi = T(\mathbf{u})|_{\Gamma_{\alpha}}$ conceives both stress tensors as Cauchy stress tensors which are defined in the spatial domain. Since the structure is generally described in the material domain, the stress tensor π obtained from the dynamic condition needs to be transformed into a corresponding stress tensor defined in the material domain. The transformation is given by Eq. (2.22).

2.6 Aggregated variational problem

With the above definitions, the fluid-structure-interaction problem in space/time can be condensed into the canonical form: Find

$$q \in \mathbb{Q}$$
: $\mathsf{P}(w,q) = \mathsf{p}(w) \quad \forall w \in \mathbb{W}.$ (2.27a)

Herein, q is the quadruple $q := (\mathbf{u}, \alpha, z, \pi)$ in the product space $\mathbb{Q} := \mathbb{U} \times \mathbb{A} \times \mathbb{Z} \times \mathbb{P}$, w is the quadruple $w := (\mathbf{v}, \eta, \lambda, \phi)$ in the product space $\mathbb{W} := \mathbb{V} \times \mathbb{H} \times \mathbb{L} \times \mathbb{F}$, and the aggregated functionals $\mathsf{P} : \mathbb{W} \times \mathbb{Q} \mapsto \mathbb{R}$ and $\mathsf{p} : \mathbb{W} \mapsto \mathbb{R}$ are defined respectively as

$$\begin{split} \mathsf{P}((\mathbf{v},\eta,\lambda,\phi),(\mathbf{u},\alpha,z,\pi)) &:= \mathsf{F}(\mathbf{v},\mathbf{u},\alpha) + \mathsf{K}(\eta,\alpha,z) + \mathsf{S}(\lambda,z,\pi) + \mathsf{D}(\phi,\mathbf{u},\alpha,\pi), \\ & (2.27\mathrm{b}) \\ \mathsf{p}(\mathbf{v},\eta,\lambda,\phi) &:= \mathsf{f}(\mathbf{v}) + \mathsf{k}(\eta) + \mathsf{s}(\lambda). \end{split}$$

Note that any fluid-structure-interaction problem can be cast in the variational form (2.27). Hence, the variational formulation (2.27) is generic. For instances of a variational statement of a fluid-structure-interaction problem conforming to (2.27) we refer to Section 3.2 and also to Ref. [70].

On account of the canonical variational form, sufficient conditions for the existence of a unique solution to the variational problem (2.27) in principle follow from the generalized nonlinear Lax-Milgram theorem; cf., e.g., Ref. [69]. However, to establish the existence of a unique solution in specific instances of (2.27) is generally intractable. Only for a few specific cases and based on restrictive assumptions such results are available; see, for instance, Ref. [23] for the interaction of a fluid with a moving rigid structure.

2.7 Discretization of the variational problem

In this section, we briefly elaborate the discretization of the aggregated variational fluid-structure equations from an abstract viewpoint. This serves as a basis to discuss some relevant implications for the numerical treatment of fluid-structure interactions, such as the computation on moving meshes, the discrete representation of the fluid-structure interface and certain conservation properties of the continuum fluid-structure system.

2.7.1 Space/time finite-element discretization

To translate the continuum variational statement (2.27) into a computable algebraic expression, the infinite-dimensional function spaces \mathbb{Q} and \mathbb{W} are replaced by finite-dimensional spaces $\widetilde{\mathbb{Q}}$ and $\widetilde{\mathbb{W}}$. Typically, the approximation spaces $\widetilde{\mathbb{Q}}$ and $\widetilde{\mathbb{W}}$ form subspaces of the original function spaces \mathbb{Q} and \mathbb{W} . In that case, the discretization is called *conforming*. If the test space $\widetilde{\mathbb{W}}$ and the trial space $\widetilde{\mathbb{Q}}$ are identified, the discretization is commonly referred to as a (Bubnov-)Galerkin method. Upon covering the computational domain with a tessellation of elements and defining the approximation spaces piecewise, i.e., per element, a Galerkin *finite-element* discretization is obtained. In this thesis, we restrict ourselves to Galerkin *space/time finite-element* methods, which yield a discretization in both space and time. In particular, we shall use the *time-discontinuous Galerkin method*, which admits discontinuities at time-slab interfaces; see, e.g., Refs. [37, 62].

A basic requirement for any finite-element discretization is convergence. Necessary conditions for convergence of the discrete solution to the continuum solution are consistency and stability of the discrete problem, and convergence of the discrete approximation space. To elaborate these conditions, let us consider a sequence of nested, increasingly fine approximations $(\{\widetilde{\mathbb{Q}}_i\}, \{\widetilde{\mathbb{W}}_i\})$ and corresponding discrete problems: Find

$$q \in \widetilde{\mathbb{Q}}_i: \quad \hat{\mathsf{P}}(w,q) = \hat{\mathsf{p}}(w) \quad \forall w \in \widetilde{\mathbb{W}}_i$$

$$(2.28)$$

for given functionals $\hat{\mathsf{P}}$ and $\hat{\mathsf{p}}$ that can be different from P and p for non-conforming discretizations. Stability requires that the discrete problems are well-posed, i.e., that small perturbations in the data cause only correspondingly small perturbations in the discrete solution. Thus, stability implies that the discrete problems are solvable. We say that the discretization (2.28) is consistent, if the continuum solution satisfies the discrete variational statement. In addition to stability and consistency, convergence of the discrete solution to the continuum solution requires convergence of the sequence $\{\widetilde{\mathbb{H}}_i\}$ ($\mathbb{H} := \mathbb{Q}, \mathbb{W}$). For conforming discretizations, we require that the sequence $\{\widetilde{\mathbb{H}}_i\}$ is asymptotically dense, i.e., $\widetilde{\mathbb{H}}_i \to \mathbb{H}$ as $i \to \infty$. For non-conforming discretizations such as discontinuous Galerkin methods, we require that the sequence $\{\widetilde{\mathbb{H}}_i\}$ encompasses the continuum function space \mathbb{H} in the asymptotic limit, i.e., $\widetilde{\mathbb{H}}_i \to \overline{\mathbb{H}} \supset \mathbb{H}$ as $i \to \infty$. For further elaboration, and for definitions and basic properties of the finite-element method, we refer to the textbooks [9, 61], for instance.

2.7.2 Implications of the discretization

For free-boundary problems in general, and fluid-structure-interaction problems in particular, space/time finite-element discretizations offer distinct advantages over conventional methods that discretize space and time separately, i.e., by different discretization methods. An important asset of space/time finite-element discretizations is that they can easily handle moving meshes, because the variational statement is expressed over the space/time domain and, thus, a movement of the domain in time is inherently accounted for in the formulation. In contrast, conventional methods typically use an arbitrary Lagrangian-Eulerian formulation for flow computations on moving meshes; see also Section 2.2. Moreover, conventional methods are typically required to comply with the so-called *Discrete Geometric* Conservation Law (DGCL) which imposes conditions on the evaluation of the position and velocity of the moving mesh; see, e.g., Refs. [17, 35] and also the review article [16]. The DGCL derives from the requirement that, for a numerical scheme on a moving mesh, the change in control volume is equal to the area swept by the moving boundaries. An equivalent formulation of this requirement is that the numerical scheme preserves the trivial solution of a uniform flow on a moving mesh. Compliance with the DGCL appears to be relevant for the stability and accuracy of a numerical method on moving meshes, but this connection still lacks a proper general theoretical foundation; see Refs. [14, 16]. We shall not further pursue the implications of the DGCL here, because space/time finite-element methods satisfy the DGCL inherently; see, e.g., Ref. [35].

Since in this thesis we adopt a time-discontinuous Galerkin discretization, let us briefly mention its advantages and complications for the numerical treatment of fluid-structure interaction. Time-discontinuous Galerkin methods facilitate an adaptive discretization in that the discrete approximation on either side of the time slab interface can in principle be arbitrary and, thus, completely unrelated. However, a discontinuous discretization of the structure induces a discontinuous displacement of the fluid boundary. This requires special care in the treatment of the fluid initial conditions associated with each time slab. For instance, an extrapolation of the fluid solution can be necessary to extend the solution to a subset of the time boundary that was not occupied by the fluid in the previous time slab. On the other hand, a time-discontinuous discretization of the fluid does not pose any difficulties for the structure discretization. Since there are no continuity requirements on the dynamic forcing, a discontinuous fluid pressure constitutes an apt boundary condition for the structure subproblem.

Another relevant aspect of the discretization is its ability to maintain certain conservation properties of the continuum fluid-structure system. In advance of Chapter 3, we note that the continuum problem conserves mass, momentum and energy at the fluid-structure interface. Under discretization, these conservation properties can be lost, i.e., the discrete problem does not necessarily inherit the conservation properties of the continuum problem. Therefore, the conservation properties of the discrete problem need to be re-established. It appears that the commonly used discretizations often lack conservation at the fluid-structure interface. More precisely, these discretizations typically satisfy conservation only in an asymptotic sense, i.e., for vanishing mesh width. This is, however, a basic consistency requirement. We defer a further elaboration of these aspects to Chapter 3.

Clearly, the discretization has implications also for the resulting algebraic system and, thus, for the solution methods. A particularity of fluid-structureinteraction discretizations is the fact that the dimensions of the approximation spaces associated with the interface variables are generally negligible in comparison to the dimensions of the approximation spaces of the fluid and the structure, because the former refer to boundary functions. This observation plays an important role in the analysis of the discretized fluid-structure system, and in the development of improved solution techniques; see Chapters 5 and 6.

2.8 Concluding remarks

In this chapter, we presented the equations governing fluid-structure interaction which comprise the equations for the fluid and the structure, augmented by kinematic and dynamic interface conditions. We elaborated the different reference frames that can be adopted for the kinematic description of a continuum. Since fluid and structure are customarily described in an Eulerian and a Lagrangian reference frame, respectively, this difference in reference frames necessitates a transformation of interface quantities from one reference frame into the other. Moreover, we discussed the difficulties pertaining to an Eulerian fluid description on a moving domain, and we compared two methodologies to conveniently incorporate a movement of the domain in the description.

To identify the generic features of fluid-structure-interaction problems, we phrased the governing equations into a generic space/time variational form, i.e., a form that does not make any stipulations on the specifics of the underlying models. As generic features we identify the interconnection of a fluid and a structure subsystem at an interface by kinematic and dynamic conditions and, moreover, the free-boundary character of the interface. The free-boundary character is manifest through the interdependence of the fluid and structure solutions and their domains of definition. The generic variational statement allows us to formulate solution methods independent of the specifics of the underlying fluid-structureinteraction problem. Hence, it enables us to establish solution methods that are, in principle, applicable to any fluid-structure-interaction problem.

We considered the discretization of the aggregated variational statement by the space/time finite-element method. Such space/time finite-element discretizations have distinct advantages over conventional methods that discretize space and time separately by different discretization methods. Whereas conventional methods commonly use an arbitrary Lagrangian-Eulerian formulation to enable the description of fluid flow on a moving domain, the space/time formulation inherently accounts for a movement of the domain in time. Consequently, space/time finite-element discretizations automatically satisfy the Discrete Geometric Conservation Law in contrast to conventional methods with a separate space and time discretization.

In this thesis, we shall restrict ourselves to prototypical model problems that conform to the presented generic variational statement. This allows us to analyse the characteristic features of fluid-structure-interaction problems with a minimum of complexity.
Chapter 3 Conservation under discretization

3.1 Introduction

The interaction of a fluid with a structure is driven by the exchange of momentum and energy at their common interface. The correct representation of this momentum and energy exchange in the numerical model is therefore crucial in the numerical solution of fluid-structure-interaction problems. However, maintaining the conservation properties at the interface in the numerical model is generally non-trivial.

In this chapter, we investigate the conservation properties in the discrete numerical model, following the elaborations in Refs. [44, 70]. To this end, we establish the conservation properties at the fluid-structure interface for the continuum system and for its discretization by the finite-element method. The conservation properties of the continuum system can be lost under discretization. Indeed, it is shown that energy conservation in the discrete model is only trivially maintained under restrictive compatibility conditions on the approximation spaces in the fluid and the structure, i.e., matching meshes and identical orders of interpolation at the interface. Clearly, this is prohibitive for practical use, as it impedes tailoring fluid and structure discretizations to their inherently different length and time scales; cf. Ref. [19]. To bypass these restrictions, a new discretization from Ref. [70] is considered, which enables conservation also for incompatible discretizations.

The ability to maintain the conservation properties at the interface is moreover subject to the choice of the discretization methods for the fluid and structure, and their pairing at the interface; see also Ref. [44]. To investigate the implications of specific discretizations for the conservation properties, we consider the trapezoidal method for the structure coupled to a finite-element discretization of the fluid equations. We show that trapezoidal time integration for the structure does in general not allow for conservative fluid-structure coupling, although the trapezoidal method for the structure separately is conservative.

To assess the relevance of maintaining the conservation properties under discretization, we conduct numerical experiments on a prototypical fluid-structureinteraction problem, viz., the piston problem from Ref. [56]. This model problem facilitates the separation of distinct sources of error. In particular, it admits an analytic solution of the structure equation and, thus, exact time integration of the structure. Fluid discretization errors can be reduced using fluid subcycling. This combination allows us to study the effect of maintaining conservation at the interface separately from fluid and structure discretization errors. To demonstrate the effectiveness of the new discretization considered in this chapter, we compare conservative and non-conservative discretizations. Our numerical results demonstrate that discretizations which ensure the conservation properties are superior to nonconservative discretizations. In particular, the conservative discretization yields a solution which is by several orders of magnitude more accurate than the solutions obtained from the non-conservative discretizations. Moreover, we show that a violation of the conservation properties can result in instability of the numerical solution; see also Refs. [44, 53, 55].

To illustrate the implications of specific pairings of fluid and structure discretizations, we compare the numerical solution obtained with, respectively, the trapezoidal and the time-discontinuous Galerkin time integration for the structure coupled to a time-discontinuous Galerkin fluid discretization. Our results indicate that the non-conservative coupling of a trapezoidal method for the structure with a time-discontinuous Galerkin discretization for the fluid induces an error which can dominate other discretization errors. In contrast, the conservative coupling of a time-discontinuous Galerkin discretization for both the fluid and the structure induces an error which is typically much smaller than other discretization errors.

The contents of this chapter are organized as follows: Section 3.2 presents a statement of the piston problem, both in strong form and in variational form. Section 3.3 establishes the conservation properties at the fluid-structure interface of the continuum system and, moreover, investigates the behaviour of the linearized system. Section 3.4 assesses the conservation properties at the interface for compatible and incompatible discretizations. Moreover, this section analyses the implications of specific pairings of fluid and structure discretization methods for the conservation properties. In Section 3.5, we present numerical experiments and results. Section 3.6 contains concluding remarks.

3.2 Problem statement

In Section 3.2.1, we present a prototypical fluid-structure interaction problem, viz., the one-dimensional piston problem from Ref. [56]. In Section 3.2.2, we rephrase the piston problem into a concise space/time variational statement, in conformity with Chapter 2.

3.2.1 The piston problem

To formulate the piston problem, let x and t denote spatial and temporal coordinates, respectively, and $\alpha(t)$ the position of the fluid-structure interface. We consider an open bounded space/time domain $\Omega_{\alpha} := \{(x,t) : 0 < t < T; 0 < x < \alpha(t)\}$. Its boundary consists of the *interface* between the fluid and the structure $\Gamma_{\alpha} := \{(x,t) : x = \alpha(t); 0 < t < T\}$ and the fixed boundary $\partial \Omega_{\alpha} \setminus \Gamma_{\alpha}$; see the illustration in Fig. 3.1 on the next page. The mathematical formulation of the piston system comprises the Euler equations in Ω_{α} in connection with a harmonic oscillator at the interface Γ_{α} . We consider the Euler equations in conservative form:

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{u})}{\partial x} = 0, \qquad (x,t) \in \Omega_{\alpha},$$
(3.1a)

with

$$\mathbf{u} := \begin{pmatrix} \rho \\ \rho v \\ E \end{pmatrix}, \quad \mathbf{f}(\mathbf{u}) := \begin{pmatrix} u_2 \\ u_2^2/u_1 + p(\mathbf{u}) \\ (p(\mathbf{u}) + u_3)u_2/u_1 \end{pmatrix}, \quad p(\mathbf{u}) := (\gamma - 1)\left(u_3 - \frac{u_2^2}{2u_1}\right),$$
(3.1b)

and γ a constant, typically 1.4. In Eq. (3.1b), ρ , v, E and p denote the density, velocity, total energy and pressure of the fluid, respectively.

Eq. (3.1) must be supplemented with appropriate initial and boundary conditions. On the fixed space/time boundary, these are prescribed by

$$\mathbf{u}(x,0) = \mathbf{u}^0(x),$$
 $0 < x < \alpha(0),$ (3.2a)

$$u_2(0,t) = 0,$$
 $0 < t < T,$ (3.2b)

with $\mathbf{u}^{0}(x)$ the given initial conditions. Condition (3.2b) translates into the impermeability condition v(0,t) = 0. The interface conditions, i.e., the conditions on Γ_{α} , are specified below.

The governing equation for the harmonic oscillator is:

$$M\ddot{z}(t) + Kz(t) = \theta(t) - \beta, \quad 0 < t < T,$$

$$(3.3)$$

subject to the initial conditions

$$z(0) = z^0,$$
 (3.4a)

$$\dot{z}(0) = \dot{z}^0$$
 (3.4b)

for certain given constants $M, K \in \mathbb{R}_+$ and $z^0, \dot{z}^0 \in \mathbb{R}$. In Eq. (3.3), z(t) designates the piston displacement from its equilibrium position and a superimposed dot denotes differentiation with respect to time. The right-hand member of Eq. (3.3) is the forcing term which is composed of the traction $\theta(t)$ exerted by the fluid on the structure through the interface, and the constant external pressure β . The forcing term acts on a surface of unit area. The Euler equations and the harmonic oscillator are connected by interface conditions. The interface conditions for the fluid-structure system can be separated into kinematic conditions

$$(\rho v)|_{\Gamma_{\alpha}} = \rho|_{\Gamma_{\alpha}} \dot{\alpha}(t), \qquad \qquad 0 < t < T, \qquad (3.5a)$$

$$\alpha(t) = \alpha^0 + z(t),$$
 $0 < t < T,$ (3.5b)

with $\alpha^0 \in \mathbb{R}_+$ a given constant, and the dynamic condition

$$p(\mathbf{u}|_{\Gamma_{\alpha}}) = \theta(t), \qquad \qquad 0 < t < T. \qquad (3.5c)$$

The conditions (3.5a) and (3.5b) express the impermeability of the interface and identify the interface position and the piston position, respectively. Condition (3.5c) implies equilibrium of the forces exerted on the interface by the fluid and the structure. Note that the conditions (3.5) are imposed on the moving boundary Γ_{α} .



Figure 3.1: The piston problem in space/time (interface region expanded for clarity).

It is noteworthy that the piston problem exhibits the distinguishing features of a free-boundary problem: Firstly, the state variables **u** and their domain of definition Ω_{α} are interdependent on account of the kinematic condition (3.5b). Secondly, the number of interface conditions exceeds by one the number of auxiliary conditions required by the initial-boundary-value problems for the fluid and the structure separately. To elucidate this statement, we note that Eqs. (3.1)– (3.2) subject to (3.5a) with α given, and Eqs. (3.3)–(3.4) subject to (3.5c) with p given are valid problems independently. Therefore, the piston problem can also be classified as a free-boundary problem.

3.2.2 Variational formulation

We rephrase the piston problem into a concise variational statement in space/time by taking L^2 inner products of the equations in Section 3.2.1 with appropriate test functions. To this end, we formally introduce the trial function spaces $\mathbb{U}(:=$ $\mathbb{U}_1 \times \mathbb{U}_2 \times \mathbb{U}_3$), \mathbb{A} , \mathbb{Z} and \mathbb{P} to accommodate \mathbf{u} , α , z and θ , respectively, and the corresponding spaces $\mathbb{V}(:=\mathbb{V}_1\times\mathbb{V}_2\times\mathbb{V}_3)$, \mathbb{H} , \mathbb{L} and \mathbb{F} to accommodate the test functions $\mathbf{v}, \eta, \lambda$ and ϕ , respectively. The dependence of the function spaces U and \mathbb{V} on the boundary-representation function α can be removed by means of an α -dependent homeomorphic transformation, i.e., a one-to-one mapping between a moving and a fixed reference domain, or, alternatively, by means of a so-called security set in combination with an extension mapping; cf. Section 2.3 and Ref. [71] for details. Moreover, we define the functionals

$$\begin{array}{ll} \mathsf{F}: \mathbb{V} \times \mathbb{U} \times \mathbb{A} \mapsto \mathbb{R}, & \mathsf{f}: \mathbb{V} \mapsto \mathbb{R}, & \mathsf{K}: \mathbb{H} \times \mathbb{A} \times \mathbb{Z} \mapsto \mathbb{R}, & \mathsf{k}: \mathbb{H} \mapsto \mathbb{R}, & (3.6) \\ \mathsf{S}: \mathbb{L} \times \mathbb{Z} \times \mathbb{P} \mapsto \mathbb{R}, & \mathsf{s}: \mathbb{L} \mapsto \mathbb{R}, & \mathsf{D}: \mathbb{F} \times \mathbb{P} \times \mathbb{U} \times \mathbb{A} \mapsto \mathbb{R}, & (3.7) \end{array}$$

$$: \mathbb{L} \times \mathbb{Z} \times \mathbb{P} \mapsto \mathbb{R}, \quad \mathsf{s} : \mathbb{L} \mapsto \mathbb{R}, \quad \mathsf{D} : \mathbb{F} \times \mathbb{P} \times \mathbb{U} \times \mathbb{A} \mapsto \mathbb{R}, \tag{3.7}$$

where F, f, K, k, S, s and D are associated with the fluid equations, the kinematic conditions, the structure equations and the dynamic conditions, respectively. For instance, the kinematic and dynamic interface condition, (3.5b) and (3.5c), translate into

$$\mathsf{K}(\eta,\alpha,z) := \int_0^T \eta(\alpha-z)dt \quad \text{and} \quad \mathsf{k}(\eta) := \int_0^T \eta\alpha^0 dt, \quad (3.8)$$

and

$$\mathsf{D}(\phi, \theta, \mathbf{u}, \alpha) := \int_0^T \phi(p(\mathbf{u}|_{\Gamma_\alpha}) - \theta) \mathrm{d}t,$$
(3.9)

respectively. For brevity, we refer to Ref. [70] for a definition of the other functionals.

With the above definitions, the piston problem can be condensed into the concise variational problem: Find $q \in \mathbb{Q}$ such that

$$\mathsf{P}(w,q) = \mathsf{p}(w) \qquad \forall w \in \mathbb{W}.$$
(3.10a)

Herein, q is the quadruple $q := (\mathbf{u}, \alpha, z, \theta)$ in the product space $\mathbb{Q} := \mathbb{U} \times \mathbb{A} \times \mathbb{Z} \times \mathbb{P}$. w is the quadruple $w := (\mathbf{v}, \eta, \lambda, \phi)$ in the product space $\mathbb{W} := \mathbb{V} \times \mathbb{H} \times \mathbb{L} \times \mathbb{F}$, and the aggregated functionals $\mathsf{P}: \mathbb{W} \times \mathbb{Q} \mapsto \mathbb{R}$ and $\mathsf{p}: \mathbb{W} \mapsto \mathbb{R}$ are defined respectively as

$$\begin{split} \mathsf{P}((\mathbf{v},\eta,\lambda,\phi),(\mathbf{u},\alpha,z,\theta)) &:= \mathsf{F}(\mathbf{v},\mathbf{u},\alpha) + \mathsf{K}(\eta,\alpha,z) + \mathsf{S}(\lambda,z,\theta) + \mathsf{D}(\phi,\theta,\mathbf{u},\alpha), \\ & (3.10\mathrm{b})\\ \mathsf{p}(\mathbf{v},\eta,\lambda,\phi) &:= \mathsf{f}(\mathbf{v}) + \mathsf{k}(\eta) + \mathsf{s}(\lambda). \end{split}$$

Note that the variational problem (3.10) conforms to the generic variational statement (2.27).

3.3 System behaviour

This section describes the behaviour of the piston problem. Section 3.3.1 assesses the conservation properties at the fluid-structure interface of the continuum system. Section 3.3.2 briefly reviews an analysis of the linearized system. Moreover, we present a non-dimensionalization which serves to establish the characteristic parameters of the piston system.

3.3.1 Conservation at the fluid-structure interface

The conditions for conservation of mass, momentum and energy at the interface of the continuum fluid-structure system can be derived from the requirement that the corresponding fluxes on either side of the interface are equal. We integrate these fluxes over the interface to express conservation in an integral sense rather than in a pointwise sense. For the considered model problem, an integral over the fluid-structure interface can be transformed into an integral over time. The conservation statements of mass, momentum and energy can then be expressed as the balance of the respective fluxes

$$\int_0^T (\rho v)|_{\Gamma_\alpha} \mathrm{d}t - \int_0^T \rho|_{\Gamma_\alpha} \dot{\alpha} \mathrm{d}t = 0, \qquad (3.11a)$$

$$\int_{0}^{T} \left((\rho v^{2})|_{\Gamma_{\alpha}} + p(\mathbf{u}|_{\Gamma_{\alpha}}) \right) \mathrm{d}t - \int_{0}^{T} (\rho v)|_{\Gamma_{\alpha}} \dot{\alpha} \mathrm{d}t = \int_{0}^{T} \theta \mathrm{d}t, \qquad (3.11b)$$

$$\int_0^T \left(E|_{\Gamma_{\alpha}} + p(\mathbf{u}|_{\Gamma_{\alpha}}) \right) v|_{\Gamma_{\alpha}} \mathrm{d}t - \int_0^T E|_{\Gamma_{\alpha}} \dot{\alpha} \mathrm{d}t = \int_0^T \theta \dot{z} \mathrm{d}t.$$
(3.11c)

In Eqs. (3.11), the left-hand-side constitutes the flux on the fluid side of the interface which is composed of the inviscid flux (first integral) and the flux induced by the movement of the fluid boundary (second integral), and the right-hand-side is the flux on the structure side of the interface. To facilitate the ensuing exposition, let us rewrite Eqs. (3.11) in the form of residuals that vanish identically provided that mass, momentum and energy conservation hold:

$$\mathsf{R}_{1}(\mathbf{u},\alpha) := \int_{0}^{T} \rho|_{\Gamma_{\alpha}}(v|_{\Gamma_{\alpha}} - \dot{\alpha}) \mathrm{d}t, \qquad (3.12a)$$

$$\mathsf{R}_{2}(\mathbf{u},\alpha,\theta) := \int_{0}^{T} (\rho v)|_{\Gamma_{\alpha}}(v|_{\Gamma_{\alpha}} - \dot{\alpha}) \mathrm{d}t + \int_{0}^{T} (p(\mathbf{u}|_{\Gamma_{\alpha}}) - \theta) \mathrm{d}t, \qquad (3.12\mathrm{b})$$

$$\mathsf{R}_{3}(\mathbf{u},\alpha,\theta,z) := \int_{0}^{T} (E|_{\Gamma_{\alpha}} + p(\mathbf{u}|_{\Gamma_{\alpha}})(v|_{\Gamma_{\alpha}} - \dot{\alpha}) \mathrm{d}t + \int_{0}^{T} (p(\mathbf{u}|_{\Gamma_{\alpha}}) - \theta) \dot{\alpha} \mathrm{d}t + \int_{0}^{T} \theta(\dot{\alpha} - \dot{z}) \mathrm{d}t, \quad (3.12c)$$

where $R_j(\cdot)$ denotes the residual pertaining to the respective conservation statement. For convenience, we have added suitable partitions of zero in Eq. (3.12c). A vanishing of the residuals is implied by the variational statement (3.10) on account of the interface conditions. This results in conservation at the fluid-structure interface. Provided that fluid and structure systems separately are conservative, conservation at the interface implies the conservation properties of the aggregated fluid-structure system.

3.3.2 Linearized-system behaviour

An analysis of the linearized piston system has also been presented in Refs. [6, 52, 56, 70]. For completeness, we provide a concise investigation of the linearized-system behaviour, but in view of the available literature we refer the reader to the above references for a further elaboration. Moreover, this section introduces a non-dimensionalization in order to establish the characteristic system parameters and the essential features of the piston problem.

In the sequel, we shall denote first-order perturbations by a prime. We assume the perturbations in the fluid to be isentropic. Under this assumption, the energy equation becomes redundant and can be discarded.

After linearization, and provided with suitable initial conditions (see Ref. [70] for the derivation), the linearized problem consistent with (3.10) admits a fundamental periodic solution:

$$\begin{pmatrix} u_1' \\ u_2' \end{pmatrix} = -\frac{\rho^0 \alpha^0 \omega}{c \sin(\omega \alpha^0 / c)} \begin{pmatrix} \cos(\omega x / c) \cos(\omega t) \\ c \sin(\omega x / c) \sin(\omega t) \end{pmatrix}, \quad z' = \alpha^0 \cos(\omega t),$$
$$\alpha' = z' \quad \text{and} \quad \theta' = c^2 u_1', \quad (3.13)$$

with c denoting the speed of sound and the system frequency ω subject to the condition

$$(1 - K/(M\omega^2))(\omega\alpha^0/c)\tan(\omega\alpha^0/c) = \rho^0\alpha^0/M.$$
(3.14)

Note that Eq. (3.14) has infinitely many roots. The solution to the linearized problem is then composed of a linear combination of the fundamental solutions (3.13).

To establish the characteristic system parameters, we translate the fundamental solutions (3.13) in dimensionless form. To this end, we define characteristic time scales for the fluid, the structure and the fluid-structure system by

$$\tau_{\rm f} := 2\alpha^0/c, \quad \tau_{\rm s} := 2\pi\sqrt{M/K} \quad \text{and} \quad \tau := 2\pi/\omega, \tag{3.15}$$

respectively. Moreover, we introduce the following time-scale ratios:

$$a := \tau_{\rm f}/\tau, \qquad b := \tau_{\rm f}/\tau_{\rm s},\tag{3.16}$$

and the fluid-to-structure mass ratio :

$$\mu := \rho^0 \alpha^0 / M. \tag{3.17}$$

With these definitions, and introducing the following dimensionless variables,

$$x_* := x/\alpha^0, \quad z_* := z/\alpha^0, \quad t_* := t/\tau, \quad v_* := v\tau/\alpha^0, \quad c_* := c\tau/\alpha^0 \text{ and}$$

 $\rho_* := \rho/\rho^0, \quad (3.18)$

Eq. (3.13) can be cast in dimensionless form:

$$\begin{pmatrix} u'_{*1} \\ u'_{*2} \end{pmatrix} = -\frac{2\pi a}{\sin(2\pi a)} \begin{pmatrix} \cos(2\pi ax_*)\cos(2\pi t_*) \\ c_*\sin(2\pi ax_*)\sin(2\pi t_*) \end{pmatrix}, \quad z'_* = \cos(2\pi t_*),$$

$$\alpha'_* = z'_* \quad \text{and} \quad \theta'_* = c^2_* u'_{*1}, \quad (3.19)$$

subject to the condition

$$(1 - b^2/a^2)2\pi a \tan(2\pi a) = \mu.$$
(3.20)

Eq. (3.20) interrelates three dimensionless parameters, viz., the characteristic time-scale ratios a and b and the characteristic mass ratio μ . Thus, only two parameters can be selected independently, the third is then implied by Eq. (3.20). Two parameters are therefore required to characterize the behaviour of the linearized system.

It is noteworthy that the parameter sets $\{(a, b)|\mu > 0\}$ can be conceived as different modes of the fluid-structure system; see Ref. [70] for details. Moreover, we remark that the mass ratio μ determines which subsystem dominates the dynamics of the fluid-structure system. In particular, for $\mu \gg 1$ ($\mu \ll 1$) the dynamics of the system is dominated by the inertia of the fluid (structure).

3.4 Conservation properties under discretization

In this section, we assess the conservation properties of the discrete approximation to the variational form of the continuum fluid-structure system. To this end, in Section 3.4.1, we formally introduce discrete approximations to the original function spaces corresponding to a Galerkin discretization of the variational problem. Because conservation in the original function spaces does not necessarily imply conservation in the approximation spaces, the conservation properties of the original variational form can be lost under discretization. Although our primary interest concerns the conservation properties at the interface, we first briefly consider the conservation properties of the isolated fluid and structure separately in Section 3.4.2. In Sections 3.4.3 and 3.4.4, we assess conservation at the fluidstructure interface for compatible and incompatible approximations, respectively; the presented theory is based on Ref. [70]. Finally, in Section 3.4.5, we investigate the implications of different pairings of fluid and structure discretization methods for the conservation properties at the interface.

3.4.1 Discrete approximations

The discrete approximation to a variational statement generally restricts the original infinite-dimensional function spaces to finite-dimensional subspaces. Here, we consider finite-dimensional approximations in the form of (piecewise) polynomials of order n or, correspondingly, degree n - 1:

$$\mathscr{P}^n := \{ f : \mathbb{R} \mapsto \mathbb{R} | f(x) \in \operatorname{span}(1, x, \dots, x^{n-1}) \}.$$
(3.21)

Discretization in the space and time dimension is based on the tensor product

$$\mathscr{P}^{\mathbf{n}} := \{ f : \mathbb{R}^2 \mapsto \mathbb{R} | f(x_1, x_2) = f(x_1) f(x_2), f(x_1) \in \mathscr{P}^{n_1}, f(x_2) \in \mathscr{P}^{n_2} \}$$
(3.22)

with index pair $\mathbf{n} := (n_1, n_2) \in \mathbb{N}^2$. We denote by $\widetilde{\mathbb{U}} := (\mathscr{P}^{\mathbf{n}_U})^3$ the approximation space associated with the product space $\mathbb{U} := \mathbb{U}_1 \times \mathbb{U}_2 \times \mathbb{U}_3$ and by \mathbf{n}_U its degree. Correspondingly, we introduce $\widetilde{\mathbb{A}}$ and $n_{\mathbb{A}}$, etc. With the above definitions, the Galerkin discretization of the variational problem (3.10) is then: Find $q \in \widetilde{\mathbb{Q}} := \widetilde{\mathbb{U}} \times \widetilde{\mathbb{A}} \times \widetilde{\mathbb{Z}} \times \widetilde{\mathbb{P}}$ such that

$$\mathsf{P}(w,q) = \mathsf{p}(w) \qquad \forall w \in \widetilde{\mathbb{Q}}.$$
(3.23)

We refer to Eq. (3.23) as the set of aggregated (discrete) equations.

We remark that the polynomials that span \mathbb{Q} have support on the domains associated with fluid, structure and interface, respectively. However, (3.23) also holds for approximations based on piecewise polynomials such as finite-element discretizations. In that case, a variational statement of the form (3.23) is associated with each element at the interface.

3.4.2 Conservation in the discretized subsystems

The conservation properties of the Galerkin finite-element discretization of the fluid and structure subsystems are implied by specific choices of the test functions. For the discretization of the fluid equations in conservation form, global conservation of mass, momentum and energy emanates from setting the test function equal to unity on the entire domain. This is straightforward for discontinuous Galerkin methods, but generally not for continuous Galerkin methods in the presence of Dirichlet boundary conditions. Because for Dirichlet boundary conditions the test function has to vanish on the boundary, Dirichlet conditions generally conflict with global conservation. However, it was recently shown for the advection-diffusion and the incompressible Navier-Stokes equations that global conservation can be attained also for the continuous Galerkin method; see Refs. [29, 30]. It is anticipated that for the continuous-in-space time-discontinuous Galerkin discretization of the fluid equations conservation can be established along similar lines.

For the isolated structure discretized with the time-discontinuous Galerkin method, conservation of momentum and energy within each time slab are implied by Eq. (3.23), because the test space associated with the structure equation contains unity and velocity. However, the weak enforcement of initial conditions causes the loss of momentum and energy conservation between two time slabs; see Ref. [70].

3.4.3 Conservation for compatible discretizations

In the sequel, we investigate the conservation properties at the fluid-structure interface, following the exposition in Ref. [70]. We recall from Section 3.3.1 that the conservation properties at the interface of the continuum fluid-structure system are implied by a vanishing of the residuals (3.12). Correspondingly, conservation at the interface under discretization (3.23) requires a vanishing of the residuals. From the conditions for mass, momentum and energy conservation, Eqs. (3.12), we deduce that a basic requirement for conservation is that $v|_{\Gamma_{\alpha}} = \dot{\alpha}$. Under the provision that this identity holds, the residuals assume the form

$$\mathsf{R}_1(\mathbf{u},\alpha) = 0,\tag{3.24a}$$

$$\mathsf{R}_{2}(\mathbf{u},\alpha,\theta) = \int_{0}^{T} (p(\mathbf{u}|_{\Gamma_{\alpha}}) - \theta) \mathrm{d}t, \qquad (3.24b)$$

$$\mathsf{R}_{3}(\mathbf{u},\alpha,\theta,z) = \int_{0}^{T} (p(\mathbf{u}|_{\Gamma_{\alpha}}) - \theta) \dot{\alpha} \mathrm{d}t + \int_{0}^{T} \theta(\dot{\alpha} - \dot{z}) \mathrm{d}t.$$
(3.24c)

Eqs. (3.24) imply that conservation at the interface is immediate if the corresponding fluid and structure variables are *identified pointwise*, i.e., if $\theta = p(\mathbf{u}|_{\Gamma_{\alpha}})$ and $\alpha = \alpha^0 + z$. We remark, however, that $p(\mathbf{u}|_{\Gamma_{\alpha}})$ is in general not polynomial which precludes strong coupling of θ and $p(\mathbf{u}|_{\Gamma_{\alpha}})$ at the interface. The identification of α and $\alpha^0 + z$ has implications for the associated approximation spaces. In particular, it requires that $\widetilde{\mathbb{A}} \supset \widetilde{\mathbb{Z}}$ or, typically, $\widetilde{\mathbb{A}} = \widetilde{\mathbb{Z}}$. We refer to the condition $\widetilde{\mathbb{A}} = \widetilde{\mathbb{Z}}$ as the *compatibility condition*, in conformity with Ref. [70]. Compliance with the compatibility condition requires that in a finite-element discretization the approximation spaces for the fluid and the structure have coincident meshes and identical orders of approximation at the interface. This, however, impedes tailoring fluid and structure discretizations to their specific resolution requirements according to their inherently different length and time scales. Therefore, the compatibility condition is prohibitively restrictive.

Conservation can also be established for specific polynomial approximations, invoking orthogonality arguments. In particular, we consider $\widetilde{\mathbb{P}} := \mathscr{P}^{n_{\mathbb{P}}}, n_{\mathbb{P}} = n_{\mathbb{Z}} - 1$ and $\widetilde{\mathbb{A}} \supseteq \widetilde{\mathbb{Z}}$: Invoking definition (3.9) of the functional associated with the dynamic interface condition, Eq. (3.24) can be rewritten as

$$\mathsf{R}_2(\mathbf{u},\alpha,\theta) = \mathsf{D}(1,\theta,\mathbf{u},\alpha), \qquad (3.25a)$$

$$\mathsf{R}_{3}(\mathbf{u},\alpha,\theta,z) = \mathsf{D}(\dot{\alpha},\theta,\mathbf{u},\alpha) + \int_{0}^{T} \theta(\dot{\alpha}-\dot{z}) \mathrm{d}t.$$
(3.25b)

Clearly, $1 \in \widetilde{\mathbb{P}}$ and, hence, \mathbb{R}_2 vanishes under discretization (3.23). Thus, conservation of momentum is trivially satisfied. That \mathbb{R}_3 also disappears can be shown in the following steps: First, we recall that on account of the definition of K and k in (3.8), Eq. (3.23) identifies α and $\alpha^0 + z$. Obviously, this also holds for $\widetilde{\mathbb{A}} \supset \widetilde{\mathbb{Z}}$, in which case (3.23) reduces the degree of α such that $\alpha \in \widetilde{\mathbb{Z}}$. Hence, $\dot{\alpha} = \dot{z}$ identically, and the second integral in (3.25b) vanishes. Next, from $\alpha \in \widetilde{\mathbb{Z}} := \mathscr{P}^{n_{\mathbb{Z}}}$ it follows that $\dot{\alpha} \in \mathscr{P}^{n_{\mathbb{Z}}-1} = \mathscr{P}^{n_{\mathbb{P}}} =: \widetilde{\mathbb{P}}$, i.e., $\dot{\alpha}$ is in the test space associated with D and, therefore, the first integral in (3.25b) also vanishes. Hence, $\mathbb{R}_3 = 0$ which implies conservation of energy. We point out, however, that energy conservation again relies on the compatibility condition.

In conclusion, conservation of mass and momentum at the interface is in general trivially maintained under discretization. However, conservation of energy at the interface is only trivially satisfied under restrictive compatibility conditions on the approximation spaces $\tilde{\mathbb{A}}$ and $\tilde{\mathbb{Z}}$.

3.4.4 Conservation for incompatible discretizations

Following Ref. [70], the variational statement can be modified such that the discretization maintains the energy conservation property without compatibility conditions on $\widetilde{\mathbb{A}}$ and $\widetilde{\mathbb{Z}}$. The idea is to replace certain orthogonality conditions on $\widetilde{\mathbb{A}}$ by coincidence conditions. To this end, additional constraints are introduced into the variational statement by replacing the variational form of the kinematic condition (3.8) by

$$\underline{\mathsf{K}}(\eta,\nu,\alpha,z) := \underline{\mathsf{k}}(\eta,\nu), \qquad \forall \eta \in \mathbb{A}, \quad \nu := (\nu_1,\nu_2) \in \mathbb{R}^2, \tag{3.26a}$$

with

$$\underline{\mathsf{K}}(\eta,\nu,\alpha,z) := \mathsf{K}(\eta,\alpha,z) - \nu_2(\alpha(T) - z(T)) + \nu_1(\alpha(0) - z(0)), \qquad (3.26b)$$

$$\underline{\mathbf{k}}(\eta,\nu) := \mathbf{k}(\eta) - \nu_2 \alpha^0 + \nu_1 \alpha^0.$$
(3.26c)

The additional constraints enforce $\alpha(0) = z(0) + \alpha^0$ and $\alpha(T) = z(T) + \alpha^0$. We remark that with a proper definition of A and Z, this is already implied by the variational problem (3.10) and, therefore, the additional constraints do not essentially alter the continuum problem. However, the additional constraints do favourably change the discretization, as will be shown in the sequel. We refer to the added constraints as *coincidence conditions* and to the modified variational problem as the *augmented problem*, in conformity with Ref. [70].

To demonstrate the role of (3.26) for energy conservation, let us consider the discretization of the augmented problem with $\widetilde{\mathbb{P}} := \mathscr{P}^{n_{\mathbb{P}}} = \mathscr{P}^{n_{\mathbb{A}}-1}$. Essentially, the coincidence conditions at t = 0 and t = T replace the orthogonality conditions corresponding to $n_{\mathbb{A}} - 1$ and $n_{\mathbb{A}}$. This necessitates a redefinition of the test space associated with η by restricting η to $\mathscr{P}^{n_{\mathbb{A}}-2}$ instead of $\mathscr{P}^{n_{\mathbb{A}}} =: \widetilde{\mathbb{A}}$.

To demonstrate that the discretization of the augmented problem is energy conservative, we consider the energy-conservation residual according to (3.25b), invoke integration by parts for the second integral

$$\int_0^T \theta(\dot{\alpha} - \dot{z}) dt = \theta(\alpha - (z + \alpha^0)) \Big|_0^T - \int_0^T \dot{\theta}(\alpha - (z + \alpha^0)) dt, \qquad (3.27)$$

and obtain

$$\mathsf{R}_{3}(\mathbf{u},\alpha,\theta,z) = \mathsf{D}(\dot{\alpha},\theta,\mathbf{u},\alpha) - \underline{\mathsf{K}}(\dot{\theta},(\theta(0),\theta(T)),\alpha,z) + \underline{\mathsf{k}}(\dot{\theta},(\theta(0),\theta(T))). \quad (3.28)$$

D in (3.28) vanishes, because with $\alpha \in \mathscr{P}^{n_{\mathbb{A}}}$ and $\widetilde{\mathbb{P}} := \mathscr{P}^{n_{\mathbb{P}}} = \mathscr{P}^{n_{\mathbb{A}}-1}$, $\dot{\alpha}$ is in the test space associated with D, i.e., $\dot{\alpha} \in \mathscr{P}^{n_{\mathbb{A}}-1} =: \widetilde{\mathbb{P}}$. Moreover, $\dot{\theta} \in \mathscr{P}^{n_{\mathbb{P}}-1} = \mathscr{P}^{n_{\mathbb{A}}-2}$ and, obviously, $(\theta(0), \theta(T)) \in \mathbb{R}^2$. Hence, $\dot{\theta}$ and $(\theta(0), \theta(T))$ reside in the test spaces associated with <u>K</u> and <u>k</u>, which implies that <u>K</u> - <u>k</u> = 0 in (3.28). This corroborates that the energy-conservation residual R₃ vanishes, and proves that the discretization of the augmented problem is energy conservative, effectively bypassing the compatibility conditions.

We remark that compliance with (3.28) does not require z or $p(\mathbf{u}|_{\Gamma_{\alpha}})$ to be polynomial. Since this result extends to non-coincident grids, the interface mesh associated with α and θ can be chosen arbitrarily and, in particular, independent of the meshes for \mathbf{u} and z. The error in the transferred energy then cancels within each element of the interface mesh.

3.4.5 Coupling of trapezoidal with time-discontinuous Galerkin time integration

To investigate the implications that particular combinations of fluid and structure discretizations can have on the conservation properties at the fluid-structure interface, we consider an alternative discretization method for the structure, viz., the trapezoidal method, coupled to a time-discontinuous Galerkin discretization of the fluid equations. Before we assess conservation at the interface, we briefly consider conservation in the isolated structure subject to a forcing term. Conservation in the isolated fluid system has been considered in Section 3.4.2.

Conservation in the isolated structure

We assess the conservation properties of the isolated structure discretized with the trapezoidal method. Trapezoidal time integration is a popular method because of its simplicity. The trapezoidal method constitutes a specific instance of a Newmark method, and it is energy conservative for an unforced structure; see, e.g., Ref. [28]. To investigate energy conservation in the presence of a forcing term, we multiply the continuum equation (3.3) by \dot{z} , integrate over the generic time interval $[t_n, t_{n+1}]$ and invoke integration-by-parts on the left-hand side, which yields the following

energy-conservation statement:

$$\left[\frac{1}{2}M\dot{z}^{2} + \frac{1}{2}Kz^{2}\right]_{t_{n}}^{t_{n+1}} = \int_{t_{n}}^{t_{n+1}} \dot{z}\theta dt.$$
(3.29)

Denoting by $E := \frac{1}{2}M\dot{z}^2 + \frac{1}{2}Kz^2$ the sum of kinetic and strain energy, we can rewrite Eq. (3.29) as

$$E_{n+1} - E_n = \int_{t_n}^{t_{n+1}} \dot{z}\theta dt.$$
 (3.30)

Manipulation of the expressions for the trapezoidal method yields a relation for (\dot{z}_{n+1}, z_{n+1}) in terms of (\dot{z}_n, z_n) , from which the change in energy between time levels can be determined as

$$E_{n+1} - E_n = \Delta t \left(\frac{\dot{z}_{n+1} + \dot{z}_n}{2}\right) \left(\frac{\theta_{n+1} + \theta_n}{2}\right).$$
(3.31)

For conciseness, the derivation of Eq. (3.31) is deferred to Appendix A.1. From a comparison of Eqs. (3.30) and (3.31) it follows that the right member of Eq. (3.30) is approximated as

$$\int_{t_n}^{t_{n+1}} \dot{z}\theta dt \approx \Delta t \left(\frac{\dot{z}_{n+1} + \dot{z}_n}{2}\right) \left(\frac{\theta_{n+1} + \theta_n}{2}\right),\tag{3.32}$$

i.e., in the trapezoidal method the integral is approximated by the midpoint rule. Only in the specific case that θ is constant and \dot{z} is linear, or vice versa, is the integral evaluated exactly by the midpoint rule. We thus infer that trapezoidal time integration does generally not conserve energy in the presence of a forcing term. Moreover, the trapezoidal method can be shown generally not to conserve momentum for forced structures; the elaboration is presented in Appendix A.2.

Conservation at the interface

More specifically, we are interested in the conservation properties of the trapezoidal time-integration method for the structure when coupled to the time-discontinuous Galerkin fluid discretization. This combination naturally generates an incompatibility between fluid and structure representations at the interface, because the trapezoidal method lacks a proper functional representation in between time levels as, in fact, does any finite-difference scheme. Indeed, the representation of a function by the trapezoidal method is restricted to the discrete time levels. This motivates us to identify the pressures $p(\mathbf{u}|_{\Gamma_{\alpha}}), \theta$ and the velocities $v|_{\Gamma_{\alpha}} = \dot{\alpha}, \dot{z}$ on both sides of the interface at discrete time levels. The coupling to the time-discontinuous Galerkin fluid discretization implies that the average acceleration in the trapezoidal method is computed from a discontinuous forcing term, θ . Displacement and velocity in the trapezoidal method, however, are continuous from

one time-slab to the next and, consequently, kinetic and strain energy of the structure are conserved from one time-slab to the next.

Let us now consider the conservation properties at the fluid-structure interface. Provided that $v|_{\Gamma_{\alpha}} = \dot{\alpha}$ holds, mass conservation of the fluid at the interface is immediate. To assess energy conservation at the interface, the change in energy within a time-slab is compared on either side of the interface. For the structure integrated by trapezoidal time integration we refer to Eq. (3.31). For the fluid we assume a discretization in primitive variables with linear-in-time shape functions for velocity v and pressure p. The change in energy within a time-slab $[t_n, t_{n+1}]$ on the fluid side of the interface can then be computed as

$$E_{n+1} - E_n = \int_{t_n}^{t_{n+1}} pv dt$$

$$= \frac{1}{6} \Delta t \left(2p_{n+1}v_{n+1} + p_{n+1}v_n + p_nv_{n+1} + 2p_nv_n \right),$$
(3.33)

with $v := v|_{\Gamma_{\alpha}}$ and $p := p(\mathbf{u}|_{\Gamma_{\alpha}})$ for brevity. Upon identification of corresponding quantities on either side of the interface, i.e., $\dot{z} = v$ and $\theta = p$, we can deduce from a comparison of Eqs. (3.31) and (3.33) that the calculated change in energy across a time-slab on the structure side is different from that on the fluid side. Therefore, the coupling of trapezoidal time integration for the structure to a time-discontinuous Galerkin discretization for the fluid does not generally conserve energy at the interface. Moreover, this pairing of discretization methods does not conserve momentum at the interface either, since the trapezoidal method does generally not conserve momentum for forced structures; cf. Appendix A.2.

3.5 Numerical experiments

To assess the relevance of maintaining the conservation properties at the fluidstructure interface, we conduct numerical experiments on the piston problem. Section 3.5.1 specifies the setup of the numerical experiments. In Section 3.5.2, we compare conservative and non-conservative discretizations. In Section 3.5.3, we investigate the implications of different pairings of fluid and structure discretizations for the conservation properties at the interface.

3.5.1 Experimental setup

Time-discontinuous Galerkin discretization of the fluid equations

For the fluid discretization we use throughout a standard time-discontinuous Galerkin method with piecewise-linear trial and test functions, which are continuous in space but discontinuous in time at the space/time slab interfaces. This separates the computational task into the sequential solution of multiple space/time slabs. Information is propagated across space/time slab interfaces by weakly enforced initial conditions. On account of the linear-in-time basis functions, the fluid discretization is formally second-order time accurate.

The conservation form of the fluid equations is discretized in primitive variables, (p, v, T) with T denoting the absolute temperature, which simplifies the implementation of boundary conditions and, in particular, the coupling to the structure; see Ref. [44] for details. For the discretization in primitive variables we refer to Ref. [25]. The formulation can be augmented with a least-squares term to improve the stability of the Galerkin formulation. For the variational formulation and discretization of the fluid equations, and for the definition of the least-squares operator we refer to Refs. [25, 62, 65].

Velocity-based time-discontinuous Galerkin structural time integration

To treat fluid and structure in a unified framework and simplify fluid-structure coupling, we use a velocity-based time-discontinuous Galerkin discretization of the structure. The discretization employs piecewise-linear test and trial functions, which are discontinuous in time at the time-slab interfaces. We shall use plus and minus signs to indicate from which side the time-slab interface is approached.

A velocity-based formulation of the structure equation (3.3) is used

$$M\dot{v} + K\left(z_{t=0} + \int_{t=0}^{t} v(\xi) \mathrm{d}\xi\right) = \theta - \beta, \qquad (3.34)$$

with v denoting the velocity of the structure. To phrase Eq. (3.34) into variational form, we multiply Eq. (3.34) by a test function, integrate the resulting expression over the time domain, and carry out integration by parts. The variational form of the initial-value problem for the structure for a single time-slab $[t_n^+, t_{n+1}^-]$ can then be stated as: Find a linear function v such that

$$M\lambda(t_{n+1}^{-})v(t_{n+1}^{-}) - M \int_{t_{n}^{+}}^{t_{n+1}} \dot{\lambda}v \, dt + K \int_{t_{n}^{+}}^{t_{n+1}} \lambda \left(\int_{t_{n}^{+}}^{t} v(\xi)d\xi\right) \, dt$$

$$= M\lambda(t_{n}^{+})v(t_{n}^{-}) - Kz(t_{n}^{-}) \int_{t_{n}^{+}}^{t_{n+1}^{-}} \lambda \, dt + \int_{t_{n}^{+}}^{t_{n+1}^{-}} \lambda(\theta - \beta) \, dt$$
(3.35)

for all linear test functions λ . Note that Eq. (3.35) incorporates the jump term, $\lambda(t_n^+)(v(t_n^+) - v(t_n^-)) = 0$, which enforces the initial condition for velocity weakly. In contrast, the initial condition for displacement, $z(t_n^+) = z(t_n^-)$, is enforced strongly. Hence, velocity is discontinuous at the time-slab interfaces, whereas displacement is continuous. By virtue of the coincidence conditions from Section 3.4.4, a continuous structure displacement implies a continuous displacement of the fluid boundary, which facilitates the treatment of the moving fluid boundary.

Analytic structural time integration

Below, we derive the analytic solution of the structure equation, which we use in Section 3.5.2 for the time integration of the structure. This allows us to exclude discretization errors in the structure.

The motion of the piston is described by the initial-value problem, Eqs. (3.3)–(3.4). As specified above, we use for the approximation of the fluid pressure shape functions which are piecewise linear in time. Accordingly, we approximate the forcing term of the structure, $\theta(t) - \beta =: \pi(t)$, by a piecewise-linear function. The variation of the forcing term within the generic time interval $t \in [0, T]$ is then of the form

$$\pi = \pi_0 + \frac{t}{T}(\pi_1 - \pi_0), \qquad (3.36)$$

where the subscripts zero and one refer to the values at the beginning and end of the time interval, respectively. For the simple right-hand-side expression (3.36), the second-order ordinary differential equation can be solved analytically. The solution of the initial-value problem, Eqs. (3.3)-(3.4), subject to Eq. (3.36), is then

$$z(t) = \left(z_0 - \frac{\pi_0}{K}\right)\cos\left(\sqrt{\frac{K}{M}}t\right) + \frac{\left(\dot{z}_0 - \frac{\pi_1 - \pi_0}{TK}\right)}{\sqrt{\frac{K}{M}}}\sin\left(\sqrt{\frac{K}{M}}t\right) + \frac{\pi_0 + \frac{t}{T}(\pi_1 - \pi_0)}{K}.$$
(3.37)

Given the values of the forcing term at either time level, π_0 and π_1 , and the initial conditions z_0 and \dot{z}_0 , the structural displacement can be extracted from Eq. (3.37).

Parameters of the discretization

To investigate the effect of the discretization error induced by the fluid-structure coupling, we have to ensure that the discretization errors in fluid and structure are sufficiently small. To this end, we use analytic structural time integration and fluid subcycling, i.e., the fluid discretization uses time steps $\Delta t_{\rm f}$, which are smaller than the time steps of the structural discretization, $\Delta t_{\rm s}$. The ratio $\kappa = \Delta t_{\rm s}/\Delta t_{\rm f}$, $\kappa \in \mathbb{N}$ is called the subcycling factor. In the numerical experiments we use a subcycling factor of $\kappa = 8$. We note that fluid subcycling induces non-matching meshes at the fluid-structure interface and, hence, directly translates into an incompatibility between fluid and structure. However, conservation can still be maintained under the conditions stated in Section 3.4.4.

To render the numerical evaluation error smaller than the errors due to discretization and (non-)conservative coupling at the interface, we set strict convergence tolerances for the iterative solution of the fluid subsystem and for the aggregated fluid-structure system. The discretized fluid-structure system is solved by subiteration, i.e., by solving fluid and structure alternately subject to complementary partitions of the interface conditions; cf. Section 4.2.2 for details. This process is repeated until convergence.

Parameters of the piston problem

The physical parameters of the piston problem are specified in terms of dimensionless parameters according to Section 3.3.2 and listed in Table 3.1.

Table 3.1:	Physical	parameters of	of the	piston	problem.
------------	----------	---------------	--------	--------	----------

z_0	\dot{z}_0	μ	a	b
10^{-2}	0	1.63	0.34	10^{-1}

The dimensionless quantities are chosen such that an oscillation period is equal to unit time. The fluid-to-structure time-scale ratio b indicates that the characteristic time-scale for the fluid is by an order of magnitude smaller than the one for the structure. This implies that the fluid behaviour can be considered as quasi-steady. Moreover, the fluid-to-structure mass ratio is of order one, which implies that fluid and structure contribute equally to the dynamics of the system. The computation is started from initial conditions which are derived from the solution of the linearized problem (3.19) with an initial structural displacement of $z_0 = 10^{-2}$.

3.5.2 Conservative versus non-conservative discretization

To investigate the effect of maintaining conservation at the interface on the numerical solution, we compare results for conservative and non-conservative discretizations, following our exposition in Ref. [44]. To distinguish the error incurred by the interface coupling from fluid and structure discretization errors, we render the latter sufficiently small by means of analytic structural time integration (cf. Section 3.5.1) and subcycling in the fluid.

Coupling methods

We consider three coupling methods, which all conserve mass at the interface, but differ in their conservation properties for energy and momentum. Nevertheless, all three discretizations are consistent and second-order time accurate, as will be shown in the sequel. Therefore, preference of one method over another can only be motivated on the basis of conservation properties.

• method A :

This method discretizes the augmented variational problem based on the coincidence conditions (3.26). Therefore, the fluid-structure coupling is fully conservative.

 $\bullet \mod B$:

This method is distinguished by the following form of the dynamic condition: $\theta(t)$ within a time interval $t \in [0, T]$ is calculated in an ad-hoc but consistent way by using a linear interpolation that coincides with $p(\mathbf{u}|_{\Gamma_{\alpha}})$ at the beginning and end of the time interval, i.e.,

$$\theta(t) = p(\mathbf{u}|_{\Gamma_{\alpha},t=0}) + \frac{t}{T} \left(p(\mathbf{u}|_{\Gamma_{\alpha},t=T}) - p(\mathbf{u}|_{\Gamma_{\alpha},t=0}) \right).$$
(3.38)

This finite-difference discretization is typically used in combination with subiteration methods; see, e.g., Ref. [1]. However, this discretization does neither conserve energy nor momentum at the fluid-structure interface.

• method C :

This method is distinguished by replacing the kinematic condition (3.5b) with a modification that is consistent with the original condition. For the discrete fluid boundary displacement, α , a quadratic function is assumed, which is determined from the requirement that α is continuous from one fluid time-slab to the next, in combination with the following orthogonality condition:

$$\int_{0}^{T} \eta \left((\alpha - \alpha_0) - (z - z_0) \right) dt = 0 \qquad \forall \eta \in \text{span}\{1, t\},$$
(3.39)

where the subscript zero refers to the values at the beginning of a time interval. Fluid boundary and structure only coincide initially at time t = 0, but can deviate henceforth. This discretization does not conserve energy. However, it conserves mass and momentum at the interface, because this discretization satisfies the interface conditions (3.5a) and (3.5c) pointwise.

As we employ fluid subcycling, interpolation of the kinematic boundary conditions for the fluid, i.e., the structural displacement and velocity, is necessary at all time levels of the fluid which do not coincide with the time levels of the structure, i.e., at $\Delta t_{\rm f}, \ldots, (\kappa - 1)\Delta t_{\rm f}$. The interpolation of the structural solution is done according to the discretized form of Eq. (3.26) for methods A and B; for method C, Eq. (3.39) is employed.

Numerical results

Fig. 3.2 compares the numerical solution for the structural displacement obtained with the coupling methods A, B and C. An oscillation period is resolved by 10 structural time steps or, equivalently, by 80 time steps in the fluid. The figure clearly shows that the amplitude of the non-conservative method B decays rapidly, whereas the amplitude of the conservative method A remains apparently unchanged. The modulation of the solution computed with method B can be attributed to interfering waves : The initial conditions correspond to a periodic solution of the continuum problem, but do not generally yield a periodic discrete solution. Moreover, we notice a considerable phase error in the solution computed with method B. The amplitude computed with method C grows in time, which indicates instability of the numerical solution.



Figure 3.2: Structural displacement versus time: conservative discretization method A (\longrightarrow) and non-conservative methods B (--) and C (\cdots).

The results of a mesh refinement study are shown in Fig. 3.3. The L^2 norm of the error in the computed structural displacement is plotted versus the size of the structural time step for $\Delta t_s = 2^{-3}, 2^{-4}, \ldots, 2^{-7}$. The reference solution was obtained with the conservative method A for a structural time step $\Delta t_s = 2^{-8}$. Fig. 3.3 confirms that the methods are second-order time accurate. Moreover, it is apparent that the conservative method A is by more than two orders of magnitude more accurate than the non-conservative methods B and C. Hence, for a given level of accuracy, the conservative method can afford much larger time steps than the non-conservative methods. We infer that maintaining conservation at the interface yields an improvement in the accuracy of the numerical solution. Note that an additional computational cost is not incurred.

3.5.3 Coupling of trapezoidal with time-discontinuous Galerkin time integration

This section illustrates the implications of specific pairings of fluid and structure discretizations for the conservation properties at the interface. We consider the coupling of the time-discontinuous Galerkin discretization of the fluid equations with the time-discontinuous Galerkin and the trapezoidal method for the structural time integration, and investigate the effect on the solution accuracy.

We recall from Section 3.4 that with the time-discontinuous Galerkin method for the structural time integration, momentum and energy at the fluid-structure interface can be conserved. However, this method loses energy in between timeslabs due to the discontinuity in velocity. With the trapezoidal discretization for



Figure 3.3: Error in the structure displacement versus structural time-step size: conservative discretization method A (--) and non-conservative methods B (--) and C (\cdots) .

the structure, neither momentum nor energy is conserved at the interface.

Grid refinement study

To assess the temporal accuracy of the different pairings of fluid and structure discretizations, we determine the observed order of temporal accuracy, ζ , by

$$\zeta = \ln\left(\frac{||z_{4\tau} - z_{2\tau}||}{||z_{2\tau} - z_{1\tau}||}\right) / \ln(2), \tag{3.40}$$

where z denotes the computed structural displacement on meshes of different structural time-step sizes (denoted by subscript τ , 2τ and 4τ , respectively). The timestep size in the fluid is set eight times smaller than the one in the structure, i.e., $\kappa = 8$. The differences are measured in the L^2 norm.

As the fluid equations are formulated and discretized in space/time, refinement of the temporal mesh width requires a corresponding refinement of the spatial mesh width to obtain second-order convergence. If this is not done, the truncation error retains a mixed term, $\Delta t_f \Delta x$, which will produce only first-order convergence. For the current computations, the non-dimensional spatial mesh width in the fluid was taken equal to the temporal mesh width.

The computations were performed on a sequence of meshes of different timestep sizes. The observed order of temporal accuracy for trapezoidal and timediscontinuous Galerkin time integration is shown in Table 3.2. The table indicates that the trapezoidal time integration gives second-order time-accurate results. This was to be expected, as both the trapezoidal method and the fluid discretization are second-order time accurate, see Ref. [28]. The time-discontinuous Galerkin discretization for the structure uses linear-in-time shape functions for the velocity. The structural displacement is obtained by integration. Therefore, this method is third-order time accurate in displacement. Indeed, we can observe third-order convergence for the Galerkin time integration, because the discretization errors in the fluid and at the interface are smaller than the discretization error in the structure due to fluid subcycling and conservative fluid-structure coupling, respectively. However, on even finer mesh sequences second-order time-accuracy is recovered, manifesting the second-order time-accuracy of the fluid discretization.

Table 3.2: Observed order of temporal accuracy ζ according to (3.40): trapezoidal versus time-discontinuous Galerkin structural time integration coupled to a time-discontinuous Galerkin discretization of the fluid.

mesh sequences	trapezoidal method	time-discont. Galerkin method		
with time-step sizes				
$2^{-3}, 2^{-4}, 2^{-5}$	1.9541	3.0150		
$2^{-4}, 2^{-5}, 2^{-6}$	2.0088	2.9925		
$2^{-5}, 2^{-6}, 2^{-7}$	2.0019	3.0035		

The different orders of the structural time integration methods also become apparent in Fig. 3.4 on the following page, in which the error is plotted versus the size of the structural time step for $\tau = 2^{-3}, 2^{-4}, \ldots, 2^{-7}$. We use the same reference solution as in Section 3.5.2. The errors corresponding to the Galerkin method and to the trapezoidal method are markedly different. This difference emanates partly from the different structure discretization and partly from the different interface treatment. As the interface coupling is formally only secondorder accurate, the observed third-order convergence behaviour of the Galerkin method indicates that conservative interface coupling introduces an error that is typically much smaller than the structural discretization error. Moreover, a comparison of Fig. 3.3 and Fig. 3.4 shows that the error in the fluid-structure coupling with the trapezoidal structural discretization is comparable to the error in the non-conservative methods. Fig. 3.4 indicates that the error incurred by non-conservative coupling at the interface can dominate the discretization errors in fluid and structure.

3.6 Concluding remarks

We established the conservation properties at the fluid-structure interface of the continuum system and of its discretization by the finite-element method on the basis of a prototypical fluid-structure system. Energy conservation in the discrete model is only trivially maintained under restrictive compatibility conditions on



Figure 3.4: Error in the structure displacement versus structural timestep size: time-discontinuous Galerkin time integration (--) and trapezoidal method (--).

the approximation spaces of fluid and structure, viz., matching meshes and identical orders of approximation at the interface. Such conditions are prohibitively restrictive in practice, since they impede tailoring fluid and structure discretizations to their specific resolution requirements. To circumvent the compatibility conditions, we considered a modified discretization which is based on orthogonality and coincidence conditions. This modified discretization enables conservation also for incompatible approximations. Moreover, we investigated the implications of specific pairings of fluid and structure discretization methods for the conservation properties at the interface. We showed that trapezoidal time integration does generally not conserve momentum and energy for forced structures. Hence, the trapezoidal method does not admit a conservative fluid-structure coupling. In contrast, with a time-discontinuous Galerkin discretization for both the fluid and the structure, the conservation properties at the interface can be maintained. However, the structure discretization by the time-discontinuous Galerkin method does not conserve momentum and energy between time slabs due to the weak enforcement of the initial conditions.

To assess the relevance of maintaining the conservation properties under discretization, we conducted numerical experiments on the one-dimensional piston model problem. A comparison of conservative and non-conservative discretizations shows that the solution delivered by a conservative discretization is by several orders of magnitude more accurate. Moreover, we demonstrated that a discretization that violates energy conservation can induce instability of the numerical solution. Our experiments with different pairings of fluid and structure discretization methods indicated that the non-conservative coupling of the trapezoidal method for the structure with a time-discontinuous Galerkin fluid discretization induces an error that can dominate the discretization errors. In contrast, conservative coupling of a time-discontinuous Galerkin method for both the fluid and the structure introduces an error that is typically much smaller than other discretization errors. We infer that maintaining the conservation properties at the interface improves the accuracy and stability of the numerical solution; an additional computational cost is not incurred.

As a prospect, let us briefly give an indication how the considered energyconservation concept extends to the multi-dimensional case. Generally, maintaining energy conservation for incompatible approximations in multiple dimensions is much more involved. However, it is anticipated that energy conservation in multiple dimensions can be ensured under appropriate orthogonality conditions on the interface spaces. In particular, the orthogonality of the pressure difference to the interface velocity, and the orthogonality of the velocity difference to the interface pressure are sufficient conditions for warranting energy conservation in multiple dimensions. We expect that maintaining conservation in the multi-dimensional case is similarly beneficial for the stability and accuracy of the numerical solution as demonstrated in our numerical experiments for the one-dimensional case.

Chapter 4 Efficiency of conventional solution methods

4.1 Introduction

Numerical solution methods for fluid-structure interaction typically employ partitioning. In a partitioned method, the fluid and structure equations are separately integrated in time subject to complementary partitions of the interface conditions, i.e., kinematic and dynamic interface conditions are enforced asynchronously; see, e.g., Refs. [1, 20, 56, 73]. If only a single fluid and structure solution per time step are carried out, such partitioned methods are commonly referred to as *looselycoupled partitioned methods*; see, e.g., Refs. [18, 55]. The benefits and deficiencies of loosely-coupled methods have been discussed in the review article by Felippa et al. [20]. Their essential disadvantage pertains to the inherent loss of the conservation properties of the continuum fluid-structure system. In particular, they only satisfy conservation in an asymptotic sense, i.e., for vanishing mesh width; this is a basic consistency requirement. Although the order of the incurred error can be improved by predictors (see Refs. [53–55]), loosely-coupled methods can never be exactly conservative. In general, they are energy increasing and, hence, numerically unstable; see, e.g., Refs. [53, 55, 56].

The deficiencies of loosely-coupled methods have motivated the investigation of methods that treat the interaction between the fluid and the structure synchronously. This can be achieved by *monolithic methods* which solve fluid and structure equations simultaneously [6, 27, 34] and, more commonly, by partitioned methods which solve the fluid-structure system by repeating within a time step alternate fluid and structure solutions until convergence [1, 36, 47]. Such partitioned methods are commonly referred to as *strongly-coupled partitioned methods*.

In contrast to monolithic methods, partitioned methods in general allow the use of separate fluid and structure solvers and, hence, enable software modularity. This constitutes an important practical asset of partitioned methods. Moreover, the solution methods can be distinguished according to their ability to maintain the conservation properties at the fluid-structure interface. A conservative discretization is necessarily fully implicit (cf. Section 3.4). Hence, maintaining conservation requires to resolve fluid and structure equations synchronously. This implies that loosely-coupled partitioned methods are inherently non-conservative. In contrast, monolithic and strongly-coupled partitioned methods can in principle maintain the conservation properties. The relevance of conservation for the stability and accuracy of the numerical method has been demonstrated in Chapter 3.

Monolithic, loosely-coupled and strongly-coupled partitioned methods are controversially advocated in literature. The stability of monolithic and partitioned methods is compared in Refs. [6, 34], for instance. However, preference of one method over another cannot be based on stability considerations only, but must also consider arguments of accuracy, computational cost, efficiency and practicality. We address these issues in this chapter and contrast the merits and drawbacks of the respective solution methods, following our elaboration in Refs. [39, 41, 44]. Moreover, we discuss techniques to further increase the efficiency of the solution methods.

We assess the properties of the respective solution methods by means of numerical experiments on the piston model problem from Section 3.2. Fluid subcycling and analytic structural time integration allow us to separate the effect of the numerical solution error from the discretization errors.

This chapter is organized as follows: In Section 4.2, we review monolithic and partitioned solution methods and discuss their advantages and disadvantages. Moreover, we discuss concepts to further reduce the computational cost of partitioned solution methods. In Section 4.3, we present numerical experiments and results. Section 4.4 contains concluding remarks.

4.2 Monolithic versus partitioned solution methods

In this section, we review and discuss the merits and drawbacks of the different solution methods. Numerical methods for fluid-structure interaction are commonly classified into monolithic and partitioned methods. Unfortunately, in literature, this nomenclature is not used in a consistent way. This gives rise to misunderstanding and confusion. Therefore, we shall specify our nomenclature first.

The discretization of the aggregated equations of fluid, structure and interface conditions, Eq. (2.27), results in a fully-implicit system of algebraic equations. Such an algebraic system can be solved with a monolithic or with a partitioned solution procedure. We refer to solution procedures that treat the coupling between fluid and structure simultaneously by operating directly on the system of aggregated equations as *monolithic methods*. Alternatively, the algebraic system can be solved by invoking partitioning, i.e., by solving fluid and structure equations separately, subject to complementary partitions of the interface conditions. Thus, *partitioned solution procedures* treat the coupling between fluid and structure *asynchronously*, which induces a lag between fluid and structure solution. From this perspective, partitioned procedures can be conceived as solving a system of *segregated equations*, which differs from the aggregated equations by the induced lag. However, by repeating within a time step alternate fluid and structure solutions until convergence, the lag can be eliminated, fluid and structure solution can be synchronized and the solution of the aggregated equations can be recovered. This iterative process is commonly referred to as *subiteration* and such partitioned procedures are generally called *strongly-coupled partitioned procedures*. Here, we shall adhere to this nomenclature.

In contrast to strongly-coupled procedures, *loosely-coupled partitioned procedures* carry out a single fluid and structure solution per time step only and, thus, do not eliminate the lag between the fluid and the structure solution. Therefore, loosely-coupled methods do not resolve the aggregated equations. The lag can be conceived as a numerical evaluation error, but it also admits a reinterpretation as a discretization error associated with the set of segregated equations. The order of the incurred error can be improved by means of predictors; see, e.g., Refs. [53–55].

In this section, we examine the advantages and drawbacks of the different solution methods specified above. To this end, Sections 4.2.1–4.2.3 introduce monolithic methods and strongly-coupled and loosely-coupled partitioned methods, respectively. Section 4.2.4 discusses prediction techniques. Section 4.2.5 contrasts strongly-coupled and loosely-coupled partitioned methods. Section 4.2.6 elaborates techniques to improve the computational efficiency of strongly-coupled methods.

4.2.1 Monolithic methods

Monolithic solution methods treat the coupled fluid and structure equations simultaneously, i.e., they directly operate on the aggregated fluid and structure equations. As this system is in general nonlinear, the solution procedure typically involves a Newton process. Below, we discuss the implications which arise from the simultaneous treatment of the coupled fluid and structure equations.

The application of Newton's method to the nonlinear variational problem of the generic form (2.27a) requires the specification of the Fréchet derivative of the functional $\mathsf{P}: \mathbb{W} \times \mathbb{Q} \mapsto \mathbb{R}$ with respect to its nonlinear argument. For any fixed $w_0 \in \mathbb{W}$, $\mathsf{P}(w_0, \cdot)$ is a nonlinear functional on \mathbb{Q} . The functional $\mathsf{P}(w_0, \cdot)$ is Fréchet differentiable at $q_0 \in \mathbb{Q}$ if there exists a bounded linear functional $\mathsf{P}'(w_0, q_0, \cdot): \mathbb{Q} \mapsto \mathbb{R}$ such that

$$\lim_{q' \to 0} \frac{|\mathsf{P}(w_0, q_0 + q') - \mathsf{P}(w_0, q_0) - \mathsf{P}'(w_0, q_0, q')|}{\|q'\|_{\mathbb{O}}} = 0.$$
(4.1)

The functional $\mathsf{P}'(w_0, q_0, \cdot)$ is then called the Fréchet derivative of P to q at (w_0, q_0) . Assuming that the functional $\mathsf{P}'(w_0, q_0, \cdot)$ exists for all $w_0 \in \mathbb{W}, q_0 \in \mathbb{Q}$, we identify P' with a functional on $\mathbb{W} \times \mathbb{Q} \times \mathbb{Q}$. With this definition of the Fréchet derivative P', and provided with an initial approximation $q_0 \in \mathbb{Q}$, Newton's method for the variational problem (2.27a) is defined as the following iterative process: for $j = 1, 2, \ldots$, find

$$q_j \in \mathbb{Q}: \quad \mathsf{P}'(w, q_{j-1}, q_j - q_{j-1}) = \mathsf{p}(w) - \mathsf{P}(w, q_{j-1}) \quad \forall w \in \mathbb{W}.$$

$$(4.2)$$

If the initial estimate q_0 is sufficiently close to the actual solution, then q_j converges to the solution of the variational problem as $j \to \infty$.

Ref. [69] points out three essential problems that arise from the synchronous treatment of fluid and structure equations. These complications derive from the free-boundary character of the problem and from the inherent interdependence between the fluid and structure state variables. To specify these problems, we expand the derivative P' according to the definition of P in (2.27b):

$$\mathsf{P}'(w,q,q') := \mathsf{F}'_{\mathbf{u}}(\mathbf{v},\mathbf{u},\alpha,\mathbf{u}') + \mathsf{F}'_{\alpha}(\mathbf{v},\mathbf{u},\alpha,\alpha') + \mathsf{S}'_{z}(\lambda,z,\pi,z') + \mathsf{S}'_{\pi}(\lambda,z,\pi,\pi') + \mathsf{K}(\eta,\alpha',z') + \mathsf{D}'_{\mathbf{u}}(\phi,\mathbf{u},\alpha,\mathbf{u}') + \mathsf{D}'_{\alpha}(\phi,\mathbf{u},\alpha,\alpha') - \langle\phi,\pi'\rangle$$
(4.3)

with $\langle \cdot, \cdot \rangle$ the standard L^2 inner product. Moreover, $\mathsf{F}'_{\mathbf{u}}$ and F'_{α} denote the Fréchet derivatives of the functional F with respect to its nonlinear arguments \mathbf{u} and α , respectively. Likewise, S'_z and S'_{π} are the Fréchet derivatives of S with respect to z and π . Eq. (4.3) also involves the L^2 inner products associated with the linear functional K . The semi-linear functional D yields the derivatives $\mathsf{D}'_{\mathbf{u}}$ and D'_{α} and the L^2 inner product $\langle \phi, \pi' \rangle$. The functionals K and D are separable in accordance with our stipulation in Section 2.5.

The first problem pertains to the derivative F'_{α} , which is commonly referred to as the *shape derivative*. This derivative is induced by the interdependence of the fluid state variables and the domain Ω_{α} on which these are defined. Discretization methods such as the finite-element method typically use boundary-fitted meshes. This implies that a perturbation of the interface position induces, in principle, a deformation of the mesh throughout the entire computational domain. This would render the computational cost incurred in the evaluation of the shape derivative F'_{α} prohibitive in practical applications. Although there exist approaches to facilitate the evaluation of the shape derivative, such as the *method of spines* [26], their applicability is typically restricted.

The second problem is rooted in the inherent interdependence between the fluid and the structure solution. This interconnection is illustrated in the connectivity Table 4.1 associated with the variational statement (2.27) and is manifest through the derivatives of the interface conditions in (4.3). This renders the Jacobian matrix \mathbf{P} associated with the discrete approximation of the operator \mathbf{P}' in (4.2) inseparable, which necessitates a simultaneous solution of fluid, structure and interface conditions. This impedes software modularity, see also Ref. [20], and constitutes a severe practical disadvantage.

The third problem emanates from the disparate time and length scales inherent in the fluid and structure subsystems, which generally results in severe ill-conditioning of the Jacobian matrix \mathbf{P} .

	α	u	π	z
Κ	×	0	0	\times
F	×	×	0	0
D	×	\times	×	0
S	0	0	×	×

 Table 4.1: Illustration of the connectivity in fluid-structure interaction.

In conclusion, the simultaneous (or *monolithic*) solution of the aggregated equations by Newton's method encounters several severe disadvantages. These disadvantages render a monolithic solution by Newton's method prohibitive for actual fluid-structure-interaction problems.

4.2.2 Strongly-coupled partitioned methods

The complications in Newton's method can be effectively circumvented by means of *partitioning*, i.e., by solving fluid and structure equations separately subject to complementary partitions of the interface conditions. *Strongly-coupled partitioned methods* repeat within a time step alternate fluid and structure solution. This iterative process is commonly called *subiteration*. Convergence of the subiteration process is contingent on the spectral radius; see the analysis in Section 6.3.3. In this chapter, we shall restrict ourselves to cases in which subiteration converges.

Provided with an initial approximation $z_0 \in \mathbb{Z}$ of the structure solution, or, in particular, of the structure displacement at the interface $(z|_{\Theta})_0$, the following steps define the subiteration process for the nonlinear variational problem (2.27): for j = 1, 2, ...

- (S1) Solve the kinematic condition: $\alpha_j \in \mathbb{A}$: $\mathsf{K}(\eta, \alpha_j, z_{j-1}) = \mathsf{k}(\eta) \quad \forall \eta \in \mathbb{H}$
- (S2) Solve the fluid: $\mathbf{u}_j \in \mathbb{U}$: $\mathsf{F}(\mathbf{v}, \mathbf{u}_j, \alpha_j) = \mathsf{f}(\mathbf{v}) \quad \forall \mathbf{v} \in \mathbb{V}$
- (S3) Solve the dynamic condition: $\pi_i \in \mathbb{P}$: $\mathsf{D}(\phi, \pi_i, \mathbf{u}_i, \alpha_i) = 0 \quad \forall \phi \in \mathbb{F}$
- (S4) Solve the structure: $z_j \in \mathbb{Z}$: $\mathsf{S}(\lambda, z_j, \pi_j) = \mathsf{s}(\lambda) \quad \forall \lambda \in \mathbb{L}$

Note that this procedure bypasses the computation of the shape derivative. Moreover, it obviates the simultaneous treatment of the fluid and the structure. In particular, subiteration reduces the complexity of solving the aggregated fluidstructure equations to a sequence of 'standard' problems. For instance, (S2) with α_j given, and (S4) with π_j given represent valid fluid and structure problems separately. Hence, subiteration enables software modularity. Upon convergence of the subiteration process, the solution of the aggregated equations (2.27) is obtained. We note that the subiteration algorithm admits a description in terms of a generic fluid-structure-interaction problem in conformity with Eq. (2.27), which implies that the method is, in principle, applicable to any fluid-structure-interaction problem. Moreover, since the algorithm can be described in a continuum variational setting, the convergence behaviour of the subiteration method is asymptotically independent of the underlying discretization, i.e., for all sufficiently fine discretizations the convergence behaviour will be identical; see also Ref. [71]. For applications of the subiteration method to fluid-structure-interaction problems see, e.g., Refs. [1, 36, 47]. We remark that subiteration is often alternatively referred to as inner-iteration, Picard iteration or successive approximation.

It is elucidating to construe this method as a defect correction method [7], following Ref. [69]. The defect correction method for a variational problem of the form (2.27a) first determines an initial estimate $q_0 \in \mathbb{Q}$ such that $\tilde{\mathsf{P}}(w, q_0) =$ $\mathsf{p}(w)$ for all $w \in \mathbb{W}$, where the functional $\tilde{\mathsf{P}} : \mathbb{W} \times \mathbb{Q} \to \mathbb{R}$ denotes a suitable approximation to P . The defect correction method then determines a sequence of iterates $\{q_i\} \in \mathbb{Q}$ such that

$$\tilde{\mathsf{P}}(w,q_j) = \mathsf{p}(w) + \tilde{\mathsf{P}}(w,q_{j-1}) - \mathsf{P}(w,q_{j-1}) \qquad \forall w \in \mathbb{W},$$
(4.4)

for j = 1, 2, ... If $\tilde{\mathsf{P}}$ is sufficiently close to P , then $\{q_j\}$ converges to the actual solution of the variational problem as $j \to \infty$. To identify the approximate functional $\tilde{\mathsf{P}}$ associated with the subiteration method, we note that the approximations generated by (S1)–(S4) comply with

$$\mathsf{K}(\eta,\alpha_j,z_{j-1}) + \mathsf{F}(\mathbf{v},\mathbf{u}_j,\alpha_j) + \mathsf{D}(\phi,\pi_j,\mathbf{u}_j,\alpha_j) + \mathsf{S}(\lambda,z_j,\pi_j) = \mathsf{k}(\eta) + \mathsf{f}(\mathbf{v}) + \mathsf{s}(\lambda),$$
(4.5)

for all admissible $(\mathbf{v}, \eta, \lambda, \phi)$. Expanding the functional K according to its definition below (2.25) and adding suitable partitions of zero yields (4.4) with the approximate functional

$$\mathsf{P}((\mathbf{v},\eta,\lambda,\phi),(\mathbf{u},\alpha,z,\pi)) := \mathsf{F}(\mathbf{v},\mathbf{u},\alpha) + \mathsf{D}(\phi,\pi,\mathbf{u},\alpha) + \mathsf{S}(\lambda,z,\pi) + \langle \eta,\alpha \rangle.$$
(4.6)

The defect can then be identified from Eqs. (2.27b) and (4.6) as

$$\mathsf{P}(w,q) - \mathsf{P}(w,q) = -\langle \eta, z \rangle. \tag{4.7}$$

Note that the approximation only involves the functional K corresponding to the kinematic interface condition. The connectivity table corresponding to the approximate operator is identical to Table 4.1, but void of the right-upper entry (K, z). This renders the connectivity table lower-triangular, which enables one to perform the inversion of the approximate operator \tilde{P} conveniently by forward substitution. The subiteration process therefore essentially corresponds to a block Gauss-Seidel iteration.

4.2.3 Loosely-coupled partitioned methods

The essential difference between strongly-coupled and loosely-coupled partitioned methods is that the latter do not repeat subiterations, but only perform a single fluid and structure solution per time step. The asynchronous enforcement of the interface conditions then implies a lag between the fluid and the structure solution. In the context of strong coupling, this lag can be conceived as a numerical evaluation error. Solving the exact (i.e., aggregated) equations approximately can be reinterpreted as solving a set of approximate (i.e., segregated) equations exactly. Thus, we can construe loosely-coupled methods as solving a set of segregated equations given by Eq. (4.5) as opposed to the aggregated equations (2.27). Accordingly, the incurred numerical evaluation error can be reinterpreted as a discretization error. Loosely-coupled methods therefore satisfy conservation only in an asymptotic sense, i.e., for vanishing mesh width; this is a basic consistency requirement. Loosely-coupled methods are typically energy-increasing and, hence, numerically unstable; see, e.g., Refs. [53, 55, 56]. To control the imbalance of energy and for reasons of accuracy, the time-step size has to be restricted.

4.2.4 Prediction techniques

To improve the order of the numerical evaluation error incurred by loosely-coupled partitioned methods, prediction techniques are proposed in, e.g., Refs. [53, 55]. Instead of integrating the fluid equations based on the position of the structure boundary in the previous time slab, a prediction can be used for the position of the structure boundary in the current time slab. Such predictions are generally based on an extrapolation of the solution from the previous time slab. Prediction techniques improve the solution accuracy and stability of loosely-coupled methods; see Refs. [53, 55]. We demonstrate this effect by numerical experiments in Section 4.3.

Alternatively, prediction techniques can be used to initialize the subiteration process in strongly-coupled partitioned methods. This can substantially reduce the number of subiterations required for convergence to sufficient accuracy, i.e., to render the numerical evaluation error in the approximation smaller than the interpolation error (discretization error). To elucidate the connection between the order of the extrapolation and the number of required subiterations, let us recall a result from Ref. [69]. To this end, we associate with the subiteration process as specified by steps (S1)-(S4) an operator C that maps one structure interface displacement onto the next according to $(z|_{\Theta})_i \mapsto (z|_{\Theta})_{i+1} = \mathsf{C}(z|_{\Theta})_i$; an elaboration is deferred to Section 5.2.1. On the basis of a linearized model problem, the subiteration-operator derivative is shown to be bounded as $\|\mathsf{C}'\| \leq 2\tau$ (in the appropriate norm), where τ denotes the computational time-step size. This implies that the numerical evaluation error in the approximation is reduced by at least a factor of 2τ per subiteration. The number of required subiterations then depends on the error in the initial approximation obtained by extrapolation. For sufficiently smooth solutions, and using an extrapolation of equal order as the interpolation, the error in the extrapolation is of the same order as the interpolation error. In that case, two subiterations are required for convergence. To elaborate this assertion, we note that with a single subiteration the numerical evaluation error in the approximation is reduced below the interpolation error only for the structure,

but not for the fluid and the interface. To reduce the global evaluation error below the interpolation error for a solution over multiple time intervals, sufficiently accurate initial conditions need to be provided. This requires a reduction of the evaluation error below the interpolation error also for the fluid and the interface. Hence, a second subiteration must be carried out. We remark that this result is based on asymptotic considerations of sufficiently small time steps.

To generalize this result to extrapolations of arbitrary order, suppose that the interpolation error is of order N, and that the solution is initialized by means of an *n*-order extrapolation with $n \leq N$. Then 2 + N - n subitrations are required to reduce the numerical evaluation error below the interpolation error.

4.2.5 Strongly-coupled versus loosely-coupled methods

Strongly-coupled and loosely-coupled partitioned methods are controversially advocated in literature; see, e.g., Refs. [1, 18, 20, 47]. Although an a-priori preference of one method over the other is not justified, a choice can be motivated by aspects of the underlying physics and by efficiency arguments. In this section, we discuss the implications of the underlying physics and, moreover, compare the efficiency of strongly-coupled and loosely-coupled partitioned methods.

To elucidate the implications of the underlying physics for the choice of a solution method, we consider, for exemplification, two problems with 'strong' and 'weak' interaction, respectively. A problem with 'strong' interaction arises, for instance, in hemodynamics, where blood flow interacts with the heart ventricles and arteries; see, e.g., Ref. [48, 63]. As fluid and structure mass are of comparable orders of magnitude, they contribute equally to the dynamics of the system. Moreover, because the fluid can be considered as incompressible, the sound speed is infinite. This implies that perturbations induced at the interface are 'felt' instantaneously throughout the fluid domain and, hence, trigger an immediate 'response' to the structure. The above characterizes hemodynamics as a problem with strong interactions between fluid and structure. For such problems, loosely-coupled partitioned methods generally do not work. Strongly-coupled partitioned methods or monolithic methods are then the only viable option; see, e.g., Refs. [26, 36, 46].

An example of 'weak' interaction is given by problems with wind-induced loads on buildings and bridges. As the structure mass is typically much larger than the fluid mass, the dynamics of the structure is hardly influenced by the fluid. The structure then dominates the dynamics of the fluid-structure system. If the fluid exercises only a minor effect on the structure, loosely-coupled partitioned methods constitute an adequate option. In the limit of vanishing forcing exerted by the fluid on the structure, the interaction reduces to a quasi one-way influence, and the structure can be solved separately. The converse, however, is generally not true. That is, if the structure mass is much smaller than the fluid mass, the effect of the structure on the fluid cannot be neglected, because the structure solution determines the domain on which the fluid is defined. An example of such a problem constitutes a flag which is flapping in the wind.

For problems for which both strongly-coupled and loosely-coupled partitioned methods constitute a viable option, preference of one method over the other can be motivated by a comparison of their stability and computational efficiency. Computational efficiency can be defined as the ratio of accuracy to computational work. Strongly-coupled solution methods have a higher computational cost per time step, but they can maintain energy conservation which renders them unconditionally stable. The computational time-step size is therefore restricted by the desired accuracy only, but not by the stability of the numerical method. In contrast, for loosely-coupled methods, the computational work per time step is comparatively small, but the inherent numerical evaluation error impedes conservation and restricts the time-step size for reasons of stability and accuracy. In particular, the stability of modes with time scales much smaller than one actually wants to resolve can impose a severe restriction on the admissible time-step size. This discussion is similar to the one on the time-step restriction of explicit and implicit time-integration methods. Typically, for loosely-coupled methods, the time step must be much smaller than for strongly-coupled methods. This raises the question whether the larger time-step size of a strongly-coupled method justifies its higher computational cost, i.e., whether multiple subiterations are computationally efficient. Ref. [20] claims that "interfield iteration [subiteration] generally costs more than cutting the time step to attain the same accuracy level". For this reason, loosely-coupled partitioned methods are strongly advocated in literature; see also Refs. [18, 55]. However, we cannot agree to the above statement, because to the best of our knowledge no evidence is provided in literature to support this claim. In fact, our numerical experiments in Section 4.3 refute this claim. Moreover, the computational cost of strongly-coupled methods can be significantly reduced, for instance, by means of the Krylov acceleration presented in Chapter 5.

4.2.6 Improving the efficiency of strongly-coupled methods

In this section, we discuss concepts to improve the computational efficiency of strongly-coupled methods by taking advantage of the fact that the subiteration process is repeated multiple times within a time step. To this end, we note that the solution of the fluid subsystem generally dominates the cost of the subiteration process, because typically the number of unknowns in the fluid exceeds by far those in the structure. This is due to the fact that the length and time scales in the fluid are typically by several orders of magnitude smaller than those in the structure and, hence, the fluid discretization requires a much finer resolution. The efficiency of the subiteration process can be significantly improved by reducing the computational expense for the fluid solution. Below, we elaborate three options.

A first, obvious option is to employ prediction techniques, in order to obtain an improved initial approximation for the iterative solution of both the fluid subsystem and the aggregated fluid-structure system; see Section 4.2.4.

A second option is based on a reduced number of iterations for the fluid solution within a subiteration cycle. We remark that generally good initial guesses for the fluid solution are available. These can be obtained either from the previous subiteration or by extrapolation from the previous time slab. With a good initial guess for the fluid iterative process, quadratic Newton convergence can in principle be obtained. Subiteration, on the other hand, corresponds to a block Gauss-Seidel process which has only a linear convergence rate. Hence, fluid convergence is asymptotically faster than subiteration convergence. Therefore, in principle, a single Newton iteration in the fluid is sufficient for the convergence of the subiteration process. Thus, instead of repeating Newton iterations in the fluid until convergence, only an approximate fluid solution is required. Upon subiteration convergence, it is however important to verify that also the fluid residual separately satisfies a given convergence criterion. The above suggests that the number of required subiterations is not significantly influenced by solving the fluid only approximately. Each subiteration, however, becomes computationally much cheaper, as fewer Newton iterations on the fluid solution are expended. Obviously, this can considerably increase the computational efficiency of subiteration-based stronglycoupled solution methods. Finally, we remark that for loosely-coupled solution methods, which by definition take only one subiteration per time step, a single Newton iteration does generally not yield a sufficiently accurate fluid solution.

A third option to reduce the computational expense of the fluid solution is based on the observation that for finite wave-propagation speeds the perturbation induced by an update of the interface position can be confined to a fluid region in the vicinity of the interface. If the perturbation decays sufficiently fast in the direction perpendicular to the interface, the fluid solution in the subiteration process can in principle be confined to a truncated fluid domain. The solution of the fluid on the entire domain is then required only once to verify convergence. This technique can potentially deliver significant computational savings, but its implementation can be involved. We shall therefore not pursue it any further.

4.3 Numerical experiments

To illustrate the different properties of strongly-coupled and loosely-coupled partitioned solution methods, we conduct numerical experiments on a prototypical fluid-structure interaction problem. In particular, we assess the stability and accuracy of the methods, and we compare their computational efficiency. Finally, we investigate prediction techniques to further improve the efficiency of stronglycoupled methods.

4.3.1 Experimental setup

The numerical experiments consider the piston problem introduced in Section 3.2. To enable a comparison with the results obtained in Chapter 3, we use the same physical parameter settings as specified in Table 3.1 on page 43.

The fluid discretization employs the time-discontinuous Galerkin method with linear-in-time shape functions and is therefore formally second-order time

accurate. For the structural time integration we use the analytic solution from Section 3.5.1.

To investigate the effect of the numerical evaluation error of loosely-coupled and strongly-coupled solution methods separately from the fluid and structure discretization errors, we must ensure that the latter are sufficiently small. Discretization errors in the structure are excluded by means of the analytic solution. Unless stated otherwise, we use for the fluid discretization subcycling with a subcycling factor $\kappa = 8$, i.e., the fluid discretization uses a time step which is eight times smaller than the time step of the structural discretization. Although the difference between the methods becomes apparent also without fluid subcycling for coarse time steps, we require fluid subcycling when investigating the accuracy of the solution methods in order to reduce the effect of fluid dissipation. In general, fluid dissipation damps the solution and, hence, counteracts the amplification caused by loosely-coupled solution methods; cf. Ref. [70]. In our numerical experiments, we use second-order structural prediction in conformity with Ref. [55].

4.3.2 Investigation of stability properties

For the numerical experiments presented in this section, the same time-step size is used for the fluid and structure discretizations. The size of the time step is indicated below the figures. Figs. 4.1-4.3 on the following pages plot the structural displacement versus time in five oscillation periods for a strongly-coupled and a loosely-coupled partitioned method with and without structural prediction. Fig. 4.1 shows that the amplitude computed with the loosely-coupled method grows in time, whereas the amplitude computed with the strongly-coupled method remains constant. This calculation uses a time step of $\tau = 0.01$. Next, we use a structural prediction according to Eq. (3.37), with values of the forcing term at the previous time level. When using a structural predictor for the looselycoupled method, the growth in amplitude is substantially reduced so that the results for the loosely-coupled method and the strongly-coupled method virtually coincide (Fig. 4.2, with the same time-step size as Fig. 4.1). This illustrates that structural prediction reduces the numerical evaluation error incurred by a looselycoupled method. However, if the time-step size is increased, e.g., by a factor of ten (Fig. 4.3), the error in the amplitude computed with the loosely-coupled method emerges again, in spite of the structural prediction. The amplification of the solution obtained with the loosely-coupled method can be attributed to artificial energy production at the interface, which can induce numerical instability of the combined fluid-structure system. In practice, this imposes a restriction on the admissible time-step size. The strongly-coupled method does not become numerically unstable. Indeed, as it maintains exactly the conservation properties at the interface, it is unconditionally stable. In this case, there is no restriction on the admissible time-step size, other than the required accuracy.

We remark that for a single-mode problem such as the one considered, stability and accuracy are closely related. Only for multiple-mode problems stability and accuracy can be distinguished clearly. For loosely-coupled methods, the stability of modes with time scales smaller than one actually wants to resolve can impose a severe restriction on the admissible time-step size. On the other hand, for strongly-coupled methods there is no such stability restriction, but the timestep size is restricted by accuracy considerations only.



Figure 4.1: Structural displacement versus time: strongly-coupled method (-) and loosely-coupled method without structural prediction (--), $\tau = 0.01$.



Figure 4.2: Structural displacement versus time: strongly-coupled method (---) and loosely-coupled method with structural prediction (--), $\tau = 0.01$.


Figure 4.3: Structural displacement versus time: strongly-coupled method (---) and loosely-coupled method with structural prediction (--), $\tau = 0.1$.

4.3.3 Investigation of accuracy properties

This subsection investigates the temporal accuracy of strongly-coupled and looselycoupled methods by a grid refinement study. We use the same settings as in Section 4.3.2. However, the subcycling ratio is set to 8.

Table 4.2 shows the observed order of temporal accuracy (determined as in Section 3.5.3) for the strongly-coupled method as well as for the loosely-coupled method with and without structural prediction. The observed order of temporal accuracy displays the expected asymptotic behaviour, i.e., second-order accuracy for the strongly-coupled method and for the loosely-coupled method with prediction, and first-order accuracy for the loosely-coupled method without prediction.

mesh sequence	loose coupling	loose coupling	strong coupling	
with time-step sizes	without prediction	with prediction		
$2^{-3}, 2^{-4}, 2^{-5}$	1.4667	1.9848	5.0188	
$2^{-4}, 2^{-5}, 2^{-6}$	1.4565	1.9891	2.6159	
$2^{-5}, 2^{-6}, 2^{-7}$	1.1958	2.0018	2.0457	

Table 4.2: Observed order of temporal accuracy according to Eq. (3.40): loosely-coupled and strongly-coupled partitioned methods.

The strongly-coupled method and the loosely-coupled method with prediction have formally the same order of accuracy; the nature and magnitude of the dominant errors, however, are inherently different. To illustrate this difference, we plot in Fig. 4.4 the L^2 norm of the error in the computed structural displacement versus the size of the structural time step for $\tau = 2^{-3}, 2^{-4}, \ldots, 2^{-7}$, using the same reference solution as in Section 3.5.2. As the strongly-coupled method and the loosely-coupled method differ only in the number of subiterations, we can infer from Fig. 4.4 that the numerical evaluation error in the solution of the loosely-coupled method is by several orders of magnitude larger than the fluid and structure discretization errors. That is, for a loosely-coupled method the numerical evaluation error can dominate other sources of error. From Fig. 4.4 it is clear that the loosely-coupled method requires smaller time steps than the strongly-coupled method for a specified error tolerance. Conversely, given a certain level of accuracy, a strongly-coupled method can allow for larger time steps than a loosely-coupled method.

As the grid refinement study uses the same settings and reference solution as in Section 3.5.2, the accuracy of the loosely-coupled method can be compared also to the accuracy of the non-conservative strongly-coupled methods B and C; see Fig. 3.3 on page 46. Although the non-conservative methods B and C are more accurate than the loosely-coupled method with prediction, the gain in accuracy does not seem to justify the computational cost associated with a strongly-coupled method. However, comparing the accuracy of a loosely-coupled and a conservative strongly-coupled method, it is clear from Fig. 4.4 that the latter is much more accurate. In fact, the conservative strongly-coupled method is by more than three orders of magnitude more accurate than the loosely-coupled method. Hence, strongly-coupled methods are much more accurate than loosely-coupled methods, provided that they maintain conservation at the interface.

4.3.4 Comparison of efficiency

In the following, we compare loosely-coupled and conservative strongly-coupled methods in terms of efficiency, which we define as the ratio of accuracy to computational work. Let us employ as a measure of accuracy the reciprocal of the error in the numerical solution and as a measure of computational work the number of fluid-structure iterations. In Fig. 4.5, the error is plotted versus the number of fluid-structure iterations on a log-log scale for the conservative strongly-coupled method and for the loosely-coupled method with structural prediction. The smaller the time step, the more iterations are required for one cycle and, correspondingly, the computational work increases. The error in a loosely-coupled method is generally larger than in a strongly-coupled method. For a given number of fluid-structure iterations the strongly-coupled method. The strongly-coupled method is at least two orders of magnitude more accurate than the loosely-coupled method. The strongly-coupled method is therefore computationally more efficient.



Figure 4.4: Error in the structure displacement versus structural time-step size: strongly-coupled method (-) and loosely-coupled method with structural prediction (-).



Figure 4.5: Error in the structure displacement versus number of fluid-structure iterations: strongly-coupled method (-) and loosely-coupled method with structural prediction (-).

4.3.5 Improving the efficiency of strongly-coupled methods

In this section, we investigate prediction techniques to reduce the required number of subiterations in strongly-coupled solution methods. The subiteration process requires an initial estimate of the structure solution which, for instance, can be obtained by prediction techniques. Here, the structural prediction is done according to Eq. (3.37), with values of the forcing term at the previous time level. The effect of such structural prediction on the number of required subiterations is displayed in Fig. 4.6, which plots the number of subiterations per time step versus time. In particular, we compare the number of subiterations for a strongly-coupled method with and without structural prediction. As the area under the curves in Fig. 4.6 corresponds to the computational work, it is clear that prediction techniques can reduce the computational cost of strongly-coupled solution methods. When increasing the time-step size, the required number of subiterations per time step increases moderately. However, the positive effect of prediction techniques on the computational cost can still be observed.

Moreover, Fig. 4.6 indicates that the computational cost of the stronglycoupled method is still three times the cost of a loosely-coupled method, which requires only a single subiteration per time step. On the other hand, the stronglycoupled method is by more than a factor of 1000 as accurate, see Fig. 4.4 on the preceding page. Preference of one method over the other depends on whether the increased accuracy also justifies the greater computational cost. For the looselycoupled method to attain the same accuracy as the considered strongly-coupled method, we need to reduce the time-step size by a factor of approximately 30 in the second-order accurate loosely-coupled method. However, the computational cost of the loosely-coupled method then increases correspondingly by a factor of 30, and it becomes 10 times as expensive as the strongly-coupled method for a comparable level of accuracy. This refutes a common belief that multiple fluidstructure iterations generally cost more than reducing the time step of a looselycoupled method to attain the same level of accuracy; see, e.g., Ref. [20]. We expect the superior efficiency of a strongly-coupled method to be even more pronounced for multiple-mode problems, as discussed in Section 4.3.2.

4.4 Concluding remarks

In this chapter, we considered monolithic methods, and strongly-coupled and loosely-coupled partitioned methods for the solution of the fluid-structure system. We elaborated that the monolithic solution by Newton's method incurs several severe disadvantages, viz., the prohibitively expensive evaluation of shape derivatives, ill-conditioning of the system Jacobian matrix and the loss of software modularity due to the simultaneous solution of fluid and structure. We explained that the complications in Newton's method can be effectively circumvented by partitioning, and we distinguished strongly-coupled and loosely-coupled partitioned methods. Strongly-coupled methods incur a greater computational cost per time step than loosely-coupled methods. However, strongly-coupled methods can maintain conservation at the fluid-structure interface, which renders them unconditionally stable. In contrast, loosely-coupled methods are typically energy increasing



Figure 4.6: Number of subiterations per time step versus time for one oscillation period: strongly-coupled method without structural prediction (--) and with structural prediction (--), $\tau = 2^{-6}$.

and, hence, numerically unstable. To control the numerical instability and for reasons of accuracy, the time-step size has to be restricted. However, this restriction can be alleviated by means of prediction techniques which mitigate the numerical instability and improve the accuracy. For strongly-coupled methods, on the other hand, the admissible time-step size is determined by accuracy considerations only.

To compare strongly-coupled and loosely-coupled methods in terms of stability, accuracy, computational cost and efficiency, we conducted numerical experiments on the piston problem. Our results illustrated the restrictions on the time step of loosely-coupled methods and the positive effect of prediction techniques. Moreover, we found that conservative strongly-coupled methods are much more accurate than loosely-coupled methods and, hence, they can afford much larger time steps for the same level of accuracy. We demonstrated that conservative strongly-coupled methods achieve a greater accuracy with fewer iterations than loosely-coupled methods and are therefore more efficient. Our numerical results also indicate that for strongly-coupled methods which are not conservative the improvement in accuracy over loosely-coupled methods does not seem to justify the additional computational cost associated with multiple subiterations. Indeed, void of the conservation properties, strongly-coupled methods appear to lose their essential advantage over loosely-coupled methods. Our results therefore warrant a clear preference for conservative strongly-coupled methods.

To improve the efficiency of strongly-coupled methods, we investigated prediction techniques which were shown to reduce the number of required subiterations. Moreover, we proposed two concepts to reduce the computational cost associated with the fluid solution in the subiteration process. The first concept propounds an approximate fluid solution within a subiteration, the second concept restricts the fluid solution to the vicinity of the fluid-structure interface. In either case, upon convergence of the subiteration process, fluid convergence needs to be verified separately.

Finally, let us briefly discuss to which extent our findings generalize. For a single-mode problem such as the one considered, stability and accuracy cannot be distinguished clearly. This is only possible for multiple-mode problems. However, we expect that for multiple-mode problems the superiority of conservative strongly-coupled methods over loosely-coupled methods is even more pronounced. This conjecture is based on the fact that for loosely-coupled methods the stability of modes with time scales smaller than one actually wants to resolve can impose a severe restriction on the admissible time-step size. For strongly-coupled methods, on the other hand, there is no such stability restriction, but the time-step size is restricted by accuracy considerations only.

Chapter 5 Interface-GMRES(R) acceleration of subiteration

5.1 Introduction

The numerical solution of fluid-structure interaction problems commonly employs subiteration, i.e., fluid and structure equations are solved alternately subject to complementary partitions of the interface conditions; see Section 4.2.2 and also Refs. [1, 36, 47], for example. This process essentially constitutes a block Gauss-Seidel iteration, which is repeated until convergence. Subiteration sidesteps a simultaneous treatment of the coupled fluid-structure equations and the concomitant difficulties. Although subiteration is a good solver for many problems, it suffers from two essential drawbacks: Firstly, subiteration converges only slowly or even diverges for problems with large computational time steps or large fluid-tostructure mass ratios. Subiteration is only conditionally stable, but even despite formal stability transient divergence can precede asymptotic convergence. These convergence difficulties can be attributed to the nonnormality of the subiteration operator; see Ref. [69]. Such non-monotonous convergence behaviour can even lead to failure of the solution method despite formal stability. Secondly, subiteration is generally employed in a sequential time-integration process and, hence, solves a sequence of similar problems. However, the method cannot exploit this property and reuse generated information, for instance, for preconditioning purposes. Therefore, subiteration is to be considered inefficient. Our objective is to overcome these drawbacks by combining subiteration with GMRES acceleration.

In this chapter, we show that subiteration can be conceived as a mapping of one interface displacement onto the next. Accordingly, the method can be characterized by recursion of a nonlinear operator on the interface displacement. The subiteration solution method therefore constitutes a fixed-point iteration. The method that we propose instead solves this nonlinear fixed-point problem by means of a *Newton-Krylov method* [10]. Such Newton-Krylov methods solve the nonlinear problem by means of an (inexact) Newton method, in which the linear systems arising in each Newton step are solved by means of a Krylov subspace method such as GMRES [59]. Newton-Krylov methods have become popular for many problems, for example in fluid dynamics, combustion and plasma physics; see the review article [33]. However, their potential for fluid-structure-interaction problems has received only sparse attention; see, e.g., Refs. [26, 38, 58]. Moreover, these approaches apply GMRES to the aggregated variables or to the Schur complement associated with the structure variables. Instead, we propose to combine the Newton-Krylov method with subiteration, which allows us to confine the Newton-Krylov method to the interface degrees-of-freedom. Accordingly, we refer to the method as *Interface-Newton-Krylov method*, or simply as *Interface-GMRES* [43,71]. It is precisely the restriction of the GMRES acceleration to the interface degrees-of-freedom which renders the storage requirements for the Krylov space and the computational cost incurred by the acceleration itself negligible.

The Newton-Krylov method consists of nested iterations and involves frequent solution of the linear system by means of the Krylov method. This offers the possibility of reusing the Krylov space in subsequent Newton iterations and time steps, which we refer to as *Interface-GMRESR*. Such reuse can yield substantial computational savings. Moreover, the GMRES acceleration enables underrelaxation without hampering convergence. Such underrelaxation can be mandatory if the subiteration method is unstable, or to mitigate nonnormality-induced divergence.

In this chapter, we treat the algorithmic description and theoretical background of the proposed numerical method in a continuum setting. We will complement this description by an analysis of the linear-algebra aspects and erroramplification properties of the method in Chapter 6.

We conduct numerical experiments on a prototypical fluid-structure-interaction problem, viz., the one-dimensional piston problem from Chapter 3, to assess the viability and potential of the Interface-GMRES method with and without the reuse option. In particular, we investigate the convergence behaviour of the method, test its robustness and assess its computational cost. A comparison to standard subiteration reveals that the Interface-GMRES(R) method is much more robust than subiteration and that it converges even if subiteration itself diverges. Moreover, the Interface-GMRES(R) method is also much more efficient than subiteration. Reuse of Krylov vectors typically yields considerable computational savings and makes the difference to subiteration even more pronounced.

Finally, we remark that the Interface-GMRES(R) acceleration preserves the modularity of the underlying subiteration method, i.e., it maintains the segregated treatment of fluid and structure; see, e.g., Ref. [20]. This renders the implementation of the acceleration method in codes that already use subiteration as a solver straightforward.

The contents of this chapter are organized as follows : Section 5.2 identifies subiteration as a fixed-point iteration. Next, this section presents the Interface-GMRES acceleration of subiteration and the optional reuse of Krylov vectors. Section 5.3 provides numerical experiments and results. Section 5.4 contains concluding remarks.

5.2 Interface-GMRES(R) acceleration

In Section 5.2.1, we identify the subiteration process as a fixed-point iteration on the interface degrees-of-freedom. Section 5.2.2 presents the GMRES acceleration for the interface fixed-point problem. Section 5.2.3 addresses the reuse of Krylov vectors.

5.2.1 Subiteration: Interface fixed-point iteration

In this section, we reconsider the subiteration process introduced in Section 4.2.2. To facilitate the ensuing exposition of the Interface-GMRES acceleration, we construe subiteration as an operator on the structure displacement at the interface, $z|_{\Theta}$. For conciseness in notation, we shall denote the interface displacement simply by z. The subiteration process can then be conceived as a mapping from one structural interface displacement to the next, i.e.,

$$\mathsf{C}: z_j \mapsto z_{j+1} = \mathsf{C} z_j, \tag{5.1}$$

where C denotes the nonlinear operator induced by subiteration; see also Refs. [69, 71]. The subiteration operator C is not explicitly available; however, its action on z_j can be evaluated by sweeping through the subiteration steps (S1)–(S4) specified in Section 4.2.2. In accordance with the mapping defined by Eq. (5.1), the subiteration process can be characterized by recursion of the nonlinear operator C on the interface displacement z and can be conceived as a fixed-point iteration. The fixed point

$$\bar{z}: \mathsf{C}\bar{z} = \bar{z} \tag{5.2}$$

corresponds to the interface component of the structure solution of the aggregated equations (2.27). Upon providing the subiteration process with the solution \bar{z} , the subiteration steps (S1)–(S4) yield the solution of the aggregated equations in a single iteration under the condition that the equations in each step are solved exactly.

Convergence of the subiteration process is contingent on the spectral radius of the subiteration-operator derivative. We defer a detailed analysis of the convergence behaviour of the subiteration method to Section 6.3.3. Here, we shall rather use subiteration as a building-block for the Interface-GMRES method.

5.2.2 Interface-GMRES acceleration of subiteration

Instead of solving the problem (5.2) by fixed-point iteration, we propose to solve it by a *hybrid Newton-Krylov method* [10]. Each Newton step requires then the solution of a linear problem for which we employ a Krylov subspace method, here GMRES [59]. Because the problem (5.2) involves only the interface degreesof-freedom, we refer to the method also as *Interface-Newton-Krylov method*, or simply as *Interface-GMRES*.

The nonlinear problem can be formulated as

$$\mathsf{R}z = 0, \tag{5.3}$$

with $\mathsf{R} := \mathsf{C} - \mathsf{I}$ the residual operator defined in conformity with Eq. (5.2). Clearly, z being a solution of Eq. (5.3) is identical to z being a fixed point of Eq. (5.2). Correspondingly, the residual of an iterate z_i is

$$r_i := \mathsf{R}z_i = (\mathsf{C} - \mathsf{I})z_i = z_{i+1} - z_i.$$
(5.4)

For a given initial guess z_0 , Newton's method generates a sequence of approximate solutions, z_i , according to

$$z_{i+1} = z_i + z'_i = z_i - [\mathsf{R}'(z_i)]^{-1}\mathsf{R}z_i,$$
(5.5)

where $\mathsf{R}'(z_i) = \partial \mathsf{R}/\partial z(z_i)$ denotes the derivative of R at z_i and z'_i a perturbation around the linearization state z_i . If the initial estimate is sufficiently close to the actual solution z of Eq. (5.3), then z_i converges to z as $i \to \infty$. Each Newton step requires the solution of a linear problem of the form ¹

$$\mathsf{R}z_0 + \mathsf{R}'(z_0)z_0' = 0, \tag{5.6}$$

in which the derivative R' is not known explicitly. However, if a Krylov method is used for the solution of problem (5.6), R' is only required in the form of a matrix-vector product. We seek approximations to z'_0 by making the following ansatz

$$z'_{0} \in \mathcal{K}^{m} := \operatorname{span}\{z_{j} - z_{0}\}_{j=1}^{j=m},$$
(5.7)

where \mathcal{K}^m is the Krylov space of dimension m associated with the linear problem (5.6). Substituting Eq. (5.7) into Eq. (5.6) and approximating the resulting matrix-vector product by finite-differences we obtain

$$R(z_0) + R'(z_0) \sum_{j=1}^{j=m} \alpha_j (z_j - z_0) = R(z_0) + \sum_{j=1}^{j=m} \alpha_j R'(z_0) (z_j - z_0)$$
$$= r_0 + \sum_{j=1}^{j=m} \alpha_j (r_j - r_0) + O(\|\sum_{j=1}^{j=m} \alpha_j (z_j - z_0)\|^2) = 0 \quad (5.8)$$

with the coefficients α_j still to be determined, and span $\{r_j - r_0\}_{j=1}^{j=m}$ is the residualsensitivity space corresponding to span $\{z_j - z_0\}_{j=1}^{j=m}$. The finite-difference approximation which underlies the second identity in Eq. (5.8) requires an additional

¹For transparency, in Eq. (5.6) the linearization state is redefined at each Newton step, i.e., $z_0 \leftarrow z_0 + z_0'$.

residual evaluation, i.e., one subiteration for each search direction. Because the subiteration operator is nonlinear, this finite-difference approximation introduces another linearization error in addition to the one incurred by Newton's method.

To determine the coefficients α_j required for the redefinition of the linearization state $z_0 \leftarrow z_0 + \sum_{j=1}^{j=m} \alpha_j (z_j - z_0)$, we solve Eq. (5.8) in a least-squares sense by minimization in the L^2 norm:

$$\bar{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} \|r_0 + \sum_{j=1}^{j=m} \alpha_j (r_j - r_0)\|, \qquad \xi := \|r_0 + \sum_{j=1}^{j=m} \bar{\alpha}_j (r_j - r_0)\|, \qquad (5.9)$$

where ξ denotes the norm of the residual of the linear problem (5.6). The latter constitutes an estimate for the norm of the residual of the nonlinear problem (5.3). Upon convergence of the Newton process, this estimate becomes increasingly sharp, because the linearization errors vanish. As the least-squares problem (5.9) is confined to the interface degrees-of-freedom, the computational cost involved is small in comparison to the cost incurred by a Newton or a GMRES step, which both require one subiteration; cf. Eqs. (5.5) and (5.8), respectively.

The Krylov space \mathcal{K}^m coincides with the span of $\{\zeta_j - z_0\}_{j=1}^{j=m}$, where ζ_j is the *j*-th subiteration iterate. As subiteration uses only the last iterate of the sequence, ζ_m , and on account of the minimal-residual property of GMRES, it follows that the residuals generated by subiteration form an upper bound for the GMRES residuals. In addition, the minimal-residual property implies that, in contrast to the norm of the subiteration residuals, the norm of the residuals of the GMRES iterates necessarily form a non-increasing sequence. However, this implies faster Newton-Krylov convergence only for problems which are sufficiently linear; for nonlinear problems this does no longer hold due to the linearization error.

Provided with an initial approximation of the structure position $z_0(t)$, Algorithm 5.2.1 on the next page summarizes the Newton-Krylov method for the solution of the nonlinear problem (5.3). To improve the robustness of the method, we generate an orthonormal basis of the Krylov space by employing Gram-Schmidt orthonormalization; see, e.g., Ref. [22]. The Gram-Schmidt orthonormalization is implemented on lines 7–10 of Algorithm 5.2.1. Moreover, we scale the Krylov vectors by an appropriate constant ν (line 10) which facilitates the subiteration process (line 12) required for the residual evaluation. This scaling can be conceived as a form of underrelaxation, because it determines the norm of the update, and it allows to combine GMRES with subiteration even if subiteration is formally unstable. In contrast, using underrelaxation in connection with subiteration separately typically renders convergence excessively slow. We remark that, on linear approximation, the Krylov space is not changed by orthonormalization and underrelaxation. A subiteration is required by each GMRES step (cf. Eq. (5.8) and line 12) and each Newton step (cf. Eq. (5.5) and line 18). The fluid solution can be extracted from the subiteration process on line 1 or 18.

The convergence tolerances for the nonlinear problem (5.3) and the linear problem (5.6) are denoted by ϵ_0 and ϵ_1 , respectively. Clearly, the nonlinear prob-

lem cannot be solved more accurately than the linear problem. On the other hand, if the linearization error is large, then a strict reduction of the linear residual does not yield a corresponding reduction of the nonlinear residual. We therefore consider a relative tolerance for the linear problem with respect to the norm of the nonlinear residual in the current Newton step i, in particular

$$\epsilon_1 = \kappa \|r_i\|,\tag{5.10}$$

where $\kappa < 1$ is a scalar that determines how accurately the linear problem is to be solved. The relative tolerance serves to prevent a waste of computational work incurred by solving the linear problem excessively accurately.

1:
$$i = 0$$
; $z_1 = Cz_0$; $r_0 = z_1 - z_0$
2: while $||r_i|| > \epsilon_0$ do
3: $j = 0$; $\xi = ||r_i||$
4: while $\xi > \epsilon_1$ do
5: $j = j + 1$
6: $z'_j = z_j - z_0$
7: for $k = 1, \dots, j - 1$ do
8: $z'_j = z'_j - z'_k(z'_j \cdot z'_k)/||z'_k||^2$
9: end for
10: $z'_j = \nu z'_j/||z'_j||$
11: $z_j = z_0 + z'_j$
12: $z_{j+1} = Cz_j$
13: $r'_j = (z_{j+1} - z_j) - r_i$
14: $\bar{\alpha} = \arg \min ||r_i + \sum_{k=1}^{k=j} \alpha_k r'_k||$
15: $\xi = ||r_i + \sum_{k=1}^{k=j} \bar{\alpha}_k r'_k||$
16: end while
17: $z_0 = z_0 + \sum_{k=1}^{k=j} \bar{\alpha}_k z'_k$
18: $i = i + 1$; $z_1 = Cz_0$; $r_i = z_1 - z_0$
19: end while

Algorithm 5.2.1: The Interface-Newton-Krylov method.

The above presentation introduces the Interface-GMRES method as an acceleration of the interface fixed-point iteration that is induced by subiteration. Alternatively, the subiteration method can be construed as a preconditioner for the aggregated system; cf. Section 6.3.4. The Interface-GMRES method then solves this preconditioned aggregated system by means of a Krylov method. As the subiteration preconditioner condenses errors into a subspace which can be associated with the interface variables (cf. Section 6.3.2), the Krylov vectors can be represented in the interface approximation space.

In contrast to approaches which apply GMRES to the aggregated variables or to the structure variables, see Refs. [26, 38, 58], the proposed Interface-GMRES method is confined to the interface variables. Therefore, the storage requirements for the Krylov space and the computational expense for the solution of the leastsquares problem are much lower. Moreover, we remark that the acceleration space for the Interface-GMRES method need not necessarily be the approximation space associated with the structure interface displacement, but the approximation space associated with any interface quantity can serve as an acceleration space. Different choices for the acceleration space are investigated in Ref. [71].

We note that, because the Interface-GMRES method allows for a description in a continuum setting, its convergence behaviour can be inferred to be asymptotically independent of the underlying discretization; see also Ref. [71]. This implies that for all sufficiently fine discretizations, the convergence behaviour of the Interface-GMRES method is identical.

Finally, to place the proposed method into context, let us briefly address its relation to so-called *iterative substructuring schemes* [36], which accelerate convergence by relaxed updates of the interface position. The relaxation parameter can be determined, for instance, by means of the steepest-descent method or Aitken's method; see Refs. [45, 46]. Such iterative substructuring schemes essentially can be viewed as a special instance of the Interface-GMRES method, in which the update of the interface position is computed based on a single search vector only.

5.2.3 Reuse of Krylov vectors: Interface-GMRESR

As each Newton step invokes the solution of a linear system by a Krylov method. the Newton-Krylov method lends itself naturally to reuse of Krylov vectors in subsequent Newton steps. In the context of the GMRES acceleration on the interface displacement, we shall refer to such reuse also as Interface-GMRESR. Reuse of Krylov vectors requires only minor modifications to Algorithm 5.2.1. These modifications are implemented by replacing the corresponding lines by those in Algorithm 5.2.2 on the following page. Essentially, reuse requires that the counter i is not reset in each Newton step. The inner loop then augments instead of overwrites the available search and residual-sensitivity space. Moreover, an additional residual estimate ξ corresponding to the reduction of the updated nonlinear residual in the available space can be added. Depending on this initial residual estimate, the search and residual-sensitivity space are further augmented or a Newton update is carried out. We remark that, once a single vector is reused, the search space does formally no longer constitute a Krylov space, which implies that the search directions do not necessarily constitute 'preferential' search directions. Nevertheless, typically, much fewer Krylov vectors need to be added to the reused space than are generated for a reconstructed Krylov space, which can result in considerable computational savings.

In addition to the reuse option within a single time step considered above, reuse is also possible over subsequent time steps. In the latter case, the available search and residual-sensitivity space are carried over from one time interval to the next. There is, however, a difference between the two reuse options. Within

1:
$$i = 0; \ j = 0; \ z_1 = Cz_0; \ r_0 = z_1 - z_0$$

3a: $\bar{\alpha} = \arg \min \|r_i + \sum_{k=1}^{k=j} \alpha_k r'_k\|$
3b: $\xi = \|r_i + \sum_{k=1}^{k=j} \bar{\alpha}_k r'_k\|$
3c: $z_{j+1} = z_1$

Algorithm 5.2.2: Modification of Algorithm 5.2.1 to enable reuse of Krylov vectors.

a single time step the nonlinear subiteration operator does not change, but only its linearization state does. In contrast, the subiteration operator does change between time steps on account of differences in initial conditions. As long as the operators in subsequent time steps are sufficiently similar, reuse will be beneficial.

Reuse can substantially enhance the efficiency of the method; however, it comes at the expense of robustness and therefore has to be exercised with some caution. The benefit and the viability of the reuse option are contingent on the similarity between the reused space and the reconstructed space. Failure of the reuse option is in principle possible, however, it appears to be rare; see Ref. [42, 71]. We will assess the potential and limitations of the reuse option based on numerical experiments in Section 5.3 and, moreover, by means of an error-amplification analysis in Chapter 6.

To place the reuse option into context, let us mention two alternative approaches for reusing computational information, viz., so-called *search space injection* [72] and *nested preconditioning* [12]. In Section 6.4.2, we establish commonalities and differences between these approaches and our methodology of reusing Krylov vectors.

5.3 Numerical experiments

To assess the properties and the potential of the proposed Interface-Newton-Krylov method, we conduct numerical experiments on the piston model problem from Section 3.2. In Section 5.3.1, we investigate the convergence behaviour of the Newton-Krylov method and compare it to the convergence behaviour of the subiteration method. In particular, we examine the effects of nonlinearity and computational time-step size on convergence. For these investigations, consideration of a single time step suffices. In Section 5.3.2, we investigate the reuse of Krylov vectors in subsequent time steps and demonstrate its potential for enhancing the efficiency of the method. We compare the Newton-Krylov method and subiteration in terms of computational expense and robustness. We consider multiple time steps and various fluid-to-structure mass ratios.

5.3.1 Convergence of Interface-GMRES(R)

First, we investigate the convergence of the basic Newton-Krylov method (Algorithm 5.2.1 on page 74) by examining the convergence of the linear and nonlinear residuals and the use of a relative tolerance for the linear residual. Moreover, we investigate the Newton-Krylov method with reuse of Krylov vectors in subsequent Newton iterations (Algorithm 5.2.2). In addition, we examine the effect of non-linearity on the convergence of the method. Finally, we compare the convergence behaviour of the Newton-Krylov method to standard subiteration for different computational time steps and fluid-to-structure mass ratios.

Experimental setup

We consider the piston problem and note that if the initial conditions are specified as $z(0) = \dot{z}(0) = 0$ and $\mathbf{u}(x,0) := (\rho, \rho v, E)(x,0) = (\rho^0, 0, E^0) =: \mathbf{u}^0$ with ρ^0 and E^0 appropriate constants and $p^0 = p(\mathbf{u}^0)$, then the obvious solution to the aggregated variational statement (3.10) is

$$q := (\mathbf{u}, \alpha, z, \pi) = (\mathbf{u}^0, \alpha^0, 0, p^0) =: q^0.$$
(5.11)

To examine the convergence behaviour of the considered solution methods we instead provide an initial approximation of the structure displacement conforming to

$$z_0(t) = z^0 + \chi(\exp(-100t^2) - 1) = \chi(\exp(-100t^2) - 1), \qquad t \in [0, \tau], \quad (5.12)$$

where τ denotes the time-step size and the parameter χ determines the error in the initial approximation. The initial approximation of the structure displacement (5.12) constitutes a curve in space/time, in accordance with the space/time finite-element discretization of the piston problem. We consider $\chi = 10^{-2}$ and $\chi = 10^{-6}$. For $\chi = 10^{-2}$ the system is pronouncedly nonlinear, whereas for $\chi = 10^{-6}$ the system is only weakly nonlinear. Starting from the initial approximation (5.12), the solution methods generate a sequence of residuals and, in principle, upon convergence retrieve the uniform solution (5.11). We monitor this sequence of residuals within a single time step.

The fluid-structure system is discretized by means of space/time finite elements. The adopted discretization is essentially identical to that in Ref. [70]. For completeness, we briefly summarize its setup. The space/time fluid domain is covered with a tessellation of quadrilateral elements. The number of elements in spatial direction is denoted by $N_{\mathbb{U}}^t$, the number of elements in temporal direction per unit time by $N_{\mathbb{U}}^t$. The structure mesh consists of $N_{\mathbb{Z}}^t$ elements per unit time. The fluid equations are discretized by means of a discontinuous Galerkin method with the approximation space consisting of piecewise tensor products of polynomials of order $P_{\mathbb{U}}$ (space,time) in conformity with Section 3.4.1. The approximation spaces admit discontinuities across element boundaries. The elements in the fluid are connected by the modified Osher scheme and weakly enforced initial conditions. The structure equation is discretized by means of a continuous Galerkin method. The approximation space of the structure consists of piecewise polynomials of order $P_{\mathbb{Z}}$. The connection between the elements in the structure is provided by strongly enforced initial conditions through Lagrange multipliers. The interface approximation spaces associated with kinematic and dynamic conditions comprise $N_{\mathbb{A},\mathbb{P}}$ elements per unit time, and consist of piecewise polynomials of order $P_{\mathbb{A}}$ and $P_{\mathbb{P}}$, respectively.

The system and discretization parameters are given in Table 5.1 and 5.2, respectively, where $c^0 := \sqrt{\gamma(\gamma - 1)E^0/\rho^0}$ denotes the speed of sound. Throughout, we shall set $N_{\mathbb{A},\mathbb{P}} = N_{\mathbb{Z}}^t$. In the following investigation, we use $\tau = 1$ and $N_{\mathbb{U}}^t = N_{\mathbb{Z}}^t = 24$. We set the convergence tolerance to $\epsilon_0 = 10^{-6} ||r_0||$, i.e., we require a reduction of the initial residual by six orders of magnitude. In addition, we specify for the Newton-Krylov method the tolerance for the GMRES iteration according to $\epsilon_1 = 10^{-7} ||r_0||$, unless stated otherwise, and an underrelaxation parameter of $\nu = 0.1$ throughout.

Table 5.1: System parameters (* indicates a variable parameter).

z^0	\dot{z}^0	α^0	$ ho^0$	c^0	K	M	au
0	0	1	6	0.5	1	1	*

Table 5.2: Discretization parameters (* indicates a variable parameter).

$N^x_{\mathbb{U}}$	$N^t_{\mathbb{U}}$	$N_{\mathbb{A}}$	$N_{\mathbb{Z}}^t$	$N_{\mathbb{P}}$	$P_{\mathbb{U}}$	$P_{\mathbb{A}}$	$P_{\mathbb{Z}}$	$P_{\mathbb{P}}$
24	*	*	*	*	(3,3)	5	5	4

Numerical results

To investigate the convergence behaviour of the Newton-Krylov method, we plot in Fig. 5.1 the residual reduction versus the iteration counter for the different variants of the Newton-Krylov method considered in the sequel and for subiteration as a reference. We begin by investigating the convergence properties of the basic Newton-Krylov method (Algorithm 5.2.1 on page 74). The solution process alternates between GMRES and Newton steps. This is reflected in Fig. 5.1 by a sequence of residual estimates followed by the true residual. Initially, there is a considerable discrepancy between the estimate and the true residual. This discrepancy is induced by the linearization errors incurred by the Newton process and the finite-difference approximation. If the linearization errors dominate the solution error in the linear problem, the difference between these errors manifests itself as a jump in the convergence curve. As the linearization error is initially larger for $\chi = 10^{-2}$ than for $\chi = 10^{-6}$, the residual reduction in the first Newton step is much smaller for $\chi = 10^{-2}$ than for $\chi = 10^{-6}$, although the linear systems are solved to the same accuracy. Hence, for $\chi = 10^{-2}$ the convergence curve exhibits larger jumps and requires more iterations for convergence; see Fig. 5.1. As upon convergence the linearization error vanishes, the estimate becomes increasingly sharp and the jumps decrease.



Figure 5.1: Residual reduction in the L^2 norm versus the iteration counter for the basic Newton-Krylov method (—), with a relative tolerance on the linear residual (– –), with reuse (…), with a relative tolerance and reuse (–) and subiteration (—) for $\chi = 10^{-6}$ (*left*) and $\chi = 10^{-2}$ (*right*) with $\tau = 1$ and $N_{\mathbb{U}}^t = N_{\mathbb{Z}}^t = 24$. The linear and nonlinear residuals of the Newton-Krylov method are indicated by \circ and *, respectively, and the residuals of subiteration by \triangle .

Fig. 5.1 also displays the residual reduction obtained by subiteration. The difference between the subiteration convergence curves for $\chi = 10^{-6}$ and $\chi = 10^{-2}$ is minute. We infer that subiteration is not significantly influenced by the non-linearity of the problem. According to Section 5.2.2, the sequence of residuals generated by subiteration constitutes an upper bound for the residuals obtained by GMRES. However, faster Newton-Krylov convergence is implied only for sufficiently linear problems; for nonlinear problems this does no longer hold due to the linearization error. Indeed, Fig. 5.1 shows that for $\chi = 10^{-6}$ the Newton-Krylov method converges faster than the subiterations. In the following, we investigate modifications of the basic Newton-Krylov method which enable faster convergence also for nonlinear problems.

As an improvement of the basic Newton-Krylov method, we consider a relative tolerance for the solution accuracy of the linear problem with respect to the current nonlinear residual (cf. Eq. (5.10)) instead of solving the linear problem to the same absolute tolerance at each Newton step. The relative tolerance serves to prevent a waste in computational work incurred by overly accurate solution of the linear problems: if the linearization error is large, then a strict reduction of the linear residual does not yield an according reduction of the nonlinear residual anyway. Here, we use a relative tolerance of $\epsilon_1 = 10^{-2} ||r_i||$, i.e., the linear problem in Newton step i is considered to be solved sufficiently accurately if the estimate is reduced by two orders of magnitude with respect to the nonlinear residual r_i . This implies that, typically, much fewer GMRES iterations are carried out than with the basic Newton-Krylov method. Fig. 5.1 exhibits that for $\chi = 10^{-2}$ still the same residual reduction is obtained after the first Newton update. This indicates that the nonlinear residual stagnates at some point. Such stagnation occurs if the linearization error dominates the solution error in the linear problem. In contrast, for $\chi = 10^{-6}$ the residual reduction after the first Newton update is less than with the basic Newton-Krylov method, because the solution error in the linear system still dominates the linearization error. As a result, for $\chi = 10^{-6}$ the Newton-Krylov method with relative tolerance requires more iterations than the basic Newton-Krylov method. We conclude that a stringent relative tolerance is suitable for weakly nonlinear problems, whereas a less stringent relative tolerance is more appropriate for nonlinear problems. The appropriateness of the specified relative tolerance can be deduced from the magnitude of the jumps. If the discrepancy between residual and estimate is minute, the linearization error is small and the update of the linearization state by Newton's method was not required; additional GMRES steps could have further reduced the solution error in the linear system. Conversely, if the jumps are large, the linear system was solved more accurately than what could be accomplished by the Newton update. The choice of the relative tolerance therefore has implications for the computational cost, as a Newton step and a GMRES step are both at the expense of one subiteration.

To investigate the reuse of Krylov vectors for the solution of the linear systems in subsequent Newton iterations (Algorithm 5.2.2 on page 76), we employ reuse in combination with the aforespecified relative tolerance and examine the convergence for the linear and nonlinear setting of the problem. From Fig. 5.1 right it can be seen that, especially for the nonlinear setting ($\chi = 10^{-2}$), the Newton-Krylov method with reuse converges considerably faster than the method without reuse. Fig. 5.1 right also conveys that, in particular for nonlinear problems, the setting of the linear-system tolerance ϵ_1 can have implications for the effectiveness of the reuse option. Comparing the curves with a relative tolerance $(\epsilon_1 = 10^{-2} ||r_i||)$ and with an absolute tolerance $(\epsilon_1 = 10^{-7} ||r_0||)$, it appears that combination of the reuse option with a suitable relative tolerance is required for the reuse option to be effective. Moreover, Fig. 5.1 shows that the reuse of Krylov vectors renders the Newton-Krylov method computationally cheaper than subiteration even for the nonlinear setting of the problem. Such a nonlinear setting is adverse for Newton-Krylov convergence, but not for subiteration convergence. On the other hand, for problems which are sufficiently linear, Newton-Krylov convergence is in general much faster than subiteration convergence. Finally, comparing the linear setting to the nonlinear one, Fig. 5.1 conveys that the differences in the convergence behaviour and cost are less pronounced for the Newton-Krylov method with a relative tolerance and reuse than for the basic method. We infer that this combination can conceal the adverse effects of nonlinearity on Newton-Krylov convergence and considerably improve the efficiency.

Finally, we compare the convergence behaviour of the Newton-Krylov method to subiteration for different computational time-step sizes. For this purpose we employ the Newton-Krylov method with reuse of Krylov vectors in subsequent Newton iterations and a relative tolerance for the linear residual of $\epsilon_1 = 10^{-2} ||r_i||$. In order to compare the convergence behaviour also for larger time steps, we restrict ourselves to small perturbations and set $\chi = 10^{-6}$, as for large initial perturbations the subiteration method can fail due to transient error growth. System and discretization parameters remain unchanged, see Tables 5.1 and 5.2 on page 78, except that, in addition to varying the time-step size, we adjust the number of elements in time such that the discretization per unit time remains the same. In particular, for $\tau = 1$ we use $N_{\mathbb{U}}^t = N_{\mathbb{Z}}^t = 24$, which is identical to the settings used for Fig. 5.1. Fig. 5.2 on the next page compares the convergence behaviour of the Newton-Krylov method and of the subiteration method for different timestep sizes. Whereas for small time steps subiteration converges properly, for large time steps subiteration converges non-monotonously. The transient divergence is caused by the nonnormality of subiteration; see Section 6.3.3 and Ref. [69] for a detailed analysis of the convergence behaviour of subiteration. Fig. 5.2 also shows that, in contrast to the norm of the subiteration residuals, the norm of the residuals of the GMRES iterates form a non-increasing sequence, which is due to the GMRES minimal-residual property. Moreover, the Newton-Krylov method does not exhibit a significant degradation of the convergence behaviour with increasing time-step size. The Newton-Krylov method therefore appears to be more robust and efficient than subiteration, in particular for large computational time steps. It is noteworthy that the Newton-Krylov method can attain convergence even for problems for which standard subiteration fails. For exemplification, we have included the convergence curve for $\tau = 16$ computed with the Newton-Krylov method. According to Ref. [69], the spectral radius of the subiteration operator is larger than one for this setting and, hence, subiteration is unstable; the curve is therefore not shown. Note that the Newton-Krylov method attains convergence despite the instability of the underlying subiteration operator. This is enabled by means of orthonormalization and underrelaxation; cf. Section 5.2.2.

Ref. [69] establishes that the convergence behaviour of the subiteration method changes similarly for an increasing fluid-to-structure mass ratio as for an increasing time-step size. We have found that the Newton-Krylov method displays a similar improvement over subiteration; the convergence curves for representative settings of the fluid-to-structure mass ratio, $\rho \alpha^0/M$, are shown in Fig. 5.3 on page 83. In particular, we varied the fluid density which translates into a variation of the fluid-to-structure mass ratio.



Figure 5.2: Residual reduction in the L^2 norm versus the iteration counter for the Newton-Krylov method with a relative tolerance and reuse and subiteration for $\tau = 2^0$ (----), $\tau = 2^1$ (--), $\tau = 2^2$ (···) and $\tau = 2^4$ (-·)(only for the Newton-Krylov method) with $N_{\mathbb{U}}^t = N_{\mathbb{Z}}^t = 24$, 48, 96 and 384, respectively. The linear and nonlinear residuals of the Newton-Krylov method are indicated by \circ and *, respectively, and the residuals of subiteration by \triangle .

5.3.2 Reuse of Krylov vectors in subsequent time steps

In this section, we investigate the Newton-Krylov method with reuse of Krylov vectors in subsequent time steps and assess the viability of the reuse option for nonlinear problems. Moreover, we compare the Newton-Krylov method in terms of robustness and computational cost to subiteration.

We consider the piston problem, provided with initial conditions that correspond to a periodic solution of the linearized system; see Section 3.3.2 and Ref. [70]. The initial conditions for the piston are set according to $z^0 = 10^{-1}$ and $\dot{z}^0 = 0$, unless stated otherwise. The system and discretization parameters are given in Table 5.3 and 5.4 on page 85, respectively. As the periodic solution is very smooth, few elements already suffice to accurately represent the solution. We remark that the ratio of time-step size over system oscillation period is such that, in combination with a sufficiently large amplitude, there is a substantial change in initial conditions between time steps, which ensures that the effect of the latter on the reuse of Krylov vectors can be assessed properly. The oscillation period of the fluid-structure system can be obtained from Ref. [70].

The iterative process is initialized each time step with an approximation



Figure 5.3: Residual reduction in the L^2 norm versus the iteration counter for the Newton-Krylov method with a relative tolerance and reuse and subiteration for $\rho \alpha^0/M = 6$ (----), 12 (--) and 24 (...) with $\tau = 1$ and $N_{\mathbb{U}}^t = N_{\mathbb{Z}}^t = 24$. The linear and nonlinear residuals of the Newton-Krylov method are indicated by \circ and *, respectively, and the residuals of subiteration by \triangle .

conforming to

$$z_0(t) = z^0 + \dot{z}^0 \ t \ \exp\left(-(t/\tau)^2\right), \qquad t \in [0,\tau].$$
(5.13)

This initial approximation is sufficiently differentiable to ensure that the initial conditions are satisfied. We remark, however, that other initial approximations are possible. We set the convergence tolerance to $\epsilon_0 = 10^{-3} ||r_0||$, i.e., in each time step we require the reduction of the initial residual by three orders of magnitude. In addition, we specify for the Newton-Krylov method the tolerance for the inner loop of the acceleration according to $\epsilon_1 = 10^{-1} ||r_i||$, i.e., relative to the current nonlinear residual, and an underrelaxation parameter of $\nu = 10^{-3}$.

To contrast the performance of the Newton-Krylov and the subiteration method, we employ three distinct settings of the model problem, which differ in the fluid density, ρ^0 , and in the initial piston deflection, z^0 . A variation in the fluid density translates into a variation in the fluid-to-structure mass ratio, $\rho^0 \alpha^0/M$. It is important to note that, according to Ref. [69], the spectral radius of the subiteration operator scales with the mass ratio. For the periodic setting of the model problem, we have chosen to vary the mass ratio rather than the timestep size in order to exclude a possible effect of reflections in the fluid; see Ref. [69]. By specifying $z^0 = 10^{-1}$ we select a nonlinear setting in order to investigate the viability of the Newton-Krylov method and the effectiveness of the reuse option under adverse conditions. In test case I, we select a setting which is favorable for fast convergence of the subiteration method, which is warranted by a sufficiently small spectral radius. Accordingly, we require a sufficiently small fluid density. Here, we set $\rho^0 = 2$ and $z^0 = 10^{-1}$. In test case II, we set $\rho^0 = 20$ and consider two different settings for z^0 , in particular, $z^0 = 10^{-1}$ and $z^0 = 10^{-3}$. Although subiteration is formally stable, with $z^0 = 10^{-1}$ it fails due to nonnormality-induced transient divergence. On the other hand, with a lower initial piston deflection of $z^0 = 10^{-3}$, failure of the subiteration method is avoided and a comparison with the Newton-Krylov method is possible. However, we remark that the latter setting is less nonlinear than the one for $z^0 = 10^{-1}$. In test case III, we set $\rho^0 = 200$ and $z^0 = 10^{-1}$. The subiteration method is unstable for this setting.

Fig. 5.4 plots the residual reduction versus the iteration counter for the subiteration and the Newton-Krylov method in time steps 1 and 10. Fig. 5.5 on page 86 plots the cumulative number of iterations versus the time-step counter, and Fig. 5.6 on page 86 plots the dimension of the Krylov space versus the timestep counter. Note that the slope of the cumulative-iteration-count curve is a measure of the computational cost per time step. For test case I, subiteration and the Newton-Krylov method converge comparatively fast in the first time step. Initially, most iterations of the Newton-Krylov method are spent on generating the Krylov space, see Fig. 5.6. However, in subsequent time steps, increasingly few Krylov vectors need to be added to the space due to reuse. This results in a decreasing number of iterations per time step and manifests in the gradually changing slope of the cumulative-iteration-count curve; see Fig. 5.5 left. Moreover, Fig. 5.6 indicates that the generation of a complete Krylov space (dimension $N_{\mathbb{Z}}^{t}P_{\mathbb{Z}} = 60$ for the considered discretization; cf. Table 5.4) is typically not required. The results for test case I demonstrate that reuse can render the Newton-Krylov method computationally cheaper than subiteration even under conditions that are favorable for the convergence of the subiteration method. For test case IIa with $z^0 = 10^{-1}$, the subiteration method fails due to nonnormality-induced transient divergence (curve not shown), whereas the Newton-Krylov method converges properly. For test case IIb with $z^0 = 10^{-3}$, subiteration converges after a period of initial divergence. In contrast, the Newton-Krylov method converges monotonously. The markedly different convergence behaviour translates into a significant discrepancy in computational cost; see Fig. 5.5 center. The convergence curves for the Newton-Krylov method for test cases IIa and IIb can hardly be distinguished; a slight difference is only apparent in the cumulative-iteration-count curves. For test case III, subiteration is unstable (curve not shown). Note that the Newton-Krylov method attains convergence despite the instability of the underlying subiteration method.

 $N^x_{\mathbb{U}}$

 $N^t_{\mathbb{U}}$

Table 5.3: System parameters for test cases I-III (* indicates a variable parameter).

z^0	\dot{z}^0	α^0	$ ho^0$	c^0	K	M	τ
*	0	1	*	0.5	1	1	1

Table 5.4: Discretization parameters for test cases I–III. $N_{\mathbb{P}}$

 $P_{\mathbb{U}}$

 $P_{\mathbb{A}}$

 $P_{\mathbb{Z}}$

 $P_{\mathbb{P}}$

 $N_{\mathbb{A}}$

 $N_{\mathbb{Z}}^t$



Figure 5.4: Residual reduction in the L^2 norm versus the iteration counter for subiteration (-----) and the Newton-Krylov method with reuse in subsequent time steps (--); the linear and nonlinear residuals of the Newton-Krylov method are indicated by \circ and $\ast,$ respectively, and the residuals of subiteration by \triangle ; time step 1 (top) and time step 10 (bottom); test case I (left), IIb (*center*) and III (*right*). In addition, for test case IIa (*center*) the Newton-Krylov method with reuse for $z^0 = 10^{-1} (\cdots)$ is included.



Figure 5.5: Cumulative number of iterations versus the time-step counter for subiteration (——) and the Newton-Krylov method with reuse in subsequent time steps (––); test case I (*left*), IIb (*center*) and III (*right*). In addition, for test case IIa (*center*) the Newton-Krylov method with reuse for $z^0 = 10^{-1}$ (···) is included.



Figure 5.6: Dimension of the Krylov space versus the time-step counter for the Newton-Krylov method with reuse in subsequent time steps (--); test case I (*left*), IIb (*center*) and III (*right*). In addition, for test case IIa (*center*) the Newton-Krylov method with reuse for $z^0 = 10^{-1}$ (· · ·) is included.

5.4 Concluding remarks

We have presented a novel solution method for fluid-structure-interaction problems, which overcomes the essential drawbacks of the customary subiteration method, viz., only conditional stability, potential convergence difficulties due to nonnormality and the inability to reuse information from previously solved similar problems. The proposed method is based on the observation that subiteration essentially constitutes a fixed-point iteration for the interface displacement. Therefore the GMRES acceleration of subiteration can be confined to the interface displacement. The GMRES acceleration operates within a Newton procedure, and accordingly, we refer to the method as Interface Newton-Krylov method, or simply as Interface-GMRES(R). Because the Krylov vectors involve only the interface degrees-of-freedom, the storage requirements for the Krylov space and the computational cost associated with the acceleration itself are negligible. The Interface-GMRES(R) acceleration of subiteration retains the segregated treatment of fluid and structure equations. This facilitates its implementation in codes that already use subiteration as a solver. Moreover, the Interface-GMRES acceleration allows for the optional reuse of the Krylov space in subsequent invocations of the GMRES method.

We assessed the convergence properties of the Interface-GMRES(R) method by numerical experiments on a prototypical fluid-structure-interaction problem. The results show that the efficiency of the Interface-GMRES(R) method depends on the convergence behaviour of the underlying subiteration method and on the nonlinearity of the problem. However, this effect can be concealed by means of a suitable relative tolerance in the inner loop of the acceleration method in combination with the reuse of the Krylov space. Our numerical results demonstrate that the Interface-GMRES(R) method is superior to the customary subiteration method in terms of robustness and efficiency. Under conditions that are favourable for the subiteration method but adverse for the Interface-GMRES(R) method, the latter still proves more efficient. Moreover, the accelerated method generally attains convergence even if the subiteration method is unstable. The reuse of the Krylov space can substantially enhance the efficiency of the method. However, it has to be exercised with some caution, as the reuse option can affect the robustness of the method. Our numerical results indicate that the combination of reuse with an appropriate relative tolerance render it a viable option also for nonlinear problems. After the initial construction of a suitable Krylov space, in general only occasional augmentations of the space are required. The reuse of the Krylov space in successive time steps enables the solution of the aggregated fluid-structure equations to sufficient accuracy in a few iterations per time step, even if subiteration separately requires tenfolds more or fails due to instability or nonnormality-induced divergence. It appears that the reuse of the Krylov space can render the computational expense of the Interface-GMRESR method comparable to loosely-coupled partitioned methods which perform only a single fluid and structure solution per time step; see Chapter 4. However, in contrast to such loosely-coupled methods, the Interface-GMRESR method leads to the solution of the aggregated equations, which enables conservation and improves stability and accuracy.

The proposed method is in principle generic. However, the specifics depend on the fluid-structure system under consideration. Further testing of this method is undertaken on the panel problem in Chapter 7.

Chapter 6 Error-amplification analysis of Interface-GMRES(R)

6.1 Introduction

In Chapter 5, we treated the theoretical background, algorithmic aspects and convergence studies of the Interface-GMRES method in a continuum setting of a generic fluid-structure-interaction problem. In the present chapter, we investigate the linear-algebra aspects of the Interface-GMRES method on the basis of properties of the error-amplification matrix for the aggregated system. To establish the error-amplification matrix, we consider a system of nonlinear-algebraic equations in conformity with discretizations of fluid-structure-interaction problems. We then investigate the iterative solution of a corresponding linear system by means of standard subiteration, and by means of the novel combined subiteration/GMRES method. By virtue of the linear-algebra setting, it is possible to derive precise expressions for the error-amplification properties of subiteration separately, and of subiteration combined with GMRES, with and without the reuse option. On the basis of the error-amplification matrix of subiteration, we can show that subiteration condenses errors into a low-dimensional subspace which can be associated with the interface degrees-of-freedom. We establish that this error-amplification matrix is highly rank-deficient. In particular, its rank is at most equal to the dimension N of the approximation space of the interface displacement. This implies that a Krylov method terminates in at most N steps, independent of the choice of the acceleration space, e.g., aggregated variables, structure variables, or interface variables. However, because computational cost and storage required by the Krylov acceleration itself increase with the dimension of the acceleration space, the acceleration on the interface variables is the most efficient. This distinguishes the considered Interface-GMRES method from approaches which apply the acceleration to the aggregated variables or to the structure variables; see Refs. [26, 38, 58]. The linear-algebra setting enables a clear explanation of the relation between the *local* GMRES acceleration (i.e., on the interface degrees-of-freedom), and the *global* error-amplification properties (i.e., for the aggregated system).

Moreover, the nonnormality of subiteration, which can induce non-monotonous convergence [69], can be traced immediately to properties of the erroramplification matrix. To elaborate the implications of nonnormality for the combined subiteration/GMRES method, we establish that the subiteration iterates span a Krylov-space corresponding to a preconditioned aggregated system. The properties of the subiteration-preconditioned GMRES method are then considered and, in particular, the relation between GMRES convergence and nonnormality of the subiteration preconditioner is addressed. We show that nonnormality can degrade the sharpness of GMRES convergence bounds.

Furthermore, we establish that the GMRES acceleration on the interface degrees-of-freedom generates an approximation to the Schur complement for the aggregated system. The GMRES acceleration and the reuse of Krylov vectors in subsequent invocations of GMRES are then assessed in terms of the approximation properties for the Schur complement, and in terms of the properties of the corresponding error-amplification matrices.

The theory is illustrated by numerical experiments on a model fluid-structureinteraction problem with a van-der-Pol oscillator. As this model problem exhibits qualitative changes in its solution behaviour in time, it is particularly suitable to investigate the effect of the reuse of Krylov vectors on the error-amplification properties of the subiteration-preconditioned GMRES method. In particular, we assess the potential and the limitations of the reuse option. Moreover, we provide convergence bounds for the methods in terms of matrix norms.

The contents of this chapter are organized as follows: Section 6.2 establishes the linear-algebra setting of the problem. Section 6.3 derives the erroramplification matrix of the subiteration method, and elaborates on the preconditioning perspective of subiteration. Section 6.4 analyses the error-amplification properties of the subiteration-preconditioned GMRES method with and without reuse in terms of the Schur-complement approximation and in terms of the erroramplification matrices. Section 6.5 provides numerical experiments and results. Section 6.6 contains concluding remarks.

6.2 Algebraic problem statement

In this section, we introduce the system of algebraic equations that emanates from a discretization of a generic fluid-structure-interaction problem; for a description of the latter we refer to Chapter 2 and Ref. [71]. We consider the system of linear-algebraic equations that arises from the application of Newton's method. This linear-algebraic system forms the basis of the error-amplification analysis for an inexact Newton method in Section 6.2.2. The subiteration method and the subiteration-preconditioned GMRES method can be construed as special instances of such an inexact Newton method.

6.2.1 The system of fluid-structure-interaction equations

The system of fluid-structure-interaction equations comprises the initial-boundaryvalue problems of the fluid and the structure, complemented by kinematic and dynamic conditions at the fluid-structure interface. We consider the generic space/ time variational formulation of such fluid-structure-interaction problems as presented in Chapter 2 and refer for a particular instance, viz., the piston problem, to Section 3.2. The space/time finite-element discretization of this generic variational statement yields a system of aggregated algebraic equations for the fluid-structure system, which we condense into the abstract form

_ /

$$\boldsymbol{R}(\mathbf{q}) = \boldsymbol{0},\tag{6.1a}$$

and, more specifically,

$$R_1(q_1, q_5) = 0, (6.1b)$$

$$R_2(q_1, q_2) = 0, (6.1c)$$

$$R_2(q_1, q_2, q_3) = 0 (6.1d)$$

$$\begin{aligned} &R_3(q_1, q_2, q_3) = 0, \\ &R_4(q_2, q_4) = 0 \end{aligned} \tag{6.1d}$$

$$\begin{array}{l}
R_4(q_3, q_4) = 0, \\
P_1(q_3, q_4) = 0, \\
P_2(q_3, q_4) = 0, \\
P_3(q_3, q_4) = 0, \\
P_4(q_3, q_4) = 0, \\
P_4(q_4, q_4) = 0, \\
P$$

$$R_5(q_4, q_5) = 0. (0.11)$$

We denote aggregated quantities by bold symbols. In particular, in (6.1), $\mathbf{R} := [R_1, R_2, R_3, R_4, R_5]^T$ and $\mathbf{q} := [q_1, q_2, q_3, q_4, q_5]^T$ denote the residual operators and variables associated with kinematic interface condition, fluid equations, dynamic interface condition, structure equation and the restriction of the structure variables to the interface, respectively. Note that the fluid and structure variables, q_2 and q_4 , are connected by the kinematic and dynamic interface conditions, R_1 and R_3 , via the fluid-interface displacement and the interface traction exerted on the structure, q_1 and q_3 , respectively. For transparency of the ensuing presentation, we have introduced an additional equation, Eq. (6.1f), given by

$$R_5(q_4, q_5) := T(q_4) - q_5 = 0, \tag{6.2}$$

where T represents the trace operator, which defines the structural displacement at the interface, q_5 , in terms of the structural variables, q_4 .

Given an initial estimate \mathbf{q}^0 , the application of Newton's method to the nonlinear system (6.1a) gives rise to a sequence of linear problems

$$\mathbf{A}(\mathbf{q}^{n+1} - \mathbf{q}^n) = -\mathbf{R}(\mathbf{q}^n), \tag{6.3a}$$

for $n = 0, 1, 2, \ldots$, which bear the particular form

$$\begin{pmatrix} A_{11} & 0 & 0 & 0 & A_{15} \\ A_{21} & A_{22} & 0 & 0 & 0 \\ A_{31} & A_{32} & A_{33} & 0 & 0 \\ 0 & 0 & A_{43} & A_{44} & 0 \\ 0 & 0 & 0 & A_{54} & A_{55} \end{pmatrix} \begin{bmatrix} q_1^{n+1} - q_1^n \\ q_2^{n+1} - q_2^n \\ q_3^{n+1} - q_3^n \\ q_4^{n+1} - q_5^n \end{bmatrix} = -\begin{bmatrix} R_1(\mathbf{q}^n) \\ R_2(\mathbf{q}^n) \\ R_3(\mathbf{q}^n) \\ R_4(\mathbf{q}^n) \\ R_5(\mathbf{q}^n) \end{bmatrix}.$$
(6.3b)

In (6.3b), A_{ij} denotes the Jacobian matrix corresponding to the residual-operator derivative $\partial R_i/\partial q_j$ and $R_i(\mathbf{q}^n)$ the residuals at the current linearization state \mathbf{q}^n . We remark that, in conformity with Eq. (6.2), $A_{55} = -I$ with I the identity matrix. If the initial estimate \mathbf{q}^0 is sufficiently close to the actual solution $\mathbf{\bar{q}}$ of Eq. (6.1a), then \mathbf{q}^n converges to $\mathbf{\bar{q}}$ as $n \to \infty$.

The consideration of Newton's method is useful mainly for conceptual reasons, as it provides the framework for the ensuing error-amplification analysis in Section 6.2.2. For the solution of fluid-structure interaction problems, however, the application of Newton's method is actually prohibitive; see also Section 4.2.1. The inherent interdependence between fluid and structure solutions induced by the interface conditions renders the matrix \mathbf{A} in (6.3) inseparable. This interconnection requires, in principle, a simultaneous solution of fluid and structure equations and, hence, causes the loss of software modularity; cf. Ref. [20]. Moreover, the disparate properties and scales inherent in the fluid and structure problems generally render the matrix A severely ill-conditioned. Finally, Eq. (6.3b) necessitates the evaluation of so-called *shape derivatives*, i.e., the derivative of the fluid equations with respect to a perturbation in the interface position, represented by the entry A_{21} . The difficulty in the evaluation of A_{21} pertains to the fact that a perturbation in the interface position in principle generates a perturbation throughout the entire fluid domain and, hence, A_{21} acts as a non-local operator. Although there exist approaches to facilitate the evaluation of the shape derivative, such as the method of spines [26], their applicability is typically restricted.

6.2.2 Error-amplification of inexact Newton methods

In order to bypass the aforementioned disadvantages of Newton's method, one generally reverts to alternative solution methods such as subiteration. For the error-amplification analysis of these methods, it is convenient to construe them, on linear approximation, as particular instances of an *inexact* Newton method. At variance with Eq. (6.3a), such inexact Newton methods determine an approximation $\tilde{\mathbf{q}}^{n+1}$ from the solution of

$$\tilde{\mathbf{A}}(\tilde{\mathbf{q}}^{n+1} - \mathbf{q}^n) = -\mathbf{R}(\mathbf{q}^n), \tag{6.4}$$

where $\hat{\mathbf{A}}$ denotes an approximate Jacobian that is 'in some sense' similar to \mathbf{A} , but easier to invert. The error induced by this approximation can be assessed as follows. The exact Newton method solves Eq. (6.3a) and, hence, on linear approximation, eliminates the error in a single step. Thus, on linear approximation, \mathbf{q}^{n+1} corresponds to the solution $\bar{\mathbf{q}}$ of (6.1a), and Eq. (6.3a) translates into

$$-\mathbf{A}\boldsymbol{\epsilon}^n = -\boldsymbol{R}(\mathbf{q}^n), \tag{6.5}$$

where $\epsilon^n := \mathbf{q}^n - \bar{\mathbf{q}}$ denotes the error in the approximation \mathbf{q}^n . Likewise, we denote by $\epsilon^{n+1} := \tilde{\mathbf{q}}^{n+1} - \bar{\mathbf{q}}$. Substituting Eq. (6.5) into Eq. (6.4) and adding suitable partitions of zero, we obtain

$$\tilde{\mathbf{A}}(\boldsymbol{\epsilon}^{n+1} - \boldsymbol{\epsilon}^n) = -\mathbf{A}\boldsymbol{\epsilon}^n,\tag{6.6}$$

which yields the error-amplification relation

$$\boldsymbol{\epsilon}^{n+1} = \mathbf{E}\boldsymbol{\epsilon}^n \,, \tag{6.7a}$$

with

$$\mathbf{E} := \mathbf{I} - \tilde{\mathbf{A}}^{-1} \mathbf{A} \tag{6.7b}$$

the error-amplification matrix of the inexact Newton method with approximate Jacobian $\tilde{\mathbf{A}}$. The inexact Newton method is formally convergent, if the spectral radius of the error-amplification matrix is smaller than unity, i.e., $\operatorname{spr}(\mathbf{E}) < 1$. With the setting $\tilde{\mathbf{A}} = \mathbf{A}$, we recover the exact Newton method, which yields $\mathbf{E} = \mathbf{0}$. We elaborate in Sections 6.3 and 6.4 that the customary subiteration method and the novel subiteration-preconditioned GMRES method with and without reuse can be associated with specific choices of $\tilde{\mathbf{A}}$. This enables a comparison of their respective error-amplification properties from a unified viewpoint.

6.3 Analysis of subiteration

This section presents a detailed analysis of the error-amplification properties of the subiteration method. To this end, we recall the basic subiteration algorithm in Section 6.3.1. In Section 6.3.2, we establish the error-amplification matrix, and we derive the precise form of its entries, which enables us to relate the erroramplification behaviour to specific entries. We show that the error-amplification matrix is rank-deficient, which provides the theoretical basis for the restriction of the GMRES acceleration to the interface degrees-of-freedom. Moreover, we consider the norm, spectral radius, and the nonnormality of the error-amplification matrix in Section 6.3.3. Nonnormality has important implications for the convergence of the subiteration method, and for the GMRES convergence bounds discussed in Section 6.4.5. Finally, in Section 6.3.4, we elaborate on the preconditioning perspective of subiteration, which provides the motivation of using it as a preconditioner to GMRES.

6.3.1 The subiteration method

For the error-amplification analysis of subiteration, we extend the definition of the subiteration algorithm in Section 4.2.2 to the linear-algebraic system given by Eq. (6.3). The subiteration method is then redefined as the following iterative procedure: Provided with an initial approximation of the structure displacement at the interface, $q_5^0(t)$, for n = 1, 2, ...

- (S1) Solve the kinematic condition: find q_1^n such that $R_1(q_1^n, q_5^{n-1}) = 0$
- (S2) Solve the fluid equations: find q_2^n such that $R_2(q_1^n, q_2^n) = 0$
- (S3) Solve the dynamic condition: find q_3^n such that $R_3(q_1^n, q_2^n, q_3^n) = 0$

- (S4) Solve the structure equations: find q_4^n such that $R_4(q_3^n, q_4^n) = 0$
- (S5) Determine the structure displacement at the interface: find q_5^n such that $R_5(q_4^n, q_5^n) = 0$

We remark that in an actual computation, the subiteration process solves the nonlinear equations. However, for our analysis we shall apply it to the linearized equations.

In accordance with Section 5.2.1, the subiteration process (S1)–(S5) can be conceived as a fixed-point iteration. For this chapter to be self-contained, we shall briefly recall from that section the relevant results pertaining to the ensuing investigation. For convenience, let us introduce the notation $z := q_5$. To facilitate the subsequent analysis of the subiteration method separately, and of the combined subiteration/GMRES method, we construe the subiteration process as a mapping from one structure interface displacement to the next, i.e.,

$$\mathsf{C}: z^n \mapsto z^{n+1} = \mathsf{C} z^n,\tag{6.8}$$

where C denotes the operator induced by the subiteration process as defined by (S1)-(S5); see Refs. [69,71] for further elaboration. Accordingly, the subiteration process can be characterized by recursion of the nonlinear operator C on the interface displacement z. Therefore subiteration constitutes a fixed-point iteration. The fixed point

$$\bar{z}: \bar{z} = \mathsf{C}\bar{z} \tag{6.9}$$

corresponds to the solution of (6.1). The nonlinear fixed-point problem (6.9) can be reformulated as

$$\mathsf{R}_{\mathrm{Sub}}\bar{z} = 0 \tag{6.10}$$

with $R_{Sub} := C - I$ the residual operator defined in conformity with (6.9) and I the identity. Clearly, \bar{z} being a solution of Eq. (6.10) is equivalent to \bar{z} being a fixed point of Eq. (6.9). The residual of an iterate is

$$r^{n} := \mathsf{R}_{\mathrm{Sub}} z^{n} = (\mathsf{C} - \mathsf{I}) z^{n} = \mathsf{C} z^{n} - z^{n} = z^{n+1} - z^{n}.$$
(6.11)

Note that, upon providing the subiteration process with the solution $\bar{q}_5 = \bar{z}$, the steps (S1)–(S5) yield the solution $\bar{\mathbf{q}} = (\bar{q}_1, \bar{q}_2, \bar{q}_3, \bar{q}_4, \bar{q}_5)$ of the aggregated nonlinear equations (6.1) in a single iteration, provided that the equations in each step are solved exactly.

6.3.2 Error-amplification analysis

This section analyses the error-amplification properties of the subiteration method. To this end, we recall from Section 6.3.1 that the subiteration method solves the nonlinear equations. For the error-amplification analysis we consider the linearization of these equations in conformity with Eq. (6.3). The subiteration algorithm

then translates into the solution of the following equations:

$$\begin{pmatrix} A_{11} & 0 & 0 & 0 & 0 \\ A_{21} & A_{22} & 0 & 0 & 0 \\ A_{31} & A_{32} & A_{33} & 0 & 0 \\ 0 & 0 & A_{43} & A_{44} & 0 \\ 0 & 0 & 0 & A_{54} & A_{55} \end{pmatrix} \begin{bmatrix} \tilde{q}_1^{n+1} - q_1^n \\ \tilde{q}_2^{n+1} - q_2^n \\ \tilde{q}_3^{n+1} - q_1^n \\ \tilde{q}_4^{n+1} - q_1^n \\ \tilde{q}_5^{n+1} - q_5^n \end{bmatrix} = - \begin{bmatrix} R_1(\mathbf{q}^n) \\ R_2(\mathbf{q}^n) \\ R_3(\mathbf{q}^n) \\ R_4(\mathbf{q}^n) \\ R_5(\mathbf{q}^n) \end{bmatrix}, \quad (6.12a)$$

which can be condensed into the form

$$\tilde{\mathbf{A}}_{\mathrm{Sub}}(\tilde{\mathbf{q}}^{n+1} - \mathbf{q}^n) = -\mathbf{R}(\mathbf{q}^n).$$
(6.12b)

From Eq. (6.12) it is apparent that we can associate a Jacobian matrix \mathbf{A}_{Sub} with the subiteration process that corresponds to the lower-triangular part of \mathbf{A} , i.e., it is identical to matrix \mathbf{A} in Eq. (6.3), but with the A_{15} block set to zero. Because $\tilde{\mathbf{A}}_{\text{Sub}}$ is lower block-triangular, the subproblems involving the inversion of the Jacobian $\tilde{\mathbf{A}}_{\text{Sub}}$ can be solved conveniently by forward substitution. Hence, the subiteration process essentially corresponds to a block Gauss-Seidel method.

Having identified the approximate Jacobian matrix $\tilde{\mathbf{A}}_{\text{Sub}}$ induced by the subiteration process, we determine the corresponding error-amplification matrix according to Eq. (6.7) as $\mathbf{E}_{\text{Sub}} := \mathbf{I} - \tilde{\mathbf{A}}_{\text{Sub}}^{-1} \mathbf{A} = \tilde{\mathbf{A}}_{\text{Sub}}^{-1} (\tilde{\mathbf{A}}_{\text{Sub}} - \mathbf{A})$ such that the error-amplification relation (6.7) translates into

$$\begin{bmatrix} \epsilon_1^{n+1} \\ \epsilon_2^{n+1} \\ \epsilon_3^{n+1} \\ \epsilon_5^{n+1} \\ \epsilon_5^{n+1} \end{bmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 & E_{15} \\ 0 & 0 & 0 & 0 & E_{25} \\ 0 & 0 & 0 & 0 & E_{35} \\ 0 & 0 & 0 & 0 & E_{45} \\ 0 & 0 & 0 & 0 & E_{55} \end{pmatrix} \begin{bmatrix} \epsilon_1^n \\ \epsilon_2^n \\ \epsilon_3^n \\ \epsilon_4^n \\ \epsilon_5^n \end{bmatrix},$$
(6.13a)

where the entries are defined as

$$E_{15} = -A_{11}^{-1}A_{15}, (6.13b)$$

$$E_{25} = A_{22}^{-1} A_{21} A_{11}^{-1} A_{15}, (6.13c)$$

$$E_{35} = -A_{33}^{-1}(A_{32}A_{22}^{-1}A_{21} - A_{31})A_{11}^{-1}A_{15}, ag{6.13d}$$

$$E_{45} = A_{44}^{-1} A_{43} A_{33}^{-1} (A_{32} A_{22}^{-1} A_{21} - A_{31}) A_{11}^{-1} A_{15}, ag{6.13e}$$

$$E_{55} = -A_{55}^{-1}A_{54}A_{44}^{-1}A_{43}A_{33}^{-1}(A_{32}A_{22}^{-1}A_{21} - A_{31})A_{11}^{-1}A_{15}.$$
 (6.13f)

Eq. (6.13) conveys that the subiteration error-amplification matrix is highly rankdeficient and, more specifically, its rank is equal to the minimum rank of the contributing block matrices in E_{55} .

It is important to note that in general the dimensions of the block matrices A_{11} , A_{33} and A_{55} associated with the interface variables are negligible compared to the dimensions of the fluid and structure block matrices, A_{22} and A_{44} , respectively, because the former refer to interface functions. Hence, the rank of the error-amplification matrix will generally be determined by the dimension of a block

matrix associated with the interface. This rank-deficiency has important consequences. In particular, it implies that the error components $\epsilon_1^n, \ldots, \epsilon_4^n$ are mapped onto zero and, hence, do not contribute to ϵ^{n+1} . The only error component that propagates from ϵ^n to ϵ^{n+1} is ϵ_5^n . More precisely, ϵ_5^n contributes to all components in ϵ^{n+1} due to the particular structure of the error-amplification matrix in (6.13). Thus, if the subiteration process is provided with the exact structure displacement $\bar{q}_5 = \bar{z}$, and, accordingly, $\epsilon_5 = 0$, then the method yields the solution to the aggregated equations in a single iteration. Moreover, from Eqs. (6.7) and (6.13), it is possible to derive a bound on the aggregated error, ϵ , in terms of the error in the structure interface displacement, ϵ_5 , as follows

$$\|\boldsymbol{\epsilon}^{n+1}\| \le \gamma \|E_{55}^n\| \|\epsilon_5^0\|, \qquad \gamma \in \mathbb{R}_+,$$
 (6.14)

with γ a constant depending on $\|\mathbf{E}_{\text{Sub}}\|$, but independent of n. If the spectral radius $\text{spr}(E_{55}) < 1$, then $\epsilon^n \to 0$ and, hence, $\mathbf{q}^n \to \bar{\mathbf{q}}$ as $n \to \infty$. Moreover, from Eq. (6.14) it is apparent that the precise convergence behaviour of subiteration is determined by the properties of E_{55} ; see Section 6.3.3 for further elaboration.

We can assign a particular meaning to the block matrix E_{55} in Eq. (6.13), which specifies the error amplification in the structure interface displacement, viz.,

$$\epsilon_5^{n+1} = E_{55}\epsilon_5^n. \tag{6.15a}$$

To this end, we linearize Eq.(6.8) around the solution $\bar{z} = \bar{q}_5$, subtract the fixedpoint equation and obtain

$$\epsilon_5^{n+1} = \mathsf{C}' \epsilon_5^n. \tag{6.15b}$$

From Eqs. (6.15a) and (6.15b), we can identify the block matrix E_{55} as the discrete representation of the subiteration-operator derivative C'.

Another observation that can be made from Eq. (6.13) is that the block matrices associated with the fluid and the structure, A_{22} and A_{44} , respectively, contribute to E_{55} only through projections onto the interface; see Eq. (6.13f). This indicates that the interior complexity of fluid and structure models yields only an indirect effect on the convergence of the subiteration method.

6.3.3 Nonnormality of subiteration

In this section, we briefly elaborate on the nonnormality of the subiteration operator, as such nonnormality has severe implications for the convergence behaviour of the subiteration method; see Ref. [69] for details. Moreover, nonnormality has implications also for the GMRES convergence bounds; see Section 6.4.5.

On linear approximation, convergence of the subiteration process requires that the spectral radius of the subiteration-operator derivative C' is strictly less than one, i.e., spr(C') < 1 or, equivalently, $spr(E_{55}) < 1$ in conformity with Eq. (6.15). Note, however, that the spectral radius only determines the asymptotic convergence behaviour of the method. The transient convergence behaviour is determined by the norm $\|C'\|$ or, equivalently, by $\|E_{55}\|$. Due to nonnormality of the subiteration-operator derivative C', the spectral radius and the norm can be disparate and, in particular, $spr(E_{55})$ can be much smaller than $||E_{55}||$. This disparity can give rise to non-monotonous convergence of the method, i.e., transient divergence can precede asymptotic convergence; cf. Ref. [24, ch. 2] for examples and, more specifically, Ref. [69]. This non-monotonous convergence behaviour was also observed in Section 5.3.1; see Fig. 5.2 on page 82 and Fig. 5.3 on page 83.

To explain such non-monotonous convergence behaviour, we note that the error in the n-th iterate can be bounded in conformity with Eq. (6.15) as

$$(\operatorname{spr}(E_{55}))^n \le \|\epsilon_5^n\| / \|\epsilon_5^0\| \le \|E_{55}^n\| \le \kappa(X)(\operatorname{spr}(E_{55}))^n,$$
 (6.16)

where $\kappa(X) := ||X|| ||X^{-1}||$ denotes the condition number of the matrix of eigenvectors, X, of E_{55} . For a normal matrix, the eigenvectors are orthonormal and, accordingly, $\kappa(X) = 1$, so that the upper and lower bounds in (6.16) coincide. For a nonnormal matrix, however, the eigenvectors are non-orthogonal and $\kappa(X)$ can be very large. In combination with the upper bound $||E_{55}^n|| \leq ||E_{55}||^n$, the bounds in (6.16) imply that if $\operatorname{spr}(E_{55}) < 1 < ||E_{55}||$ and $\kappa(X)$ is large, then the initial error can be amplified by many orders of magnitude before it eventually decreases at an asymptotic rate determined by $spr(E_{55})$; see Fig. 6.1 on the following page for an illustration. Clearly, this induces a severe degradation in the robustness and efficiency of the subiteration method. Moreover, the transient error growth can even cause failure of the iterative method despite formal stability. We remark that, if the spectral radius is sufficiently small, then the detrimental effect of nonnormality is concealed, i.e., the subiteration process converges monotonously. Ref. [69] establishes for a linear model problem that the spectral radius of the subiterationoperator derivative scales in the asymptotic limit with the time-step size. Hence, in order to control nonnormality-induced divergence, the computational time step can be reduced. Alternatively, the spectral radius can also be decreased by means of underrelaxation; see, for instance, Ref. [40]. However, either option generally renders the method inefficient.

6.3.4 Subiteration preconditioning

For the ensuing presentation of the combined subiteration/GMRES method in Section 6.4, it is elucidating to construe subiteration as a preconditioner for GMRES, following the concepts of Ref. [74].

The subiteration iterates computed according to the recursion (6.12) span a Krylov space corresponding to a left-preconditioned aggregated system:

$$\mathbf{q}^{n+1} - \mathbf{q}^n \in \operatorname{span}\{\tilde{\mathbf{A}}^{-1}\mathbf{r}^0, (\tilde{\mathbf{A}}^{-1}\mathbf{A})\tilde{\mathbf{A}}^{-1}\mathbf{r}^0, \dots, (\tilde{\mathbf{A}}^{-1}\mathbf{A})^n\tilde{\mathbf{A}}^{-1}\mathbf{r}^0\} = \mathcal{K}^{n+1}(\tilde{\mathbf{A}}^{-1}\mathbf{A}, \tilde{\mathbf{A}}^{-1}\mathbf{r}^0), \quad n = 0, 1, \dots, \quad (6.17)$$

where for ease of notation we have dropped the subscript in $\tilde{\mathbf{A}}_{\text{Sub}}$ and the tilde in $\tilde{\mathbf{q}}^{n+1}$, and denoted the Krylov space of dimension n+1 by \mathcal{K}^{n+1} . Moreover,



Figure 6.1: Illustration of the convergence behaviour of subiteration, exhibiting nonnormality-induced transient divergence: error reduction plotted versus iteration counter n; convergence of $||E_{55}^n||$ and convergence bounds.

in Eq. (6.17), we have invoked that, on linear approximation, the residual of the linear problem (6.3)

$$\mathbf{r}^{n} := \mathbf{A}\mathbf{q}^{n} - \mathbf{A}\mathbf{q}^{n-1} + \mathbf{R}(\mathbf{q}^{n-1}), \quad n = 1, 2, \dots$$
 (6.18)

is identical to the residual of the nonlinear problem (6.1), i.e., $\mathbf{r}^n = \mathbf{R}(\mathbf{q}^n)$. With this identity and provided with $\mathbf{r}^0 = \mathbf{R}(\mathbf{q}^0)$, the residual of the linear problem can be defined recursively as

$$\mathbf{r}^{n} := \mathbf{A}\mathbf{q}^{n} - \mathbf{A}\mathbf{q}^{n-1} + \mathbf{r}^{n-1}, \quad n = 1, 2, \dots,$$
 (6.19)

and Eq. (6.12) can be rewritten as

$$\mathbf{q}^{n+1} - \mathbf{q}^n = -\tilde{\mathbf{A}}^{-1}\mathbf{r}^n, \quad n = 0, 1, \dots$$
 (6.20)

The assertion that the subiteration iterates span a Krylov space in conformity with (6.17) follows straightforwardly by induction. Clearly, this assertion holds for n = 0:

$$\mathbf{q}^{1} - \mathbf{q}^{0} = -\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0} \in \operatorname{span}\{\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}\}$$
(6.21)

in conformity with Eq. (6.20). The induction makes use of the premise

$$\mathbf{q}^{n} - \mathbf{q}^{n-1} \in \operatorname{span}\{\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}, \dots, (\tilde{\mathbf{A}}^{-1}\mathbf{A})^{n-1}\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}\}.$$
 (6.22)
Starting from the left-hand-side of Eq. (6.17) and invoking Eqs. (6.19), (6.20) and (6.22), it holds that

$$\mathbf{q}^{n+1} - \mathbf{q}^{n} = -\tilde{\mathbf{A}}^{-1}\mathbf{r}^{n} = -\tilde{\mathbf{A}}^{-1}(\mathbf{A}\mathbf{q}^{n} - \mathbf{A}\mathbf{q}^{n-1} + \mathbf{r}^{n-1})$$

$$= -\tilde{\mathbf{A}}^{-1}(\mathbf{A}[\mathbf{q}^{n-1} - \tilde{\mathbf{A}}^{-1}\mathbf{r}^{n-1}] - \mathbf{A}\mathbf{q}^{n-1} + \mathbf{r}^{n-1})$$

$$= -\tilde{\mathbf{A}}^{-1}\mathbf{r}^{n-1} + (\tilde{\mathbf{A}}^{-1}\mathbf{A})\tilde{\mathbf{A}}^{-1}\mathbf{r}^{n-1}$$

$$\in \operatorname{span}\{\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}, \dots, (\tilde{\mathbf{A}}^{-1}\mathbf{A})^{n-1}\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}\}$$

$$+ (\tilde{\mathbf{A}}^{-1}\mathbf{A})\operatorname{span}\{\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}, \dots, (\tilde{\mathbf{A}}^{-1}\mathbf{A})^{n-1}\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}\}$$

$$= \operatorname{span}\{\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}, \dots, (\tilde{\mathbf{A}}^{-1}\mathbf{A})^{n-1}\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}, (\tilde{\mathbf{A}}^{-1}\mathbf{A})^{n}\tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}\}$$

$$= \mathcal{K}^{n+1}(\tilde{\mathbf{A}}^{-1}\mathbf{A}, \tilde{\mathbf{A}}^{-1}\mathbf{r}^{0}). \quad (6.23)$$

This completes the proof.

Our error-amplification analysis in Section 6.3.2 conveys that the eigenvalues of $\tilde{\mathbf{A}}^{-1}\mathbf{A}$ exhibit a much more favourable distribution than those of \mathbf{A} . In particular, Eq. (6.13) imparts that most eigenvalues of $\tilde{\mathbf{A}}^{-1}\mathbf{A}$ are 1. This makes subiteration a good preconditioner for the aggregated equations. Thus, instead of using subiteration as a solver, we will employ it as a preconditioner for GMRES; see Section 6.4.

The rank-deficiency of $\tilde{\mathbf{A}}^{-1}\mathbf{A}$ and, moreover, the fact that the error-amplification behaviour is essentially determined by the properties of E_{55} (see Eq. (6.14)), have important implications for the Krylov method. In particular, it follows that a Krylov method with subiteration preconditioning solves the linear system (6.3) in at most N steps, where $N := \operatorname{rank}(\tilde{\mathbf{A}}^{-1}\mathbf{A}) = \operatorname{rank}(E_{55})$. Moreover, on account of the rank-deficiency and the particular structure of $\tilde{\mathbf{A}}^{-1}\mathbf{A}$, the Krylov vectors need not contain all degrees-of-freedom, but only those associated with the interface degrees-of-freedom. This enables an efficient storage of the Krylov space and, moreover, renders the computational cost of the least-squares problem low in contrast to approaches which apply GMRES to the aggregated equations [26, 58], or to the Schur complement pertaining to the structure [38].

6.4 Analysis of Interface-GMRES(R)

In this section, we present a detailed error-amplification analysis of the combined subiteration/GMRES method. To this end, we introduce in Section 6.4.1 the GM-RES method with subiteration preconditioning. Moreover, we establish the Schur complement associated with the interface displacement. The Schur complement plays a central role in our investigation in that it enables us to analyse the considered solution methods in terms of their approximation properties for the Schur complement. Thus, the linear-algebra setting enables a clear explanation of the relation between the *local* GMRES acceleration (i.e., on the interface degrees-of-freedom), and the *global* error-amplification properties (i.e., for the aggregated system). The subiteration-preconditioned GMRES method allows for optional reuse

of Krylov vectors in subsequent invocations of GMRES. This reuse option is considered in Section 6.4.2. Next, we analyse the error-amplification properties of the subiteration-preconditioned GMRES method, first without the reuse option (Section 6.4.3), and subsequently with reuse (Section 6.4.4). Finally, in Section 6.4.5, we consider the convergence behaviour of subiteration-preconditioned GMRES, and we derive sharp upper bounds for the GMRES residual.

6.4.1 Subiteration-preconditioned GMRES

On account of the fact that the subiteration iterates span a Krylov space, and that subiteration exhibits favourable error-amplification properties, subiteration constitutes an apt preconditioner for GMRES. The combined subiteration/GMRES method comprises the following steps: firstly, one subiteration to condense the errors into a low-dimensional subspace and, moreover, to obtain a particular form of the Schur complement, and secondly, the application of GMRES to the Schurcomplement equation, which yields the solution of the structure interface displacement. Finally, another subiteration is required to compute the remaining components of the solution vector.

The Schur complement associated with the interface constitutes a key element in the analysis of the subiteration-preconditioned GMRES method. Therefore, let us first introduce the notion of the Schur complement associated with the structure interface displacement q_5 and the linear-algebraic system (6.3b). To this end, we translate the system (6.3b) into Schur form

$$\begin{pmatrix} A_{11} & 0 & 0 & 0 & A_{15} \\ A_{21} & A_{22} & 0 & 0 & 0 \\ A_{31} & A_{32} & A_{33} & 0 & 0 \\ 0 & 0 & A_{43} & A_{44} & 0 \\ 0 & 0 & 0 & 0 & S \end{pmatrix} \begin{bmatrix} q_1^{n+1} - q_1^n \\ q_2^{n+1} - q_2^n \\ q_3^{n+1} - q_1^n \\ q_4^{n+1} - q_5^n \end{bmatrix} = -\begin{bmatrix} R_1(\mathbf{q}^n) \\ R_2(\mathbf{q}^n) \\ R_3(\mathbf{q}^n) \\ R_4(\mathbf{q}^n) \\ R_S(\mathbf{q}^n) \end{bmatrix}$$
(6.24a)

with the Schur-complement matrix S and the corresponding right-hand-side vector $R_S(\mathbf{q}^n)$ defined respectively as

$$S := -A_{55}^{-1}A_{54}A_{44}^{-1}A_{43}A_{33}^{-1}(A_{32}A_{22}^{-1}A_{21} - A_{31})A_{11}^{-1}A_{15} - I$$
(6.24b)

and

$$R_{S}(\mathbf{q}^{n}) := A_{55}^{-1} \left(-R_{5}(\mathbf{q}^{n}) + A_{54}A_{44}^{-1} \left(R_{4}(\mathbf{q}^{n}) + A_{43}A_{33}^{-1} \left(-R_{3}(\mathbf{q}^{n}) + A_{32}A_{22}^{-1} \left(R_{2}(\mathbf{q}^{n}) - A_{21}A_{11}^{-1}R_{1}(\mathbf{q}^{n}) \right) + A_{31}A_{11}^{-1}R_{1}(\mathbf{q}^{n}) \right) \right). \quad (6.24c)$$

Note that the fifth equation in (6.24a) decouples. This equation constitutes the Schur-complement equation, viz.,

$$S(q_5^{n+1} - q_5^n) = -R_S(\mathbf{q}^n).$$
(6.25)

The expression of the right-hand-side vector $R_S(\mathbf{q}^n)$ in (6.24c) can be significantly simplified when provided with residuals $R_1(\mathbf{q}^n), \ldots, R_5(\mathbf{q}^n)$ based on a \mathbf{q}^n that has been generated by subiteration. To substantiate this assertion, we recall from Section 6.3.1 that, provided with an initial approximation of the structure interface displacement q_5^{n-1} , the subiteration steps (S1)–(S5) can be conceived as a sequence of mappings, viz., $q_5^{n-1} \mapsto q_1^n \mapsto q_2^n \mapsto q_3^n \mapsto q_4^n \mapsto q_5^n$. For convenience, and without loss of generality, let us shift the indices according to $q_5^n \mapsto q_1^n \mapsto q_2^n \mapsto$ $q_3^n \mapsto q_4^n \mapsto \hat{q}_5^{n+1}$. The 'hat' symbol serves to indicate that \hat{q}_5^{n+1} constitutes only an intermediate value which will be replaced by the subsequent invocation of GMRES. Note that, in conformity with the mapping above and Eq. (6.1), the residuals $R_1(\mathbf{q}^n), \ldots, R_4(\mathbf{q}^n)$ are identically zero, and $R_5(\mathbf{q}^n) = R_{\text{Sub}}(q_5^n) = \hat{q}_5^{n+1} - q_5^n$ in accordance with Eqs. (6.2) and (6.11). In summary, one subiteration generates a vector $\mathbf{q}^n = [q_1^n, q_2^n, q_3^n, q_4^n, q_5^n]$ such that the corresponding residual vector is $\mathbf{R}(\mathbf{q}^n) = [0, 0, 0, 0, \hat{q}_5^{n+1} - q_5^n]$. Provided with the residual vector $\mathbf{R}(\mathbf{q}^n)$ and noting that $A_{55} = -I$, Eq. (6.24c) simplifies considerably. Using the notation $z := q_5$ introduced in Section 6.3.1, the Schur-complement equation (6.25) can then be written as

$$S(z^{n+1} - z^n) = -(\hat{z}^{n+1} - z^n).$$
(6.26a)

With the definition (6.13f), the expression for the Schur complement (6.24b) can be specified as

$$S = E_{55} - I.$$
 (6.26b)

Recall from Section 6.3.2 that the matrix E_{55} constitutes the discrete representation of the subiteration-operator derivative C'. Hence, Eq. (6.26a) can be identified as a Newton iteration for the fixed-point residual equation (6.10), viz.,

$$R'_{\rm Sub}(z^{n+1} - z^n) = -R_{\rm Sub}(z^n), \tag{6.27}$$

and $S = R'_{Sub}$. Note that one subiteration is required to provide the residual $R_{Sub}(z^n)$ that forms the right-hand-side of the Schur-complement equation (6.26a).

The Schur-complement matrix S in (6.26a) is generally not known explicitly. However, if a Krylov method is used to solve the linear system (6.26a), then S is only required in the form of matrix-vector products, which can be approximated by finite differences. To this end, we note that on linear approximation

$$S(z_j^{n+1} - z^n) = r_j^{n+1} - r^n, (6.28)$$

where j indicates the counter for the GMRES iterations. In Eq. (6.28), the action of the Schur-complement matrix S on the increment vector $(z_j^{n+1} - z^n)$ yields the corresponding residual sensitivity $(r_j^{n+1} - r^n)$. For a Krylov method that makes use of the finite-difference approximation (6.28), we thus require a space of *search directions* in the form of updates around the current linearization state z^n , and the corresponding space of *residual sensitivities* in the form of increments around r^n . We then seek an approximation to the solution of Eq. (6.26a) from the search space \mathcal{K}^m , according to

$$(z^{n+1} - z^n) \in \mathcal{K}^m := \operatorname{span}\{z_j^{n+1} - z^n\}_{j=1}^{j=m},$$
(6.29)

where the z_j^{n+1} are generated by successive subiterations. For a formal proof that \mathcal{K}^m constitutes a Krylov space we refer to Ref. [71]. Denoting the vectors of search directions by $u_j = z_j^{n+1} - z^n$ with $j = 1, \ldots, m$, we collect them in a matrix $U_m := [u_1, \ldots, u_m]$. Moreover, we require the space of residual sensitivities, \mathcal{R}^m , that corresponds to the space of search vectors, viz.,

$$(r^{n+1} - r^n) \in \mathcal{R}^m := \operatorname{span}\{r_j^{n+1} - r^n\}_{j=1}^{j=m}.$$
(6.30)

Denoting the residual-sensitivity vectors by $v_j = r_j^{n+1} - r^n$ with $j = 1, \ldots, m$, we collect them in a matrix $V_m := [v_1, \ldots, v_m]$. In conformity with Eq. (6.28), V is the image of U under S, i.e.,

$$SU = V. (6.31)$$

We remark that the generation of the search and residual-sensitivity space implicitly builds an approximation to the Schur-complement matrix. If the spaces Uand V are complete and have full rank, then S can be obtained from Eq. (6.31) by $S = VU^{-1}$.

Provided with the linearization state z^n and the corresponding residual r^n , the generation of one pair of search direction and corresponding residual sensitivity, $(u_j, v_j) = (z_j^{n+1} - z^n, r_j^{n+1} - r^n)$, is at the expense of one subiteration. The subiteration is required for the evaluation of the residual $r_j^{n+1} := R_{\text{Sub}}(z_j^{n+1}) = z_{j+1}^{n+1} - z_j^{n+1}$. At the same time, it also generates a new search direction $u_{j+1} = z_{j+1}^{n+1} - z^n$. In practice, though, the search direction must be orthogonalized with respect to the previously generated search directions to avoid ill-conditioning of the search space.

To construct an approximation to the solution of the linear problem (6.26a), we make the following ansatz

$$z^{n+1} - z^n = \sum_{j=1}^{j=m} \alpha_j u_j, \tag{6.32}$$

with coefficients α_j that are determined from the requirement that the update (6.32) minimizes the residual of the Schur-complement equation (6.26a), viz.,

$$\bar{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} \|\boldsymbol{r}^n + S \sum_{j=1}^{j=m} \alpha_j u_j\| = \arg\min_{\boldsymbol{\alpha}} \|\boldsymbol{r}^n + \sum_{j=1}^{j=m} \alpha_j v_j\|,$$
(6.33)

in conformity with the finite-difference approximation (6.28) and Eq. (6.31). Hence, the considered Krylov method is a minimal-residual method. As the spaces \mathcal{K}^m and \mathcal{R}^m are generally not complete, Eq. (6.33) is solved in a least-squares sense (see, e.g., Ref. [22, ch. 5.3]), i.e., the coefficients $\bar{\alpha}$ are obtained from the solution of the normal equation

$$V^T V \bar{\boldsymbol{\alpha}} = -V^T r^n. \tag{6.34}$$

The computational cost involved in solving Eq. (6.34) is small in comparison to that of a subiteration, because the least-squares problem (6.33) is confined to the interface degrees-of-freedom. With coefficients $\bar{\alpha}$ determined from the solution of Eq. (6.34), the norm of the residual of the Schur-complement equation (6.26a) is given by

$$\xi_m := \|r^n + \sum_{j=1}^{j=m} \bar{\alpha}_j v_j\|.$$
(6.35)

If the residual norm ξ_m in (6.35) satisfies a given tolerance, the coefficients $\bar{\alpha}$ are used to determine the corresponding approximate solution z^{n+1} from Eq. (6.32). Note that in the actual nonlinear process, the residual norm given by Eq. (6.35) constitutes an estimate for the norm of the nonlinear residual; cf. Ref. [71].

Having established the equations solved by subiteration and GMRES, we can now establish the aggregated system that is solved by the combined subiteration/GMRES method. To this end, we recall that the subiteration process acting on the aggregated system (6.12) and the GMRES method acting on the Schurcomplement equation (6.26a) correspond to the mappings $q_5^n \mapsto \tilde{q}_1^{n+1} \mapsto \ldots \mapsto \tilde{q}_4^{n+1} \mapsto \hat{q}_5^{n+1}$ and $\hat{q}_5^{n+1} \mapsto q_5^{n+1}$, respectively. The combined subiteration/GMRES method then corresponds to the combined mapping $q_5^n \mapsto q_5^{n+1}$. Upon combining the Schur-complement equation (6.26a) with the system (6.12a) associated with subiteration, we obtain

$$\begin{pmatrix} A_{11} & 0 & 0 & 0 & 0 \\ A_{21} & A_{22} & 0 & 0 & 0 \\ A_{31} & A_{32} & A_{33} & 0 & 0 \\ 0 & 0 & A_{43} & A_{44} & 0 \\ 0 & 0 & 0 & A_{54} & -A_{55}S \end{pmatrix} \begin{bmatrix} \tilde{q}_{1}^{n+1} - q_{1}^{n} \\ \tilde{q}_{2}^{n+1} - q_{2}^{n} \\ \tilde{q}_{3}^{n+1} - q_{4}^{n} \\ q_{5}^{n+1} - q_{5}^{n} \end{bmatrix} = - \begin{bmatrix} R_{1}(\mathbf{q}^{n}) \\ R_{2}(\mathbf{q}^{n}) \\ R_{3}(\mathbf{q}^{n}) \\ R_{4}(\mathbf{q}^{n}) \\ \tilde{q}_{5}^{n+1} - q_{5}^{n} \end{bmatrix} .$$
(6.36)

Clearly, the system (6.36) retains the lower block-triangular structure induced by the subiteration method. As subiteration discards the A_{15} entry in (6.3b), two iterations are required to compute the aggregated solution vector: Upon solution of Eq. (6.36) only the component q_5^{n+1} , obtained from the Schurcomplement equation, corresponds to the solution of (6.3b). To determine the remaining components of the solution vector, one additional subiteration needs to be carried out; see also Section 6.3.1.

As we already mentioned in the beginning of this section with respect to the combined subiteration/GMRES method, prior to an invocation of GMRES, a subiteration needs to be carried out, which serves two purposes. Firstly, it condenses all error components into a low-dimensional subspace associated with the interface; cf. Section 6.3.2. Secondly, the subiteration process provides the right-hand-side residual for the Schur-complement equation. So far we have exclusively considered the solution of the linear-algebraic system (6.3b). However, in practice, the system of equations (6.1) is nonlinear. The application of Newton's method to the nonlinear system gives rise to a sequence of linear systems. We solve these linear systems by subiteration-preconditioned GMRES. Hence, we can alternatively classify the considered solution technique as a *Newton-Krylov method* [10], in which subiteration acts as a preconditioner.

Having established the subiteration-preconditioned GMRES method, let us briefly address some algorithmic aspects of the method; for details see Chapter 5 and also Ref. [71]. For computational efficiency, it is customary to set the convergence tolerance of the linear problem (6.3) relative to the norm of the actual residual of the nonlinear problem (6.1). Moreover, it is important to endow the method with Gram-Schmidt orthonormalization and underrelaxation. The former improves the robustness, the latter facilitates the subiteration process and allows the combination of GMRES with subiteration even if subiteration is formally unstable.

6.4.2 Reuse of Krylov vectors

We presented the optional reuse of Krylov vectors in Section 5.2.3. Next, let us briefly discuss the effect of reuse on the Schur complement. We recall from Section 6.4.1 that search and residual-sensitivity space implicitly approximate the exact Schur-complement matrix S; cf. Eq. (6.31). Likewise, reusing the search and residual-sensitivity space implies an approximation \tilde{S} to the exact S. With such an approximation \tilde{S} , the exact Schur-complement equation (6.26a) translates into

$$\tilde{S}(\tilde{q}_5^{n+1} - q_5^n) = -R(q_5^n), \tag{6.37}$$

which yields an approximation \tilde{q}_5^{n+1} to q_5^{n+1} . The system solved by the subiterationpreconditioned GMRES method with reuse can still be cast in the form of Eq. (6.36), but with the Schur-complement matrix S replaced by \tilde{S} . The effect of reuse and augmentation of the spaces on the error-amplification behaviour can thus be investigated in terms of the approximation properties for the Schur complement.

To place the reuse option into context, we briefly consider two alternative approaches for reusing computational information, viz., so-called *search space injection* [72] and *nested preconditioning* [12]. In particular, we establish commonalities and differences between these approaches and our methodology of reusing Krylov vectors. We consider first the *search space injection* from Ref. [72]. A commonality with our approach consists in storing a search space and its image under the operator, which allows for straightforward reuse and augmentation of these spaces. An important difference to Ref. [72] is that in our problem the operator is not known explicitly, but its action on a vector has to be evaluated by subiteration. Moreover, in contrast to [72], we consider the reuse of Krylov vectors in subsequent invocations of GMRES for which the left-hand-side matrix as well as the right-hand-side vector have changed. Finally, at variance with [72], in our approach the Krylov space need only contain the interface degrees-of-freedom and, hence, storage is not an issue. Next, we consider *nested preconditioning* from Ref. [12]. Likewise, this approach solves the linear problems by a Krylov method. Upon solution of the linear system, the generated Krylov space is condensed into a preconditioning matrix. However, an incomplete Krylov space would translate into a rank-deficient preconditioning matrix. To ensure that the preconditioner is invertable, some 'ad hoc' terms are added to this matrix in [12]. The preconditioning matrix is then applied to the subsequent linear system, which is again solved by a Krylov method. Successive solution of the linear systems thus generates a sequence of preconditioners, which are applied in a 'nested' way to the subsequent linear problems. The disadvantages of this approach pertain to the 'ad hoc' terms, which can interfere with the performance of subsequent preconditioners. Moreover, due to storage requirements, the number of preconditioners must be limited by restarts. These disadvantages can be avoided by methods in conformity with search space injection and reuse of Krylov vectors. Finally, we remark that Ref. [12] restricts reuse to subsequent Newton iterations, but does not consider reuse in subsequent time steps. However, it is in particular the reuse in subsequent time steps which can substantially reduce the computational cost in transient fluid-structure-interaction problems.

6.4.3 Error-amplification analysis

In this section, we derive the error-amplification matrix for the combined subiteration/GMRES method. In our analysis, the Schur complement from Section 6.4.1 plays a central role: the solution methods under consideration can all be conceived as instances of an inexact Newton method with a specific approximation to the Schur-complement matrix. Thus, an error-amplification analysis based on the Schur complement allows for a direct comparison of the subiteration process separately, the subiteration method with GMRES acceleration, and the subiteration method with GMRES acceleration and reuse, from a unified viewpoint.

As the subiteration-preconditioned GMRES method can be conceived as an instance of an inexact Newton method in conformity with Section 6.2.2, its erroramplification matrix is given in a general form by Eq. (6.7). The implied approximation $\tilde{\mathbf{A}}$ to the exact Jacobian \mathbf{A} of the Newton process is given by Eq. (6.36). However, to analyse the general case with an approximation to the Schur complement, for instance due to reuse of Krylov vectors, we consider the Jacobian matrix from Eq. (6.36) with the Schur complement S replaced by a complete and full-rank approximation \tilde{S} . Substituting the expressions for \mathbf{A} and $\tilde{\mathbf{A}}$ in Eq. (6.7) yields, after some straightforward manipulations, the error-amplification relation for the subiteration-preconditioned GMRES method:

$$\begin{bmatrix} \epsilon_1^{n+1} \\ \epsilon_2^{n+1} \\ \epsilon_3^{n+1} \\ \epsilon_4^{n+1} \\ \epsilon_5^{n+1} \end{bmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 & E_{15} \\ 0 & 0 & 0 & 0 & E_{25} \\ 0 & 0 & 0 & 0 & E_{35} \\ 0 & 0 & 0 & 0 & E_{45} \\ 0 & 0 & 0 & 0 & I - \tilde{S}^{-1}S \end{pmatrix} \begin{bmatrix} \epsilon_1^n \\ \epsilon_2^n \\ \epsilon_3^n \\ \epsilon_4^n \\ \epsilon_5^n \end{bmatrix}.$$
(6.38)

Eq. (6.38) expresses the error-amplification under a combined subiteration / GM-RES step with an approximation \tilde{S} to the Schur complement matrix S. In analogy to our observations in Section 6.3.2, we find that the error-amplification matrix in (6.38) is highly rank-deficient. In particular, the only non-zero eigenvalues are the ones of the lower block-diagonal matrix $I - \tilde{S}^{-1}S$, which expresses the mapping of one error in the structure interface displacement onto the next.

In fact, Eq. (6.38) is of a general form and encompasses the following three special cases: Firstly, upon setting S = -I, Eq. (6.38) reverts to Eq. (6.13a), and we recover the error-amplification matrix of the subiteration method. By virtue of the result from Section 6.4.1 that S constitutes the discrete approximation to $R'_{Sub} := \mathsf{C}' - \mathsf{I}$, the identity $\tilde{S} = \tilde{S}_{Sub} := -I$ conveys that \tilde{S}_{Sub} implies the most trivial approximation to R'_{Sub} . Secondly, upon setting $\tilde{S} = S$, the lower blockdiagonal entry is zero, and we obtain the error-amplification matrix of subiteration combined with 'exact' GMRES, i.e., no approximations are made to search and residual-sensitivity space. This is not to be confused with the error-amplification matrix of the exact Newton method which is identically zero; cf. Section 6.2.2. In the case of subiteration combined with GMRES, the matrix in Eq. (6.38) is nilpotent of index 2, i.e., $E^n \equiv 0, n \geq 2$; see, e.g., Ref. [32, ch. 1]. This implies that, to eliminate all error components and, hence, to obtain the complete vector of solution components, two applications of Eq. (6.38) are required. This is in agreement with the elaboration in Section 6.4.1, according to which one step of subiteration-preconditioned GMRES eliminates the error in q_5 , and a second step eliminates the error in the remaining components of the solution vector. Thirdly, Eq. (6.38) specifies the error-amplification matrix of subiteration combined with GMRES for the case that an approximation S to the exact Schur-complement matrix S is used. Such an approximation can be obtained, for instance, by reusing the search and residual-sensitivity spaces generated for the solution of the linear system (6.26a) in a previous Newton or time step. The error-amplification behaviour for the combined subiteration/GMRES method with reuse will be discussed in more detail in Section 6.4.4.

The above elaboration assumes that the Schur-complement approximation \tilde{S} is complete and has full rank and is thus invertible. In practice, the generation of a complete space V $(N \times N)$ with $N := \operatorname{rank}(\tilde{\mathbf{A}}^{-1}\mathbf{A}) = \operatorname{rank}(E_{55})$ is generally too expensive. However, it is usually also not required. An incomplete space \tilde{V} $(N \times k)$ with k < N, however, implies rank-deficiency of the Schur-complement approximation \tilde{S} . In that case, \tilde{S} can be determined as follows: With the solution

 $\bar{\alpha}$ from the normal equation (6.34), we consider the Newton update $\bar{z}^{n+1} = z^n + \bar{\alpha}U$. With $r^n = S\epsilon_5^n$ and upon subtracting the solution \bar{z} , we can identify the Schur-complement approximation in analogy with Eq. (6.38) as

$$\tilde{S}^{-1} = (\tilde{U}(\tilde{V}^T \tilde{V})^{-1} \tilde{V}^T).$$
(6.39)

We remark, however, that incompleteness of \tilde{U} and \tilde{V} implies that \tilde{S} has N-k eigenvalues equal to zero and, hence, $I - \tilde{S}^{-1}S$ has N - k eigenvalues equal to 1. This entails that the corresponding error components cannot be reduced in the available space. The precise error-amplification behaviour then depends specifically on the error in the initial approximation

6.4.4 Analysis of the reuse of Krylov vectors

Reuse of Krylov spaces enhances the efficiency of the method at the expense of robustness. In this section, we discuss the limitations of the reuse option based on the effect of reuse on the error-amplification behaviour of GMRES. We restrict our considerations to the case that search and residual-sensitivity space are complete and have full rank. For rank-deficient spaces the elaboration in Section 6.4.3 applies likewise.

With reuse, the residual-sensitivity space \tilde{V} is no longer equal to the image of the search space \tilde{U} under the current Schur complement S, i.e., $\tilde{V} \neq S\tilde{U}$. This is due to the fact that $\tilde{V} := \tilde{S}\tilde{U}$ may have been generated with, possibly multiple, different Schur-complement matrices, here collectively denoted by \tilde{S} . The benefit and limitations of the reuse option therefore depend pivotally on the similarity between S and \tilde{S} . If the Schur complement changes significantly in subsequent Newton iterations and time steps, the disparity between S and \tilde{S} can be substantial. Clearly, this is detrimental for the effectiveness of the reuse option. Reuse fails, if the spectral radius of the error-amplification matrix exceeds 1, i.e., if $\operatorname{spr}(I - \tilde{S}^{-1}S) > 1$, in conformity with Eq. (6.38). The reused space is then simply too inaccurate. In Section 6.5.4, we consider the properties of $I - \tilde{S}^{-1}S$ for a fluid-structure-interaction model problem.

For completeness let us briefly mention another 'failure mechanism' of the reuse option. If the reused space \tilde{V} is rank-deficient, then certain error components cannot be eliminated in \tilde{V} . If these error components are contained in the initial approximation, then the residual estimate stalls at some point, i.e., $\xi_m \rightarrow 0$ for $m \rightarrow N$ with $N := \operatorname{rank}(\tilde{\mathbf{A}}^{-1}\mathbf{A}) = \operatorname{rank}(E_{55})$. In that case, the algorithm breaks down; cf. Ref. [71] for further elaboration.

6.4.5 GMRES convergence bounds

To elaborate the implications of the nonnormality of the subiteration operator for GMRES convergence, we recall from Section 6.3.3 that nonnormality of E_{55} is associated with non-orthogonality of its eigenvectors. By virtue of the fact that E_{55} and the Schur complement $S = E_{55} - I$ have the same eigenvectors, S is

also nonnormal. This has consequences for the solution of the Schur-complement equation (6.26a) by GMRES. Nonnormality typically implies a severe degradation in the sharpness of the usual GMRES convergence bounds, which renders such bounds unreliable; cf. Ref. [24, ch.3]. To substantiate this statement, let us recall from Ref. [24] the following sequence of GMRES convergence bounds:

$$||r_k|| = \min_{p \in \mathcal{P}_k} ||p(S)r_0|| = \min_{p \in \mathcal{P}_k} ||Xp(\Lambda)X^{-1}r_0|| \le \kappa(X)\min_{p \in \mathcal{P}_k} ||p(\Lambda)|| ||r_0||, \quad (6.40a)$$

which implies

$$||r_k||/||r_0|| \le \kappa(X) \min_{p \in \mathcal{P}_k} \max_{\lambda \in \mathcal{L}} |p(\lambda)|, \qquad (6.40b)$$

where $\|\cdot\|$ denotes the standard 2-norm, \mathcal{P}_k is the set of k-th-order polynomials, defined by $\mathcal{P}_k(z) := 1 - \sum_{i=1}^{i=k} \alpha_k z^k$, \mathcal{L} is the spectrum of S and $\Lambda := \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$ is the diagonal matrix of eigenvalues of S, X is the matrix of right eigenvectors of S, and $\kappa(X) := \|X\| \|X^{-1}\|$ is the eigenvector-matrix condition number. Non-normality of S implies that the eigenvectors of S are non-orthogonal. Hence, the corresponding eigenvector matrix X can be severely ill-conditioned, in which case $\kappa(X)$ can be very large. It is then not clear whether GMRES indeed converges slowly, or whether the bound is simply a large overestimate of the actual residual norm. This renders the bounds (6.40) virtually useless. For nonnormal matrices, a sharp convergence upper bound follows from

$$||r_k|| / ||r_0|| \le \max_{||u||=1} \min_{p \in \mathcal{P}_k} ||p(S)u||.$$
(6.41)

This bound is sharp in the sense that there exists an r_0 for which the right-hand side value is actually attained. Eq. (6.41) delineates the worst-case convergence, independent of the specifics of the initial residual.

6.5 Numerical experiments

To illustrate the theoretical results from Sections 6.3 and 6.4, we conduct numerical experiments on a model fluid-structure-interaction problem. A concise setup of the model problem is provided in Section 6.5.1. Section 6.5.2 presents results for the subiteration method separately. In Section 6.5.3, we assess the convergence of the subiteration-preconditioned GMRES method and determine convergence upper bounds. In Section 6.5.4, we investigate the viability of reusing Krylov spaces by examining the properties of the corresponding error-amplification matrix.

6.5.1 Experimental setup

We conduct numerical experiments on a fluid-structure system consisting of the Euler equations of gas dynamics in one spatial dimension, in connection with a nonlinear structure, viz., the van-der-Pol oscillator, at the interface. For completeness, we shall briefly recall the Euler equations from Section 3.2.1. The Euler

equations in conservative form are given by

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{u})}{\partial x} = 0, \quad t \in (0, T), \quad x \in (0, \alpha(t)), \tag{6.42a}$$

with $\alpha(t)$ the position of the fluid-structure interface and

$$\mathbf{u} := \begin{pmatrix} \rho \\ \rho v \\ E \end{pmatrix}, \quad \mathbf{f}(\mathbf{u}) := \begin{pmatrix} u_2 \\ u_2^2/u_1 + p(\mathbf{u}) \\ (p(\mathbf{u}) + u_3)u_2/u_1 \end{pmatrix}, \quad p(\mathbf{u}) := (\gamma - 1)\left(u_3 - \frac{u_2^2}{2u_1}\right),$$
(6.42b)

and γ a constant, typically $\gamma = 1.4$. In Eq. (6.42b), ρ , v, E and p denote the density, velocity, total energy and pressure of the fluid, respectively.

Eq. (6.42) is subject to initial and boundary conditions $u_2(0,t) = 0$, $u_2(\alpha(t)) = u_1(\alpha(t))\dot{\alpha}(t)$ and $\mathbf{u}(x,0) = \mathbf{u}^0(x)$, where $\mathbf{u}^0(x)$ denotes prescribed initial conditions. The initial conditions are determined from the linearized model problem in Ref. [70]. However, we remark that for the considered small initial deflections of the oscillator, the influence of the initial conditions is small.

The structure model consists of the van-der-Pol oscillator, viz.,

$$M\ddot{z} + Kz + \eta\beta^{-2}(z^2 - \beta^2)\dot{z} = \pi - p^0, \qquad (6.43)$$

with z := z(t) the structure displacement from its equilibrium position and M, K, β and η suitable constants. The right member of Eq. (6.43) is composed of the stress $\pi := \pi(t)$ exerted by the fluid on the structure through the interface and the constant external pressure p^0 . The ordinary differential equation (6.43) is supplemented with the initial conditions $z(0) = z^0$ and $\dot{z}(0) = \dot{z}^0$. The fluid and the structure are connected by the dynamic interface condition $p(\mathbf{u}(\alpha(t), t)) = \pi(t)$ and the kinematic interface condition $\alpha(t) = \alpha^0 + z(t)$, where α^0 denotes the reference length of the fluid domain.

With $\eta = 0$, Eq. (6.43) simplifies to the equation of a harmonic oscillator considered in Ref. [71]. The associated fluid-structure system admits periodic solutions; cf. the linearized-system analysis in Section 3.3.2 and Ref. [70]. In contrast, for $\eta > 0$, the fluid-structure system given by Eqs. (6.42)–(6.43) behaves distinctly different: For non-vanishing initial conditions, the amplitude of z(t) in the vander-Pol equation increases to approximately 2β . As a consequence, the behaviour of the system at later times is in general distinctly different from its initial behaviour and, correspondingly, the subiteration operator changes. Therefore this model problem is particularly suitable to test the reuse of Krylov spaces under adverse conditions.

The fluid-structure system is discretized by means of space/time finite elements. The adopted discretization is essentially identical to that in Section 5.3. We consider the fluid-structure system for representative settings of the system parameters. The system and discretization parameters are listed in Table 6.1 and 6.2 on the next page, respectively. The discretization is sufficiently fine to ensure that the results are essentially mesh independent. The computational time step τ is chosen of the order of the period of the linearized system. We consider three representative settings of the parameter η which determines the nonlinearity of the structure model. In particular, we choose $\eta = 2$, 0.5 and 0, corresponding to strong, medium and vanishing structural nonlinearity, respectively. We remark that the parameter η moreover determines the initial-growth rate of the displacement, i.e., for larger η the initial growth of z(t) is more pronounced.

Table 6.1: System parameters (* indicates a variable parameter).

z^0	\dot{z}^0	α^0	$ ho^0$	c^0	K	M	η	β	τ
10^{-4}	0	1	20	0.5	1	1	*	0.025	8

Table 6.2: Discretization parameters.

$N^x_{\mathbb{U}}$	$N^t_{\mathbb{U}}$	$N_{\mathbb{A}}$	$N_{\mathbb{Z}}$	$N_{\mathbb{P}}$	$P_{\mathbb{U}}$	$P_{\mathbb{A}}$	$P_{\mathbb{Z}}$	$P_{\mathbb{P}}$
12	4	4	4	4	(2, 2)	4	4	3

6.5.2 Convergence of subiteration

In this section, we consider the convergence of the standard subiteration method. To this end, we investigate the properties of the corresponding error-amplification matrix and, in particular, of the lower diagonal block matrix E_{55} which yields the only non-zero eigenvalues of the error-amplification matrix; cf. Section 6.3.2.

For reference, Fig. 6.2 plots the structure displacement versus time for $\eta = 0, 0.5, 2$. The figure illustrates the significant change in the solution behaviour of the system over time. For $\eta > 0$, the amplitude of the structure oscillation increases by a factor of 500, viz., from an initial deflection $z^0 = 10^{-4}$ to approximately $5 \cdot 10^{-2}$ when the oscillation has settled into a (quasi-)periodic regime. For the larger value of η , the amplitude increases more rapidly on account of the stronger nonlinearity of Eq. (6.43) and, accordingly, the oscillation settles into its periodic regime faster. In contrast, for $\eta = 0$, the amplitude remains constant and equal to the initial deflection $z^0 = 10^{-4}$. For $\eta = 2$, a slight drift in the oscillation mean is visible in Fig. 6.2. Mesh refinement indicates that this drift is caused by discretization error and, moreover, that this discretization error does not significantly change the results presented in the sequel.

Fig. 6.3 on page 112 plots the spectral radius of the error-amplification matrix of subiteration versus time, viz., $(E_{55})_j := S_j + I$ in accordance with (6.26b), with S_j the Schur-complement matrix pertaining to time step j. The Schur complement



Figure 6.2: Structure displacement versus time for $\eta = 0.0(--)$, $\eta = 0.5(--)$ and $\eta = 2.0(\cdots)$.

is computed as $S = VU^{-1}$ with, in particular,

$$V^{T} := \nu^{-1} \left\{ \begin{pmatrix} R_{\text{Sub}} (u_{0} + \nu(u_{1} - u_{0})) \\ \vdots \\ R_{\text{Sub}} (u_{0} + \nu(u_{n} - u_{0})) \end{pmatrix} - \begin{pmatrix} R_{\text{Sub}} (u_{0}) \\ \vdots \\ R_{\text{Sub}} (u_{0}) \end{pmatrix} \right\}, \quad (6.44)$$

i.e., the image space V is determined by finite differencing of the nonlinear residual operator of the subiteration process, $R_{\rm Sub}$, as defined below (6.10). Throughout, we used $\nu = 10^{-8}$.

For $\eta = 2$ there is an initial growth in the spectral radius, before it decreases to an essentially constant value of approximately 2. This initial growth is apparently related to the strong initial growth of the solution; see Fig. 6.2. For $\eta = 0.5$, the initial growth of the spectral radius is absent, and the spectral radius decreases to approximately 3 in the periodic regime. In either case, however, the spectral radius is larger than one and, hence, the subiteration process diverges. For reference, we also computed the spectral radius for $\eta = 0$. This spectral radius is approximately 4. Clearly, for $\eta = 0$ the spectral radius remains essentially constant in time. For a detailed analysis of the convergence behaviour of subiteration for $\eta = 0$, we refer to Refs. [69, 71].



Figure 6.3: Spectral radius of the error-amplification matrix of subiteration, $\operatorname{spr}(E_{55})$, versus time for $\eta = 0.0$ (\circ), $\eta = 0.5$ (\triangle) and $\eta = 2$ (\Box).

6.5.3 GMRES convergence bounds

To assess the convergence of the subiteration-preconditioned GMRES method, we compute

$$\sigma_k := \sup_{u \in \mathbb{R}^n} J(u) \quad \text{with} \quad J(u) := \inf_{\alpha \in \mathbb{R}^k} J(u, \alpha), \quad J(u, \alpha) := \left\| u \right\|^{-1} \left\| \left(I - \sum_{i=1}^{i=k} \alpha_i S^i \right) u \right\|$$

Let us note that the convergence behaviour of the GMRES method pertaining to a particular initial residual r_0 is given by

$$\|r_k\| = \min_{\alpha} \|(I - \sum_{i=1}^{i=k} \alpha_i S^i) r_0\|.$$
(6.46)

For an illustration of the convergence behaviour of the method for the piston problem ($\eta = 0$) in specific instances we refer to Ref. [71]. One easily infers that Eq. (6.45) provides an upper bound to the residual reduction, i.e., $||r_k||/||r_0|| \leq \sigma_k$ for all $r_0 \in \mathbb{R}^n$. Moreover, the bound in (6.45) is sharp in the sense that there exists an initial residual for which the bound is actually attained, in contrast to the usual GMRES convergence bounds; see Section 6.4.5 and also Ref. [24, ch.3].

To determine the supremum over all u, we employ a steepest-ascent method (or *hill-climbing* algorithm), in combination with a line-search technique. Essentially, the steepest-ascent method determines the gradient J_u of J(u) with respect to u, and updates u according to $u + \theta J_u$, with θ a small number determined by the line-search strategy. This process is repeated until $J_u = 0$ and, accordingly, u corresponds to a supremum of J(u). As the functional J(u) is nonconvex, the process is started from multiple initial guesses, and the maximum over all local extrema is determined, to ensure that the global supremum is computed. Instead of computing the functional J(u) according to its definition in (6.45), we first orthonormalize the Krylov space $[u, Su, S^2u, \ldots, S^ku]$ to avoid ill-conditioning. Although the orthonormalization changes the coefficients of the GMRES polynomial, it does not change σ_k . Such a change of basis can be straightforwardly incorporated in the functional, but the expression for the corresponding gradient is prohibitively complicated. Therefore, in the numerical procedure we apply the gradient for the non-orthonormalized functional. Consequently, the accuracy of the gradient deteriorates with increasing dimension of the Krylov space on account of the loss of digits, and the supremum is typically not attained. As a result, for large dimensions of the Krylov space, the steepest-ascent process yields a reasonably sharp lower estimate of the upper bound σ_k .

Fig. 6.4 on the next page displays the upper bound σ_k versus the dimension of the Krylov space k for $\eta = 0, 0.5, 2$ in the first time step (*left*), i.e., in the time interval $0 \le t \le \tau$, and in the final time step (*right*), viz., $(j-1)\tau \le t \le j\tau$ with j = 50. For reference and validation purposes, the figure also plots the supremum of the functional J(u) over 10^6 random vectors u, versus the Krylovspace dimension k. Let us first note that the computed value of σ_k in all cases exceeds the supremum over the 10^6 random vectors. Hence, although the accuracy of the gradient deteriorates with increasing k and, accordingly, it cannot be ensured that the steepest-ascent algorithm converges to the supremum, it appears that the algorithm nonetheless yields a sharp estimate. Fig. 6.4 conveys that there is a lower bound to the dimension of the Krylov space in which a reduction of the residual can be ascertained. In particular, for dimensions k less than approximately 10 it holds that $\sigma_k \approx 1$. Furthermore, the figure shows that the convergence behaviour in the first time step deteriorates with increasing η . Conjecturally, this can be attributed to the growth of the solution in the first time step. For $\eta = 0$, the convergence behaviour is essentially identical for all time steps. For $\eta = 0.5, 2$ the convergence behaviour in the final time step is better than in the initial time step. In contrast to the growth regime, it appears that in the periodic regime the convergence behaviour improves with increasing η .

6.5.4 Analysis of the reuse of Krylov vectors

In this section, we investigate the viability of reuse of the Krylov space in subsequent invocations of GMRES. To this end, we consider the corresponding erroramplification matrix and, in particular, the lower diagonal block matrix $I - \tilde{S}^{-1}S$ which yields the only non-zero eigenvalues. We then examine its properties in terms of spectral radius, norm and pseudospectra.

We consider the reuse of the Krylov space generated in the first time step in the subsequent time steps. Fig. 6.5 on page 116 plots the spectral radius of the error-amplification matrix versus time. The spectral radius expresses the quality of the Krylov space of the first time step as an approximation to the Krylov spaces of the subsequent time steps. In particular, if the spectral radius is large, then the first Krylov space constitutes a poor approximation and reuse is ineffective. In



Figure 6.4: Convergence of the subiteration-preconditioned GMRES method: σ_k according to (6.45) (+) and the supremum of J(u) over 10^6 random vectors u (o) versus k for $\eta = 0$ (...), 0.5 (--), 2 (----), in the initial time step (*left*) and the final time step (*right*). y-axis in \log_{10} -scale.

fact, if the spectral radius exceeds 1, then reuse leads to divergence of the method. The figure indicates that the quality of the reused space degrades in the growth regime of the solution, reflecting the change in the solution behaviour and the corresponding change in the subiteration operator. As the solution settles into its periodic regime, the spectral radius becomes essentially constant, modulo minor erratic variations. Moreover, Fig. 6.5 indicates that for large η the deterioration of the approximation properties of the reused space is more severe. This can be attributed to the stronger nonlinearity of the structure associated with a larger value of η . Even for $\eta = 2$, however, the spectral radius remains smaller than one and, hence, the reuse option is formally convergent. Conjecturally, the spectral radius of the error-amplification matrix with reuse can exceed one, thus causing the reuse option to fail. However, for the considered test case such failure appears to be rare, and we have not observed failure of the reuse option despite testing it under various extreme conditions.

Next, we consider error bounds for the GMRES method with reuse in conformity with

$$(\operatorname{spr}(I - \tilde{S}^{-1}S))^n \le \|\epsilon_5^n\| / \|\epsilon_5^0\| \le \|(I - \tilde{S}^{-1}S)^n\| \le \kappa(X)(\operatorname{spr}(I - \tilde{S}^{-1}S))^n, (6.47)$$

where ϵ_1^n denotes the error in the interface displacement after *n* iterations, and $\kappa(X) := \|X\| \|X^{-1}\|$ is the condition number of the matrix of eigenvectors, *X*, of $I - \tilde{S}^{-1}S$. To this end, we plot in Fig. 6.6 on page 116 the norm of powers of the error-amplification matrix along with the bounds given by Eq. (6.47). The figure considers reuse of the Krylov space generated in the first time step in time step 50 for $\eta = 0$ (*left*), $\eta = 0.5$ (*center*) and $\eta = 2$ (*right*). In either case, the spectral radius is less than one, and the method with reuse is formally convergent. However, the condition numbers of the eigenvector matrices are very large, viz., $\kappa(X) = 3.1 \cdot 10^5$ for $\eta = 0$, $\kappa(X) = 6.9 \cdot 10^7$ for $\eta = 0.5$ and $\kappa(X) = 1.2 \cdot 10^8$

for $\eta = 2$. Moreover, the norm of the error-amplification matrix is smaller than one for $\eta = 0$, but it exceeds one for both $\eta = 0.5$ and $\eta = 2$. In combination with the upper bound $||(I - \tilde{S}^{-1}S)^n|| \leq ||I - \tilde{S}^{-1}S||^n$, the bounds in (6.47) then imply that for $\eta = 0$ convergence is monotonous, whereas for $\eta = 0.5, 2$ the initial error can be amplified by several orders of magnitude before it eventually decreases at an asymptotic rate determined by $\operatorname{spr}(I - \tilde{S}^{-1}S)$. For $\eta = 2$, the phase of transient divergence is much more pronounced, and the initial error can be amplified by more than five orders of magnitude before asymptotic convergence sets in. The transient divergence is indicative of the nonnormality of the GMRES method with reuse which is induced by the nonnormality of the underlying subiteration operator; cf. the discussion in Sections 6.3.3 and 6.4.5. We remark that the GMRES method without reuse converges monotonously; see Ref. [24] for the general theory, and Fig. 6.4 for an illustration.

Nonnormality also manifests itself in sensitivity of the spectrum to perturbations in the error-amplification matrix, e.g., due to discretization, linearization and round-off errors. To illustrate this effect, we plot in Fig. 6.7 on page 117 the ε -pseudospectra¹ of the error-amplification matrices considered in Fig. 6.6. A perturbation of the error-amplification matrix according to $(I - \tilde{S}^{-1}S) + E'$ with $||E'|| \leq \varepsilon$ can shift the eigenvalues anywhere within the corresponding ε -contour line; see [68]. For normal matrices the ε -pseudospectrum consists of all points in the complex plane at distance at most ε from the spectrum. For nonnormal matrices, however, the ε -pseudospectrum can be much larger. Although Fig. 6.7 indicates that as a result of nonnormality the spectrum of the error-amplification matrix is sensitive to perturbations, the presented results do not change significantly with mesh refinement or variations of the finite-differencing parameter ν involved in the evaluation of the image space according to Eq. (6.44).

6.6 Concluding remarks

We presented an error-amplification analysis of the subiteration-preconditioned GMRES method for fluid-structure-interaction problems. We considered the linearalgebra aspects of the subiteration method separately, and of the subiterationpreconditioned GMRES method, on the basis of properties of the error-amplification matrix for the aggregated fluid-structure system. We showed that the subiteration iterates span a Krylov space corresponding to a preconditioned aggregated system. The analysis of the error-amplification matrix of subiteration establishes that subiteration condenses errors into a low-dimensional subspace which can be associated with the interface degrees-of-freedom. Therefore, the GMRES acceleration of subiteration can be confined to the interface degrees-of-freedom and, accordingly, the method is also referred to as *Interface-GMRES*. The restriction of the acceleration to the interface degrees-of-freedom renders the storage require-

 $^{^1{\}rm The}$ pseudospectra plots in Fig. 6.7 were computed using the EigTool package by T.G. Wright, M. Embree and L.N. Trefethen; see [75] for further information.



Figure 6.5: Reuse of the Krylov space generated in the first time step $0 \le t \le \tau$ in subsequent time steps $(j-1)\tau \le t \le j\tau$ with $1 \le j \le 50$: Spectral radius of the error-amplification matrix versus time for $\eta = 0.0$ (\circ), $\eta = 0.5$ (\bigtriangleup) and $\eta = 2$ (\Box).



Figure 6.6: Reuse of the Krylov space generated in the first time step in time step 50 for $\eta = 0.0$ (*left*), $\eta = 0.5$ (*center*) and $\eta = 2$ (*right*): $||(I - S_1^{-1}S_{50})^n||(\circ)$, $||I - S_1^{-1}S_{50}||^n (\Box)$, $(\text{spr}(I - S_1^{-1}S_{50}))^n (\triangle)$ and $\kappa(X)(\text{spr}(I - S_1^{-1}S_{50}))^n (+)$ versus *n*. y-axis in \log_{10} -scale.

ments for the Krylov space and the computational cost of the acceleration itself low. The error-amplification analysis elucidates the connection between the *local* GMRES acceleration (i.e., on the interface degrees-of-freedom), and the *global* error-amplification properties (i.e., for the aggregated system).

An attractive feature of the subiteration-preconditioned GMRES method is that it enables the optional reuse of Krylov vectors in subsequent invocations of GMRES. This can substantially enhance the efficiency of the method, at the expense of robustness. We analysed the implications of reuse for the erroramplification behaviour. A pivotal element in the analysis is the observation that the GMRES acceleration on the interface degrees-of-freedom generates an approx-



Figure 6.7: Spectra (•) and $L^2 \varepsilon$ -pseudospectra of $(I - S_1^{-1} S_{50})$ in the complex plane for $\eta = 0.0$ (*left*), $\eta = 0.5$ (*center*) and $\eta = 2$ (*right*) and $\varepsilon = 10^{-9}$ (---), $\varepsilon = 10^{-8}$ (--), $\varepsilon = 10^{-7}$ (...), $\varepsilon = 10^{-6}$ (--), $\varepsilon = 10^{-5}$ (---), $\varepsilon = 10^{-4}$ (--) and $\varepsilon = 10^{-3}$ (...).

imation to the Schur complement for the aggregated system. The Schur complement allows for a concise expression of the error-amplification matrices, and enables the assessment of the GMRES acceleration and of the reuse of Krylov vectors in terms of the approximation properties for the Schur complement.

We illustrated the developed theory by numerical experiments on a model fluid-structure-interaction problem. In particular, we assessed convergence of the GMRES method with and without reuse in terms of spectral radius, norm and sharp convergence upper bounds. These bounds indicate that a minimum Krylovspace dimension is required to ensure a reduction of the residual. Moreover, the results show that significant changes in the solution due to nonlinearity can result in a degradation of the convergence behaviour.

Next, we investigated the viability of reusing Krylov vectors in subsequent invocations of GMRES. For the considered numerical experiments, the method with reuse is convergent, despite significant changes in the solution behaviour in time. Conjecturally, the spectral radius of the error-amplification matrix can exceed one, thus causing the reuse option to fail. However, such failure appears to be rare, and we have not observed it in the numerical investigations. This indicates that reuse constitutes a viable option, which renders it attractive for reducing the computational cost. However, the method with reuse appears to be affected by the nonnormality of the underlying subiteration operator. Thus, the method can exhibit transient divergence, whereas without reuse convergence is monotonous.

The presented error-amplification analysis of the subiteration-preconditioned GMRES methods is in principle generic. However, the specifics depend on the fluid-structure system under consideration.

Chapter 7 Assessment of Interface-GMRES(R) on the panel problem

7.1 Introduction

To demonstrate the versatility of the Interface-GMRES(R) method proposed in Chapter 5, in the present chapter we assess the solution method on an alternative model problem, viz., the panel problem; see, e.g., Ref. [55]. The panel problem is distinctly different from the piston problem considered in Chapter 5. In contrast to the one-dimensional piston problem, for the two-dimensional panel problem the interface degrees-of-freedom pertain to both space and time. This distinction is relevant for further testing of the Interface-GMRES(R) method, as the method operates on the interface degrees-of-freedom. Another relevant feature that distinguishes the panel problem from the piston problem relates to the aspect of parameter-dependent stability behaviour. Many fluid-structure-interaction problems can display instabilities such as flutter and divergence for certain parameter settings, whereas other parameter settings yield a stable behaviour. The piston problem does not have this feature, as it is (marginally) stable for all parameter settings. The panel problem, on the other hand, can exhibit instabilities such as flutter and divergence; see, e.g., Refs. [15, 21] and [5, ch.9]. An investigation of the convergence behaviour of the Interface-GMRES(R) method for different stability regimes is therefore relevant.

In our numerical experiments on the panel problem, we investigate the convergence behaviour of the Interface-GMRES(R) method, test its robustness and assess its computational cost. We consider the method with and without reuse of the Krylov space, and compare its performance to standard subiteration. We also investigate the effect of changes in the solution behaviour due to flutter on the convergence of the Interface-GMRES(R) method and on the effectiveness of the reuse option. Moreover, we assess the effect of the initial conditions on the system behaviour and on the convergence of Interface-GMRES(R). Besides strongly-

coupled solution methods such as subiteration and Interface-GMRES(R), we also consider loosely-coupled partitioned methods for the computation of a stable and an unstable fluid-structure system. Strongly-coupled methods can maintain the conservation properties at the fluid-structure interface, which renders them unconditionally stable. In contrast, loosely-coupled methods are typically energy increasing and, hence, numerically unstable; cf. Section 4.2.3.

The contents of this chapter are organized as follows: Section 7.2 presents a statement of the panel problem. Moreover, we introduce the characteristic system parameters and describe some distinct properties of the system. Section 7.3 presents numerical experiments and results. Section 7.4 contains concluding remarks.

7.2 **Problem statement**

Section 7.2.1 presents a statement of the panel problem. Section 7.2.2 establishes the characteristic system parameters and identifies some particular features of the system.

7.2.1 The panel problem

Below, we present a concise description of the panel problem, for an elaboration we refer to Ref. [55]. The upper side of the panel is exposed to an airstream, and its lower side to a cavity with still air; see Fig. 7.1 on page 122 for an illustration. We consider a panel with an infinite aspect ratio, which renders the problem essentially two-dimensional. The motion of the structure can then be described by the beam equation. Let x, y and t be spatial and temporal coordinates, respectively, $\alpha(x, t)$ the y-coordinate position of the fluid-structure interface and L the length of the beam. The mathematical formulation of the fluid-structure system comprises the Euler equations on $\Omega_{\alpha} := \{(x, y, t) : -\infty < x < \infty; \alpha(x, t) < y < \infty; 0 < t < T\}$ in connection with the beam equation at the interface $\Gamma_{\alpha} := \{(x, y, t) : 0 < x < L; y = \alpha(x, t); 0 < t < T\}$. We consider the Euler equations in conservative form:

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{u})}{\partial x} + \frac{\partial \mathbf{g}(\mathbf{u})}{\partial y} = 0, \quad (x, y, t) \in \Omega_{\alpha},$$
(7.1a)

with

$$\mathbf{u} := \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ E \end{pmatrix}, \quad \mathbf{f}(\mathbf{u}) := \begin{pmatrix} \rho u \\ \rho u^2 + p(\mathbf{u}) \\ \rho u v \\ (p(\mathbf{u}) + E)u \end{pmatrix}, \quad \mathbf{g}(\mathbf{u}) := \begin{pmatrix} \rho v \\ \rho u v \\ \rho v^2 + p(\mathbf{u}) \\ (p(\mathbf{u}) + E)v \end{pmatrix},$$
$$p(\mathbf{u}) := (\gamma - 1) \left(E - \frac{1}{2} \rho (u^2 + v^2) \right), \quad (7.1b)$$

and $\gamma = 1.4$. In Eq. (7.1b), ρ , u, v, E and p denote the density, the x- and y-component of the velocity, the total energy and the pressure of the fluid, respectively.

Eq. (7.1) must be supplemented with appropriate initial and boundary conditions. On $\partial \Omega_{\alpha} \backslash \Gamma_{\alpha}$ these are prescribed by

$$\mathbf{u}(x, y, 0) = \mathbf{u}^{0}(x, y), \qquad -\infty < x < \infty, \quad \alpha(x, 0) < y < \infty, \quad (7.2a)$$

$$(\rho v)(x, 0, t) = 0,$$
 $x < 0, x > L, 0 < t < T,$ (7.2b)

with $\mathbf{u}^0(x, y)$ the given initial conditions. Condition (7.2b) translates into the impermeability condition. Moreover, 'farfield boundary conditions' are imposed for $x \to \pm \infty$ and for $y \to \infty$. The interface conditions, i.e., the conditions on Γ_{α} , are specified below.

The governing equation for the beam is:

$$M\frac{\partial^2 z}{\partial t^2} + D\frac{\partial^4 z}{\partial x^4} = -\pi + \beta, \qquad 0 < x < L, \quad 0 < t < T,$$
(7.3)

where z designates the beam displacement from its equilibrium position, and the constants $M, D \in \mathbb{R}_+$ denote the mass and the bending stiffness of the beam, respectively. The right-hand member of Eq. (7.3) is the forcing term which is composed of the traction π exerted by the fluid on the structure through the interface, and the constant pressure β in the cavity underneath the panel. The cavity pressure is equal to the freestream pressure. Eq. (7.3) is subject to the initial and boundary conditions

$$z(x,0) = z^{0}(x), \qquad \qquad \frac{\partial z}{\partial t}(x,0) = \dot{z}^{0}(x), \qquad 0 < x < L, \qquad (7.4a)$$

$$z(0,t) = z(L,t) = 0, \qquad \frac{\partial z}{\partial x}(0,t) = \frac{\partial z}{\partial x}(L,t) = 0, \qquad 0 < t < T, \qquad (7.4b)$$

with $z^0(x), \dot{z}^0(x)$ the given initial conditions. The boundary conditions (7.4b) state that the beam is clamped on both sides.

The Euler equations and the beam equation are connected at the interface Γ_α by the kinematic conditions

$$\begin{aligned} (\rho v)|_{\Gamma_{\alpha}} &= \rho|_{\Gamma_{\alpha}} \frac{\partial \alpha}{\partial t}(x,t) + (\rho u)|_{\Gamma_{\alpha}} \frac{\partial \alpha}{\partial x}(x,t) , \qquad 0 < x < L , \quad 0 < t < T , \quad (7.5a) \\ \alpha(x,t) &= z(x,t) , \qquad \qquad 0 < x < L , \quad 0 < t < T , \quad (7.5b) \end{aligned}$$

and the dynamic condition

$$p(\mathbf{u}|_{\Gamma_{\alpha}}) = \pi(x, t), \qquad 0 < x < L, \quad 0 < t < T.$$
 (7.5c)

The condition (7.5a) constitutes a 'slip' boundary condition, which translates into the tangency of the flow to the moving beam and renders the interface impermeable. The condition (7.5b) identifies the interface position and the beam position. The condition (7.5c) implies equilibrium of the forces exerted on the interface by the fluid and the structure. Note that the interface conditions are imposed on the moving boundary Γ_{α} .



Figure 7.1: Illustration of the panel problem with infinite aspect ratio.

7.2.2 System behaviour

In this section, we establish the characteristic parameters of the panel problem and identify some distinct properties of the system.

To derive the characteristic parameters of the fluid-structure system, we cast the governing equations from Section 7.2.1 in dimensionless form. To this end, we introduce the nondimensional coordinates $\hat{x} = x/L$, $\hat{y} = y/L$ and $\hat{t} = t/(L/C_0)$, the nondimensional velocity $(\hat{u}, \hat{v}) = C_0^{-1}(u, v)$, the nondimensional density $\hat{\rho} = \rho/\rho_0$, the nondimensional pressure $\hat{p} = p/(\rho_0 C_0^2)$ and the nondimensional total energy $\hat{E} = E/(\rho_0 C_0^2)$, where C_0 denotes the speed of sound and ρ_0 is the reference density. The dimensionless form of the governing equations for the fluid is essentially identical to the dimensional form (7.1) with the provision that the dimensional state vector **u** is replaced by its nondimensional counterpart $\hat{\mathbf{u}} := (\hat{\rho}, \hat{\rho} \hat{u}, \hat{\rho} \hat{v}, \hat{E})^T$. Likewise, the nondimensional form of the kinematic and dynamic interface conditions is identical to the dimensional form with the dimensional quantities replaced by their nondimensional counterparts. The structure equation can be written in nondimensional form as

$$\frac{\partial^2 \hat{z}}{\partial \hat{t}^2} + \lambda^2 \frac{\partial^4 \hat{z}}{\partial \hat{x}^4} = \mu(-\hat{\pi} + \hat{\beta}), \qquad 0 < \hat{x} < 1, \quad 0 < \hat{t} < T/(L/C_0), \tag{7.6}$$

with

$$\lambda = \frac{LC_0^{-1}}{M^{1/2}L^2D^{-1/2}}, \quad \mu = \frac{\rho_0 L}{M}.$$
 (7.7a)

The parameter λ can be identified as the ratio of characteristic time scales of the fluid and the structure, and the parameter μ constitutes the ratio of characteristic

fluid mass to characteristic structure mass. In addition to the parameters specified in Eq. (7.7a), a third nondimensional parameter is required to specify the state of the fluid-structure system. This parameter is given by the Mach number

$$Ma = \frac{V_0}{C_0}, \qquad (7.7b)$$

which appears in the specification of the freestream velocity V_0 in the farfield boundary condition of the fluid.

Provided with the characteristic system parameters, let us briefly discuss some particular properties of the fluid-structure system. A distinct property of the panel problem is its ability to exhibit parameter-dependent stability behaviour. That is, the fluid-structure system can display instabilities such as flutter and divergence for certain parameter settings, whereas other parameter settings yield stable behaviour. For an elaboration of these phenomena we refer to the textbooks by Bisplinghoff [5, ch.9], Dowell [15] and Fung [21]. Instability of the fluid-structure system is a property that is shared by many fluid-structure-interaction problems and that is of significant practical importance. Since flutter and divergence can induce the failure of the structure, the analysis and prediction of such instabilities plays a crucial role in engineering design. For instance, in aerospace engineering, flutter and divergence impose constraints on the allowable operating conditions of aircraft. Hence, they need to be controlled by an adequate design; see, e.g., Ref. [11].

7.3 Numerical experiments

To assess the versatility of the Interface-GMRES(R) method and, in particular, its convergence behaviour and performance for multi-dimensional problems, we conduct numerical experiments on the panel problem. Section 7.3.1 specifies the setup of the numerical experiments. The first test case in Section 7.3.2 investigates the convergence behaviour of Interface-GMRES(R), with and without the reuse option, for representative settings of the system parameters. For reference purposes, we compare the results to the standard subiteration method. The second test case in Section 7.3.3 compares strongly-coupled and loosely-coupled methods for the computation of a stable and an unstable fluid-structure system. Moreover, we investigate the effect of instability on the convergence behaviour of subiteration and Interface-GMRES(R) and, in particular, on the reuse of the Krylov space. The third test case in Section 7.3.4 examines the effect of the initial conditions on the system behaviour and on the convergence of the Interface-GMRES(R) method. Since the Interface-GMRES(\mathbf{R}) method forms a special instance of a Newton-Krylov method, we shall refer to it in the sequel also as Newton-Krylov method, emphasizing the conjugation of Newton's method with a Krylov subspace method; cf. Section 5.2.

7.3.1 Experimental setup

We consider the panel problem from Section 7.2. The infinite-dimensional domain with $x \to \pm \infty$ and $y \to \infty$ is modeled by a truncated domain. In particular, in the x-direction inflow and outflow fluid boundary conditions are prescribed, and in the y-direction the domain is bounded by a solid wall at a distance of one from the panel. The distance of the solid wall to the panel is sufficiently large to ensure that the wall does not significantly influence the solution and the convergence behaviour of the solution methods.

To specify the initial conditions for the fluid-structure system, we note that the n-th mode shape of the beam assumes the generic form

$$g_n(x) = A_{1,n} \sin(\varphi_n x) + A_{2,n} \cos(\varphi_n x) + A_{3,n} \sinh(\varphi_n x) + A_{4,n} \cosh(\varphi_n x), 0 < x < L.$$
(7.8)

The φ_n are determined such that Eq. (7.8) admits a nontrivial solution for $g_n(x)$, subject to the homogeneous boundary conditions (7.4b). The $A_{1,n}, \ldots, A_{4,n}$ are then the coefficients that correspond to that nontrivial solution. In our experiments, we use initial conditions for the beam according to its first mode shape and, alternatively in Section 7.3.4, its second mode shape; see Table 7.1 for a specification of the parameters. That is, we set $z^0(x) = g_n(x)$ (n = 1, 2) and, moreover, $\dot{z}^0(x) = 0$. The initial conditions for the fluid are determined as the steady-state solution of the flow over a beam that is deflected according to the specified mode shape. The system parameters are given in Table 7.2, where τ denotes the length of the solution time interval. With Ma = 1.5, the flow is supersonic.

Table 7.1: Parameters for the first and second mode shape of the clamped beam according to Eq. (7.8).

n	$A_{1,n}$	$A_{2,n}$	$A_{3,n}$	$A_{4,n}$	φ_n
1	-0.4956	0.5044	0.4956	-0.5044	4.7300
2	-0.5002	0.4998	0.5002	-0.4998	7.8532

Table 7.2: System parameters (* indicates a variable parameter).

Case	λ	μ	Ma	au
Ι	0.25	*	1.5	0.05
II	*	10	1.5	*
III	*	10	1.5	0.05

The fluid-structure system is discretized by the space/time finite-element method with piecewise-polynomial base functions that are discontinuous in time and continuous in space. For the approximation space of the structure, we require in addition continuous differentiability in space. As base functions for the structure discretization we use Legendre polynomials. C^1 -continuity is then enforced by means of Lagrange multipliers. The base functions for the fluid are of modal type in conformity with Ref. [31, ch.3].

The time-discontinuous Galerkin discretization implies that displacement and velocity of the structure are discontinuous from one time slab to the next. However, since the fluid-boundary representation assumes a continuous displacement, the discontinuity in the structure displacement needs to be controlled. To render the discontinuity in the structure displacement and velocity negligible, we use polynomials of sufficiently high order for the approximation space of the structure.

Let us allude to the fact that the considered discretization does not maintain the conservation properties at the fluid-structure interface. To render the error pertaining to the conservation properties negligible, we choose a discretization for fluid and structure that is sufficiently fine.

The discretization parameters are given in Table 7.3, where the polynomial order of the approximation spaces associated with \mathbf{u} , α , z and π are, respectively, $(P_{\mathbb{U}}^x, P_{\mathbb{U}}^y, P_{\mathbb{U}}^t), (P_{\mathbb{A}}^x, P_{\mathbb{A}}^t), (P_{\mathbb{Z}}^x, P_{\mathbb{Z}}^t)$ and $(P_{\mathbb{P}}^x, P_{\mathbb{P}}^t)$, and the number of elements, N, is denoted accordingly. The number of elements in the x-direction is specified over the length of the panel. The discretization time step is equal to the length of the solution time interval. The discretization is sufficiently fine to ensure that the results are essentially mesh independent.

Table 7.3: Discretization parameters for test cases I-III.

$N_{\mathbb{U}}$	$N_{\mathbb{A}}$	$N_{\mathbb{Z}}$	$N_{\mathbb{P}}$	$P_{\mathbb{U}}$	$P_{\mathbb{A}}$	$P_{\mathbb{Z}}$	$P_{\mathbb{P}}$
(16, 24, 1)	(16, 1)	(16, 1)	(16, 1)	(3,3,3)	(2, 2)	(7, 7)	(3,3)

In each time slab, we provide an initial approximation of the structure displacement based on a linear extrapolation of the initial conditions conforming to

$$z_0(x,t) = z^0(x) + \dot{z}^0(x)t, \qquad 0 \le x \le L, \quad 0 \le t \le \tau.$$
(7.9)

The initial approximation of the structure displacement (7.9) constitutes a surface in space/time, in accordance with the space/time finite-element discretization of the problem.

We set the convergence tolerance to $\epsilon_0 = 10^{-4} ||r_0||$, i.e., we require a reduction of the initial residual by four orders of magnitude. In addition, we specify for the Newton-Krylov method the tolerance for the GMRES iteration according to $\epsilon_1 = 10^{-1} ||r_i||$, i.e., we use a relative tolerance for the convergence in the inner loop of the acceleration; cf. Section 5.2.2. Moreover, the underrelaxation parameter is set to $\nu = 10^{-2} ||r_0||$ for the Newton-Krylov method with reuse and to $\nu = ||r_i||$ for the method without reuse.

7.3.2 Assessment of Interface-GMRES(R) convergence

The first test case serves to assess the convergence behaviour of the Newton-Krylov method with reuse on the panel problem. For reference purposes, we compare the method to standard subiteration and to the Newton-Krylov method without reuse in terms of convergence behaviour and computational cost. We consider three distinct settings of the problem with physical parameters as given in Table 7.2 on page 124 and $\mu = 1,50,100$. The nondimensionalization in Section 7.2.2 conveys that the spectral radius of the subiteration-operator derivative scales with μ ; see also Ref. [69].

Fig. 7.2 plots the displacement of the beam in space/time. For all considered settings, the oscillation of the structure attenuates with time, indicating that the fluid-structure system is stable. Moreover, it is apparent that the beam deflection is downwind according to the direction of the flow. The convergence behaviour of the Newton-Krylov method with and without reuse and of the subiteration method is displayed in Fig. 7.3 on page 128 for time steps 1 and 50, respectively. In addition, we plot in Figs. 7.4 and 7.5, respectively, the dimension of the Krylov space and the cumulative number of iterations versus the time-step counter. The cumulative number of iterations specifies the total number of iterations required for convergence up to and including the time step under consideration. Fig. 7.3 illustrates that, initially, most iterations of the Newton-Krylov method are spent on generating the Krylov space. However, in subsequent time steps, increasingly fewer Krylov vectors need to be added to the space due to reuse; see also Fig. 7.4. This results in a decreasing number of iterations per time step and manifests in the gradually changing slope of the cumulative-iteration-count curve; see Fig. 7.5. In contrast, the number of iterations required by subiteration hardly changes in subsequent time steps. We infer that reuse can render the Newton-Krylov method computationally cheaper than subiteration even under conditions that are favorable for the convergence of subiteration; see Figs. 7.3 and 7.5 left with $\mu = 1$. Subiteration convergence deteriorates significantly with increasing μ , in contrast to Newton-Krylov convergence. Hence, a discrepancy in computational cost for larger μ emanates. For $\mu = 100$, subiteration diverges. Note that the Newton-Krylov method attains convergence despite the instability of the underlying subiteration method.

For reference, we have included in Figs. 7.3 and 7.5 the results for the Newton-Krylov method without reuse of the Krylov space. A comparison to the method with reuse clearly demonstrates the significant savings in computational cost that can be obtained by reusing the Krylov space.

To put our results into context, we remark that for an initial amplitude of the beam deflection of approximately 10^{-4} the system behaviour is close to linear. Preliminary studies indicate that for nonlinear system behaviour corresponding to larger initial amplitudes the performance of the Newton-Krylov method degrades only moderately. Moreover, we remark that our results are in good agreement with the results obtained on the piston model problem; see Section 5.3.2.



Figure 7.2: Test case I: Space/time displacement of the beam (colour bars) for system parameters according to Table 7.2 and $\mu = 1$ (*left*), $\mu = 50$ (*center*) and $\mu = 100$ (*right*).

7.3.3 Loosely-coupled versus strongly-coupled methods

In the second test case, we compare loosely-coupled and strongly-coupled partitioned methods for the computation of a stable and an unstable fluid-structure system. Moreover, we investigate the effect of a change in the solution behaviour due to flutter on the convergence of strongly-coupled solution methods. As instances of a strongly-coupled method, we consider subiteration and the Newton-Krylov method with and without reuse. The loosely-coupled method carries out only a single fluid and structure solution per time step; cf. Section 4.2.3. Throughout we obtain the initial approximation of the structure in the new time slab by linear extrapolation of the initial conditions in conformity with Eq. (7.9). We consider the fluid-structure system with parameters according to Table 7.2 on page 124 and two representative settings of λ . For $\lambda = 0.1$ the system undergoes flutter, whereas for $\lambda = 0.25$ the system is stable. The discretization parameters are specified in Table 7.3 on page 125.

Fig. 7.6 on page 130 plots the numerical solution of the beam displacement in space/time for the unstable system (left figure) and the stable system (right), each computed with a strongly-coupled method (top) and a loosely-coupled method (middle) for a time-step size $\tau = 0.05$. In addition, Fig. 7.6 bottom plots the solution obtained with a loosely-coupled method for $\tau = 0.01$. The solution behaviour for $\lambda = 0.1$ and $\lambda = 0.25$ is distinctly different. For $\lambda = 0.1$, the oscillation amplifies, which indicates flutter. For $\lambda = 0.25$, the oscillation attenuates, indicating stability of the fluid-structure system. Comparing for $\tau = 0.05$ the numerical solutions obtained with the loosely-coupled and the strongly-coupled method, we



Figure 7.3: Test case I: Residual reduction in the L^2 norm versus iteration number in time steps 1 (*top*) and 50 (*bottom*) for the Newton-Krylov method with reuse (—) and without reuse (–) and for subiteration (…); residual estimates and true residuals of the Newton-Krylov method are indicated by \circ and \Box , respectively, and residuals of subiteration by Δ ; $\mu = 1$ (*left*), $\mu = 50$ (*center*) and $\mu = 100$ (*right*). y-axis in \log_{10} -scale.

notice that the numerical solution of the loosely-coupled method exhibits wiggles. Since these wiggles are absent in the numerical solution of the strongly-coupled method, one can infer that they are caused by loose coupling. The wiggles increase with time and eventually cause the failure of the computation. Moreover, the figure displays that the loosely-coupled method yields an amplification of the solution for both the stable and the unstable system. Specifically, the loosely-coupled method yields a numerically unstable solution of the physically stable fluid-structure system; cf. also the discussion in Section 4.2.3. To control the numerical instability of loosely-coupled methods, the time-step size can be reduced; see Fig. 7.6 bottom with $\tau = 0.01$. The figure indicates that a reduction in the time-step size mitigates the numerical instability, but it does not remove it.

To investigate the effect of physical instability on the convergence behaviour of the Newton-Krylov method and, in particular, on the reuse of the Krylov space,



Figure 7.4: Test case I: Dimension of the Krylov space versus the time-step counter for the Newton-Krylov method with reuse in subsequent time steps; $\mu = 1$ (*left*), $\mu = 50$ (*center*) and $\mu = 100$ (*right*).



Figure 7.5: Test case I: Cumulative number of iterations versus the time-step counter for the Newton-Krylov method with reuse (—) and without reuse (—) and for subiteration (…); $\mu = 1$ (*left*), $\mu = 50$ (*center*) and $\mu = 100$ (*right*).

we consider the stable and the unstable system setting and plot in Fig. 7.7 on page 131 the cumulative number of iterations versus the time-step counter for the Newton-Krylov method and for subiteration as a reference. In addition, Fig. 7.7 right plots the dimension of the Krylov space versus the time-step counter. We remark that these figures plot upto a time step of n = 200 corresponding to computational time t = 10, whereas Fig. 7.6 plots only upto n = 100 (t = 5). Fig. 7.7 left displays a slight change in slope of the cumulative-iteration-count curve of the Newton-Krylov method with reuse for the unstable system setting. To explain this change in slope, we consider the evolution of the Krylov-space dimension plotted in Fig. 7.7 right. The figure exhibits that, after the initial construction of a sufficiently large Krylov space, the dimension of the space remains essentially constant upto a time step of approximately 100. Henceforth, the dimension of the Krylov space further increases in the case of the unstable system, which indicates that additional Krylov vectors need to be added to the space to attain convergence. However, this effect appears to be minor and reuse remains beneficial, see Fig. 7.7.



Figure 7.6: Test case II: Space/time displacement of the beam (colour bars): Solution computed with a strongly-coupled method and time-step size $\tau = 0.05$ (*top*), with a loosely-coupled method and time-step size $\tau = 0.05$ (*middle*) and with a loosely-coupled method and time-step size $\tau = 0.01$ (*bottom*) for system parameters according to Table 7.2 with $\lambda = 0.1$ (*left*) and $\lambda = 0.25$ (*right*).



Figure 7.7: Test case II: Cumulative number of iterations versus the timestep counter for the Newton-Krylov method with reuse in subsequent time steps (\Box) and without reuse (\circ) and for subiteration (\triangle) (*left*) and dimension of the Krylov space versus the time-step counter for the Newton-Krylov method with reuse (*right*); system parameters according to Table 7.2 with $\lambda = 0.1$ (——) and $\lambda = 0.25$ (——).

7.3.4 Effect of the initial conditions

The third test case investigates the effect of the initial conditions on the system behaviour and on the convergence of subiteration and the Newton-Krylov method. To this end, we consider the fluid-structure system provided with initial conditions corresponding to the first and second mode shape of the structure, respectively; cf. Section 7.3.1. For comparison purposes, the system and discretization parameters are chosen identical to those in Section 7.3.3; see also Tables 7.2 on page 124 and 7.3 on page 125 with $\tau = 0.05$. In particular, we consider $\lambda = 0.1, 0.25$. For $\lambda = 0.1$ the system undergoes flutter, whereas for $\lambda = 0.25$ the system is stable. The discretization is sufficiently fine to ensure that the results do not change significantly under further refinement.

The space/time displacement of the beam is plotted in Fig. 7.8 on the next page for initial conditions corresponding to the first and second mode shape and for $\lambda = 0.1$ and $\lambda = 0.25$. The figure indicates that, for a given λ , the initial behaviour of the solutions corresponding to different initial conditions is different, but their long-term behaviour is similar. We thus infer that the initial conditions determine only the initial behaviour of the system and that their effect on the system behaviour diminishes with time. This can be attributed to the fact that the long-term behaviour is determined by the least-stable modes of the system which are, of course, independent of the initial conditions. The long-term behaviour of the system essentially depends on the settings of the system parameters, which completely characterize the system and, hence, its spectral decomposition.

The convergence behaviour of subiteration and of the Newton-Krylov method for the system with initial conditions according to the second mode shape of the structure does not differ significantly from the one for the first mode shape and,



Figure 7.8: Test case III: Space/time displacement of the beam (colour bars) for initial conditions corresponding to the first mode shape of the structure (*left*) and the second mode shape (*right*); system parameters according to Table 7.2 with $\lambda = 0.1$ (*top*) and $\lambda = 0.25$ (*bottom*).

thus, essentially corresponds to the curves given in Fig. 7.7 on the preceding page for initial conditions according to the first mode shape. This is not unexpected, since the solution of the respective fluid-structure systems, and in particular their long-term behaviour, are very similar.

7.4 Concluding remarks

To assess the versatility of the Interface-GMRES(R) method, we investigated its convergence behaviour on the panel problem. This model problem is distinctly different from the piston problem considered in Chapter 5 on account of the fact that the interface degrees-of-freedom pertain to both space and time. Moreover, in case of the panel problem, the fluid-structure system exhibits parameter-dependent stability behaviour. In particular, it admits instabilities such as flutter and divergence. Our numerical results show that the Interface-GMRESR method is superior to the customary subiteration method in terms of robustness and efficiency: The Interface-GMRES(R) method converges even if the standard subiteration method diverges; and the reuse of the Krylov space considerably enhances the efficiency of the method and renders the Interface-GMRESR method computationally cheaper than subiteration even for settings which are favorable for subiteration convergence. These results are in good agreement with the ones obtained for the piston problem in Chapter 5.

Next, we considered loosely-coupled and strongly-coupled partitioned solution methods for a stable fluid-structure system and a system undergoing instability in the form of flutter. The loosely-coupled method can yield a numerically unstable solution for a physically stable system. To control the numerical instability, loosely-coupled methods are generally confined to small time steps. Moreover, our numerical results indicate that a change in the solution behaviour due to flutter can affect the effectiveness of reuse of the Krylov space. However, this effect appears to be minor and reuse remains beneficial.

Finally, we investigated the influence of the initial conditions on the system behaviour and on the convergence of the solution methods. Our results indicate that the initial conditions determine only the initial system behaviour and that their effect on the system behaviour diminishes with time. The long-term behaviour is determined by the least-stable modes of the system, and it essentially depends on the characteristic system parameters. Our results moreover convey that the convergence behaviour of subiteration and Interface-GMRES(R) is essentially determined by the characteristic system parameters rather than by the precise form of the initial conditions.

Although our results give an indication of the performance of the Interface-GMRES(R) method, we do not suggest that they provide final and conclusive information on the method in all respects. Further investigations are required, in particular with respect to the relation between the convergence behaviour of the method and the stability and nonlinearity of the problem.
Chapter 8 Conclusions

The concern of this thesis is to devise improved numerical methods for the computation of fluid-structure-interaction problems which are of great relevance in many engineering disciplines. The need for improved numerical techniques derives from the inadequacy of the customary numerical methods, viz., the inability of the usual discretization methods to maintain conservation at the fluid-structure interface, and the lack of robustness and efficiency of the standard solution methods.

For the development and analysis of novel numerical techniques we presented a generic space/time variational formulation of fluid-structure-interaction problems. Since this form does not make any stipulations on the specifics of the underlying models, it allows us to identify the generic features of fluid-structure interaction. As such generic features we identified the interconnection of a fluid and a structure subsystem by kinematic and dynamic interface conditions and, moreover, the free-boundary character of the interface. Furthermore, this generic variational statement enables us to formulate solution methods that are independent of the specifics of the underlying problem and, hence, in principle applicable to any fluid-structure-interaction problem. Free-boundary problems in general, and fluid-structure-interaction problems in particular, exhibit a deformation of the domain. We elaborated that this conflicts with the Eulerian fluid description, but that it can be treated straightforwardly by space/time methods. These space/time methods inherently account for a deformation of the domain in time, because they express the variational statement over the space/time domain. Therefore, we adopted the space/time finite-element method in this work.

The continuum fluid-structure-interaction problem possesses conservation properties, which can be lost under discretization. To assess conservation in the discrete numerical model, we established on the basis of a prototypical model problem the conservation properties of the continuum variational problem and of its discretization by the finite-element method. Conservation at the fluid-structure interface is only trivially maintained under restrictive compatibility conditions on the approximation spaces of the fluid and the structure, such as matching meshes and identical orders of interpolation at the interface. However, these conditions are prohibitively restrictive for practical use. To enable conservation also for incompatible approximations, we considered an alternative discretization based on coincidence and orthogonality conditions. The ability to maintain conservation at the interface also pertains to the pairing of discretization methods employed for the fluid and the structure. We showed that the trapezoidal method does generally not conserve momentum and energy for forced structures. Hence, it does not maintain momentum and energy conservation at the interface. In contrast, the conservation properties can be maintained by a space/time finite-element discretization that complies with the coincidence and orthogonality conditions; an additional computational expense is not incurred. Our numerical results demonstrate that maintaining conservation results in a much more accurate solution, whereas a violation of the conservation properties can induce instability of the numerical solution.

We considered monolithic methods, and strongly-coupled and loosely-coupled partitioned methods for the solution of the fluid-structure system. The monolithic (simultaneous) solution of fluid and structure by Newton's method incurs several severe disadvantages such as the loss of software modularity. The complications in Newton's method can be effectively circumvented by partitioning, i.e., by solving fluid and structure alternately subject to complementary partitions of the interface conditions. This iterative process is commonly referred to as subiteration. Whereas strongly-coupled partitioned methods repeat subiterations until convergence, loosely-coupled partitioned methods carry out only a single fluid and structure solution per time step. Therefore, strongly-coupled methods incur a greater computational cost per time step than loosely-coupled methods. However, strongly-coupled methods can maintain conservation at the fluid-structure interface, which renders them unconditionally stable. In contrast, loosely-coupled methods are typically energy increasing and, hence, numerically unstable. To control the numerical instability and for reasons of accuracy, the time-step size has to be restricted. For strongly-coupled methods, on the other hand, the admissible time step is determined by accuracy considerations only. Our numerical results demonstrate that conservative strongly-coupled methods are much more accurate than loosely-coupled methods and, hence, they can afford much larger time steps for the same level of accuracy. Moreover, the results show that conservative strongly-coupled methods achieve a greater accuracy with fewer iterations than loosely-coupled methods and are therefore computationally more efficient. Our results also indicate that for strongly-coupled methods which are not conservative the improvement in accuracy over loosely-coupled methods does not seem to justify the computational cost associated with multiple subiterations.

Although subiteration constitutes an apt solver for many problems, it suffers from three essential drawbacks, viz., only conditional stability, potential convergence difficulties due to nonnormality, and the inability to reuse information from previously solved similar problems. To overcome these deficiencies, we proposed to combine subiteration with GMRES acceleration. Since the combined subiteration/GMRES method allows for an algorithmic description in a generic, continuum setting, it is essentially applicable to any fluid-structure-interaction problem and, moreover, its convergence behaviour is asymptotically independent of the underlying discretization. Because subiteration can be conceived as a fixed-point iteration on the interface variables, the GMRES acceleration can be confined to the interface degrees-of-freedom. Thus, the corresponding Interface-GMRES acceleration requires only negligible computational resources. An additional asset of the presented acceleration method is that it enables the optional reuse of Krylov vectors in subsequent invocations of GMRES, which we refer to as GMRESR. Such reuse can considerably enhance the efficiency of the method; however, it comes at the expense of robustness and, therefore, it has to be exercised with some caution. The implementation of the Interface-GMRES(R) method in codes which use subiteration as a solver is straightforward, because the method retains the modularity of the underlying subiteration method. Our numerical results on the piston problem show that the Interface-GMRES(R) acceleration of subiteration leads to a substantial improvement in robustness and efficiency. Even under conditions that are favourable for the conventional subiteration method and adverse for the Interface-GMRES(R) accelerated method, the accelerated method still proves more efficient. The reuse of Krylov vectors usually results in significant computational savings in that, after the initial construction of a suitable Krylov space, only occasional augmentations of the space are required. The aggregated fluid-structure system is then typically resolved in a few iterations per time step, even if the subiteration method separately requires tenfolds more or fails due to instability or nonnormality-induced divergence.

We complemented the algorithmic description of the Interface-GMRES(R)method by an analysis of its error-amplification properties. We showed that subiteration constitutes an appreconditioner for the aggregated equations, since it condenses errors into a low-dimensional subspace that can be associated with the interface degrees-of-freedom. Because the rank of the subiteration error-amplification matrix is at most equal to the dimension N of the approximation space of the interface variables, the GMRES acceleration of subiteration terminates in at most Nsteps, independent of the acceleration space, e.g., aggregated variables, structure variables, or interface variables. However, as the computational expense and storage required by the Krylov acceleration itself increase with the dimension of the acceleration space, the acceleration on the interface variables is the most efficient. Thus, the linear-algebra setting elucidates the connection between the local GM-RES acceleration (i.e., on the interface variables) and the global error-amplification properties (i.e., for the aggregated system). Moreover, we showed that the nonnormality of the subiteration preconditioner induces a degradation in the sharpness of the usual GMRES convergence bounds. To assess the linear-algebra properties of the Interface-GMRES method with and without reuse, we exploited the fact that the acceleration on the interface degrees-of-freedom generates an implicit approximation to the Schur complement of the aggregated equations. This Schur complement allows for a concise expression of the error-amplification properties and, moreover, enables us to analyse the convergence of the respective methods in terms of norms, spectral radii and sharp convergence upper bounds. Conjecturally, the spectral radius of the error-amplification matrix associated with reuse can exceed one, which implies failure of the reuse option. However, such failure appears to be rare, and we have not observed it in our numerical investigations. This indicates that reuse constitutes a viable option. However, the method with reuse appears to be affected by the nonnormality of the underlying subiteration operator. Thus, the method can exhibit transient divergence, whereas without reuse convergence is monotonous.

We demonstrated the versatility of the Interface-GMRES(R) method through numerical experiments on the panel problem. This model problem is distinctly different from the piston problem in that the interface degrees-of-freedom pertain to both space and time. Moreover, in case of the panel problem, the fluid-structure system exhibits parameter-dependent stability behaviour. In particular, it admits instabilities such as flutter and divergence for certain parameter settings. Our numerical results show that the Interface-GMRESR method is superior to the customary subiteration method in terms of robustness and efficiency: The Interface-GMRES(R) method converges even if the standard subiteration method diverges; and the reuse of the Krylov space considerably enhances the efficiency of the method and renders Interface-GMRESR computationally much cheaper than subiteration even for settings which are favorable for subiteration convergence. These results are in good agreement with the ones obtained for the piston problem. Our numerical results indicate that a change in the solution behaviour due to flutter can affect the effectiveness of reuse of the Krylov space in Interface-GMRESR. However, this effect appears to be minor, and reuse remains beneficial. The results also indicate that the convergence behaviour of subiteration and Interface-GMRES(R) is essentially determined by the characteristic system parameters rather than by the specifics of the initial conditions. Finally, we demonstrated that loosely-coupled partitioned methods can yield a numerically unstable solution for a physically stable fluid-structure system. To control the numerical instability, loosely-coupled methods are generally confined to small time steps.

Chapter 9 Recommendations

In this thesis, we investigated numerical methods for fluid-structure-interaction problems. Based on the results of this investigation, we formulate below several recommendations for future research and applications.

In Chapter 3, we considered a discretization that maintains conservation at the fluid-structure interface under incompatible approximations, i.e., non-matching meshes and/or different orders of approximation at the interface. This concept was established on the basis of a one-dimensional fluid-structure-interaction model problem. The relevance of conservation for the stability and accuracy of the numerical solution was demonstrated by numerical experiments on the one-dimensional piston problem. An important continuation of this research consists in the extension of this concept to multi-dimensional problems. It is anticipated that conservation in multiple dimensions can be established under the appropriate orthogonality conditions on the approximation spaces of certain interface variables. In particular, conservation requires the orthogonality of the pressure difference to the interface velocity, and the orthogonality of the velocity difference to the interface pressure. We expect that for multi-dimensional problems maintaining the conservation properties is similarly beneficial.

In Chapter 4, we assessed the stability, accuracy and efficiency of stronglycoupled and loosely-coupled partitioned methods. Since the piston problem considered in our numerical experiments exhibits only a single mode, stability and accuracy cannot be distinguished clearly. This is only possible for multiple-mode problems. Thus, we recommend complementing the investigation by a comparison of strongly-coupled and loosely-coupled methods on a multiple-mode problem such as the panel problem from Ref. [55]. We expect that the superiority of conservative strongly-coupled methods over loosely-coupled methods will be even more pronounced, because for loosely-coupled methods the stability of modes with time scales smaller than one actually wants to resolve can impose a severe restriction on the admissible time-step size. To further improve the efficiency of strongly-coupled methods, we considered in Section 4.2.6 two concepts to reduce the computational cost associated with the fluid solution. To this end, we proposed an approximate fluid solution within the subiteration process and, moreover, the confinement of the fluid solution to a truncated domain in the vicinity of the interface. Whereas the former approach is straightforward in terms of implementation, the latter approach can be involved. In view of the considerable computational savings that an approximate fluid solution can deliver, we recommend testing this approach on different problems. Although principally an approximate fluid solution suffices for the convergence of the subiteration process, the specifics of this approach depend on the considered fluid-structure system.

In Chapter 5, we proposed to accelerate the subiteration process by means of Interface-GMRES(R), and we demonstrated the performance of the accelerated method by numerical experiments. To further enhance the efficiency of the Interface-GMRES(R) method, we propose to combine it with an approximate fluid solution, in conformity with the concept from Section 4.2.6. Provided that the approximate fluid solution is sufficiently accurate, it does not essentially hamper the convergence of the Interface-GMRES(R) method. Upon convergence, the obtained solution needs to be verified by evaluating its residual.

In Chapter 6, we presented an error-amplification analysis of the Interface-GMRES(R) method. Although the analysis is in principle generic, the numerical results are specific for the investigated model problem. Therefore, we recommend conducting corresponding numerical experiments on a multi-dimensional problem which likewise exhibits a significant change in the solution behaviour, for instance due to physical instability.

In Chapter 7, we assessed the versatility of the Interface-GMRES(R) method on the panel problem. Although our results give an indication of the performance of the Interface-GMRES(R) method, we do not suggest that they provide final and conclusive information on the method in all respects. Additional investigations are required, in particular with respect to the parameter settings for which the panel problem is stable and, if it is unstable, which type of instability it displays. Our numerical results indicated that flutter can affect the effectiveness of reuse of the Krylov space. However, to establish a relation between the convergence behaviour of the method and the stability of the problem, further research is needed. Another relevant topic for future research pertains to nonnormality. We conjecture that the subiteration method can exhibit nonnormality also for the panel problem, although we have not observed it in our numerical experiments. This conjecture remains to be confirmed, and it is important to assess which role nonnormality plays for the Interface-GMRES(R) acceleration of subiteration.

Bibliography

- J.J. Alonso and A. Jameson, Fully-implicit time-marching aeroelastic solutions, AIAA 0056 (1994), 1–13.
- [2] J.D. Anderson, Fundamentals of aerodynamics, 3rd ed., McGraw-Hill, New York, 2001.
- [3] G.K. Batchelor, An introduction to fluid dynamics, Cambridge University Press, Cambridge, 1967.
- [4] R.L. Bisplinghoff and H. Ashley, *Principles of aeroelasticity*, Dover, New York, 1975.
- [5] R.L. Bisplinghoff, H. Ashley, and R.L. Halfman, *Aeroelasticity*, Dover, New York, 1996.
- [6] F.J. Blom, A monolithical fluid-structure interaction algorithm applied to the piston problem, Comput. Methods Appl. Mech. Engrg. 167 (1998), 369–391.
- [7] K. Böhmer, P.W. Hemker, and H.J. Stetter, *The defect correction approach*, Computing 5 (1984), 1–32.
- [8] H. Braess and P. Wriggers, Arbitrary Lagrangian Eulerian finite element analysis of free surface flow, Comput. Methods Appl. Mech. Engrg. 190 (2000), 95–109.
- [9] S.C. Brenner and L.R. Scott, The mathematical theory of finite element methods, 2nd ed., Springer, New York, 2002.
- [10] P.N. Brown and Y. Saad, Hybrid Krylov methods for nonlinear systems of equations, SIAM J. Sci. Stat. Comput. 11 (1990), 450–481.
- [11] A.M. Cunningham, Buzz, buffet and LCO on military aircraft the aeroelastician's nightmares, Proceedings of the International Forum on Aeroelasticity and Structural Dynamics (IFASD) (Amsterdam, The Netherlands, 4-6 June 2003), National Aerospace Laboratory (NLR), The Netherlands, 2003.

- [12] S. Deparis, Numerical analysis of axisymmetric flows and methods for fluidstructure interaction arising in blood flow simulation, Ph.D. thesis, Ecole Polytechnique Fédérale de Lausanne, Switzerland, 2004.
- [13] J. Donea, P. Fasoli-Stella, and S. Giuliani, Lagrangian and Eulerian finite element techniques for transient fluid-structure interaction problems, Trans. 4th Int. Conf. on Structural Mechanics in Reactor Technology, San Francisco, 1977, Paper B1/2.
- [14] J. Donea, A. Huerta, J.-Ph. Ponthot, and A. Rodríguez-Ferran, Arbitrary Lagrangian-Eulerian methods, Encyclopedia of Computational Mechanics (E. Stein, R. de Borst, and T.J.R. Hughes, eds.), John Wiley & Sons, 2004.
- [15] E.H. Dowell, Aeroelasticity of plates and shells, Monographs and textbooks on mechanics of solids and fluids, Noordhoff International Publishing, Leiden, 1975.
- [16] C. Farhat, CFD-based nonlinear computational aeroelasticity, Encyclopedia of Computational Mechanics (E. Stein, R. de Borst, and T.J.R. Hughes, eds.), John Wiley & Sons, 2004.
- [17] 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 (2001), 669–694.
- [18] C. Farhat and M. Lesoinne, Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelastic problems, Comput. Methods Appl. Mech. Engrg. 182 (2000), 499–515.
- [19] C. Farhat, M. Lesoinne, and P. LeTallec, Load and motion transfer algorithms for fluid/structure interaction problems with non-matching discrete interfaces: Momentum and energy conservation, optimal discretization and application to aeroelasticity, Comput. Methods Appl. Mech. Engrg. 157 (1998), 95–114.
- [20] C.A. Felippa, K.C. Park, and C. Farhat, *Partitioned analysis of coupled mechanical systems*, Comput. Methods Appl. Mech. Engrg. **190** (2001), 3247– 3270.
- [21] Y.C. Fung, An introduction to the theory of aeroelasticity, Dover, New York, 1993.
- [22] G.H. Golub and C.F. van Loan, *Matrix computations*, Johns Hopkins, 1996.
- [23] C. Grandmont and Y. Maday, Existence for an unsteady fluid-structure interaction problem, Math. Modelling and Numerical Analysis 34 (2000), no. 3, 609–636.
- [24] A. Greenbaum, Iterative methods for solving linear systems, Frontiers in Applied Mathematics, vol. 17, SIAM, Philadelphia, 1997.

- [25] G. Hauke and T.J.R. Hughes, A comparative study of different sets of variables for solving compressible and incompressible flows, Comput. Methods Appl. Mech. Engrg. 153 (1998), 1–44.
- [26] M. Heil, An efficient solver for the fully coupled solution of large-displacement fluid-structure interaction problems, Comput. Methods Appl. Mech. Engrg. 193 (2004), 1–23.
- [27] B. Hübner, E. Walhorn, and D. Dinkler, A monolithic approach to fluidstructure interaction using space-time finite elements, Comput. Methods Appl. Mech. Engrg. 193 (2004), 2087–2104.
- [28] T.J.R. Hughes, The finite element method linear static and dynamic finite element analysis, Dover Publishers, New York, 2000.
- [29] T.J.R. Hughes, G. Engel, L. Mazzei, and M.G. Larson, *The continuous Galerkin method is locally conservative*, J. Comput. Phys. **163** (2000), 467–488.
- [30] T.J.R. Hughes and G.N. Wells, Conservation properties for the Galerkin and stabilised forms of the advection-diffusion and incompressible Navier-Stokes equations, Comput. Methods Appl. Mech. Engrg. 194 (2005), 1141–1159.
- [31] G.E. Karniadakis and S.J. Sherwin, Spectral/hp element methods for CFD, Oxford University Press, New York, 1999.
- [32] T. Kato, Perturbation theory for linear operators, corr. pr. 2nd ed., Grundlehren der Mathematischen Wissenschaft, no. 132, Springer, Berlin, 1984.
- [33] D.A. Knoll and D.E. Keyes, Jacobian-free Newton-Krylov methods: a survey of approaches and applications, J. Comput. Phys. 193 (2004), 357–397.
- [34] E. Kuhl, S. Hulshoff, and R. de Borst, An arbitrary Lagrangian Eulerian finite-element approach for fluid-structure interaction phenomena, Int. J. Num. Meth. Engrg. 57 (2003), 117–142.
- [35] M. Lesoinne and C. Farhat, Geometric conservation laws for flow problems with moving boundaries and deformable meshes, and their impact on aeroelastic computations, Comput. Methods Appl. Mech. Engrg. 134 (1996), 71–90.
- [36] P. LeTallec and J. Mouro, Fluid-structure interaction with large structural displacements, Comput. Methods Appl. Mech. Engrg. 190 (2001), 3039–3067.
- [37] A. Masud and T.J.R. Hughes, A space-time Galerkin/least-squares finite element formulation of the Navier-Stokes equations for moving domain problems, Comput. Methods Appl. Mech. Engrg. 146 (1997), 91–126.

- [38] H.G. Matthies and J. Steindorf, *Partitioned strong coupling algorithms for fluid-structure interaction*, Computers and Structures **81** (2003), 805–812.
- [39] C. Michler, S.J. Hulshoff, E.H. van Brummelen, and R. de Borst, A monolithic approach to fluid-structure interaction, Computers and Fluids 33 (2004), 839– 848.
- [40] C. Michler and E.H. van Brummelen, GMRES acceleration of subiteration for fluid-structure interactions, Moving Boundaries 2003: Proceedings of the Seventh International Conference on Computational Modelling of Free and Moving Boundary Problems (Santa Fe, New Mexico, USA, November 4-6 2003) (A.A. Mammoli and C.A. Brebbia, eds.), WIT Press, Wessex, 2003, pp. 109–118.
- [41] C. Michler, E.H. van Brummelen, S.J. Hulshoff, and R. de Borst, Efficiency of monolithic and partitioned methods for fluid-structure interaction, Proceedings of the International Forum on Aeroelasticity and Structural Dynamics (IFASD) (Amsterdam, The Netherlands, 4-6 June 2003), National Aerospace Laboratory (NLR), The Netherlands, 2003.
- [42] C. Michler, E.H. van Brummelen, and R. de Borst, Error-amplification analysis of subiteration-preconditioned GMRES for fluid-structure interaction, Comput. Methods Appl. Mech. Engrg., (Accepted for publication),(2005).
- [43] _____, An interface Newton-Krylov solver for fluid-structure interaction, Int. J. Num. Meth. Fluids 47 (2005), no. 10–11, 1189–1195, Special issue containing the Proceedings of the Eighth ICFD Conference on Numerical Methods for Fluid Dynamics (Oxford, U.K., March 29 - April 1, 2004).
- [44] C. Michler, E.H. van Brummelen, S.J. Hulshoff, and R. de Borst, The relevance of conservation for stability and accuracy of numerical methods for fluid-structure interaction, Comput. Methods Appl. Mech. Engrg. 192 (2003), 4195–4215.
- [45] D.P. Mok, Partitionierte Lösungsansätze in der Strukturdynamik und der Fluid-Struktur-Interaktion, Ph.D. thesis, Universität Stuttgart, Germany, 2001, Available at: http://elib.uni-stuttgart.de/opus/volltexte/ 2001/941/.
- [46] D.P. Mok, W.A. Wall, and E. Ramm, Accelerated iterative substructuring schemes for instationary fluid-structure interaction, Proceedings of the First MIT Conference on Computational Fluid and Solid Mechanics (K.J. Bathe, ed.), Elsevier, June 12–15 2001.
- [47] S.A. Morton, R.B. Melville, and M.R. Visbal, Accuracy and coupling issues of aeroelastic Navier-Stokes solutions on deforming meshes, AIAA 1085 (1997), 252–262.

- [48] F. Nobile, Numerical approximation of fluid-structure interaction problems with application to haemodynamics, Ph.D. thesis, Ecole Polytechnique Fédérale de Lausanne, Switzerland, 2001.
- [49] W.F. Noh, CEL: A time-dependent two-space-dimensional coupled Eulerian-Lagrangian code, Methods in Comput. Physics (B. Alder, S. Fernbach, and M. Rotenberg, eds.), vol. 3, Academic Press, New York, 1964, pp. 117–179.
- [50] R.W. Ogden, Non-linear elastic deformations, Ellis Horwood Limited, Chichester, 1984.
- [51] R. Ohayon and C. Felippa (eds.), Special issue on fluid-structure interaction, Comput. Methods Appl. Mech. Engrg. 190 (2001), no. 24-25, 2977–3292.
- [52] S. Piperno, Simulation numérique de phénomènes d'interaction fluidestructure, Ph.D. thesis, Ecole Nationale des Ponts et Chaussées, France, 1995, (partly in French).
- [53] _____, Explicit/implicit fluid/structure staggered procedures with a structural predictor and fluid subcycling for 2D inviscid aeroelastic simulations, Int. J. Num. Meth. Fluids 25 (1997), 1207–1226.
- [54] S. Piperno and C. Farhat, Design and evaluation of staggered partitioned procedures for fluid-structure interaction simulations, Tech. Report 97-107, CERMICS-INRIA, 1997.
- [55] _____, Partitioned procedures for the transient solution of coupled aeroelastic problems - part II: Energy transfer analysis and three-dimensional applications, Comput. Methods Appl. Mech. Engrg. 190 (2001), 3147–3170.
- [56] S. Piperno, C. Farhat, and B. Larrouturou, Partitioned procedures for the transient solution of coupled aeroelastic problems - part I: Model problem, theory and two-dimensional applications, Comput. Methods Appl. Mech. Engrg. 124 (1995), 79–112.
- [57] O. Pironneau, Optimal shape design for elliptic systems, Computational Physics, Springer, Berlin, 1984.
- [58] S.M. Rifai, Z. Johan, W.-P. Wang, J.-P. Grisval, T.J.R. Hughes, and R.M. Ferencz, Multiphysics simulation of flow-induced vibrations and aeroelasticity on parallel computing platforms, Comput. Methods Appl. Mech. Engrg. 174 (1999), 393–417.
- [59] Y. Saad and M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Comp. 7 (1986), 856– 869.
- [60] J. Salençon, Handbook of continuum mechanics, Springer, Berlin, 2001.

- [61] C. Schwab, p- and hp- finite element methods, Oxford University Press Inc., New York, 1998.
- [62] F. Shakib, T.J.R. Hughes, and Z. Johan, A new finite element formulation for computational fluid dynamics: X. The compressible Euler and Navier-Stokes equations, Comput. Methods Appl. Mech. Engrg. 89 (1991), 141–219.
- [63] J.M.A. Stijnen, J. de Hart, P.H.M. Bovendeerd, and F.N. van de Vosse, Evaluation of a fictitious domain method for predicting dynamic response of mechanical heart valves, J. of Fluids and Structures 19 (2004), 835–850.
- [64] R. Témam and Alain Miranville, Mathematical modeling in continuum mechanics, Cambridge University Press, Cambridge, U.K., 2001.
- [65] T.E. Tezduyar, M. Behr, and J. Liou, A new strategy for finite element computations involving moving boundaries and interfaces – the DSD/ST procedure: I. the concept and the preliminary numerical tests, Comput. Methods Appl. Mech. Engrg. 94 (1992), 339–351.
- [66] T.E. Tezduyar, M. Behr, S. Mittal, and J. Liou, A new strategy for finite element computations involving moving boundaries and interfaces – the DSD/ST procedure: II. computation of free-surface flows, two-liquid flows, and flows with drifting cylinders, Comput. Methods Appl. Mech. Engrg. 94 (1992), 353– 371.
- [67] T.E. Tezduyar and Y. Osawa, Fluid-structure interactions of a parachute crossing the far wake of an aircraft, Comput. Methods Appl. Mech. Engrg. 191 (2001), 717–726.
- [68] L.N. Trefethen, Pseudospectra of linear operators, SIAM Rev. 39 (1997), no. 3, 383–406.
- [69] E.H. van Brummelen and R. de Borst, On the nonnormality of subiteration for a fluid-structure-interaction problem, SIAM J. Sci. Comput., (Accepted for publication), (2004). Preprint available at : http://www.em.lr.tudelft. nl/~brummelen/publications.html.
- [70] E.H. van Brummelen, S.J. Hulshoff, and R. de Borst, *Energy conservation under incompatibility for fluid-structure interaction problems*, Comput. Methods Appl. Mech. Engrg. **192** (2003), 2727–2748.
- [71] E.H. van Brummelen, C. Michler, and R. de Borst, Interface-GMRES(R) acceleration of subiteration for fluid-structure-interaction problems, (Submitted for publication), (2005). Also available as: Tech. Report DACS-05-001, Delft University of Technology, The Netherlands, 2005. Available at: http://www.em.lr.tudelft.nl/downloads/DACS-05-001.pdf.

- [72] M.E. Verbeek, Repairing near-singularity for dense EMC problems by adaptive basis techniques, Num. Lin. Alg. Appl. 7 (2000), 617–634.
- [73] W.A. Wall, Fluid-Struktur-Interaktion mit stabilisierten Finiten Elementen, Ph.D. thesis, Universität Stuttgart, Germany, 1999, Available at: http:// elib.uni-stuttgart.de/opus/volltexte/2000/623/.
- [74] T. Washio and C.W. Oosterlee, Krylov subspace acceleration for nonlinear multigrid schemes, Electronic Transactions on Numerical Analysis 6 (1997), 271–290.
- [75] T.G. Wright, *Matlab pseudospectra GUI*, Available at : http://web.comlab. ox.ac.uk/projects/pseudospectra/eigtool/.

Appendix A Conservation properties of trapezoidal time integration

A.1 Assessment of energy conservation

Considering the trapezoidal discretization of a structure which is subject to external forces, the objective is to determine the change in energy between time levels that is implied by the trapezoidal discretization.

Provided with displacement d, velocity v and acceleration a at time level t_n , the trapezoidal method determines the corresponding values at time level t_{n+1} from

$$d_{n+1} = d_n + \Delta t \frac{v_{n+1} + v_n}{2},$$
 (A.1a)

$$v_{n+1} = v_n + \Delta t \frac{a_{n+1} + a_n}{2},$$
 (A.1b)

$$Ma_{n+1} + Kd_{n+1} = \theta_{n+1}$$
 (A.1c)

for the harmonic oscillator given by Eq. (3.3). Upon replacing the average acceleration in Eq. (A.1b) by Eq. (A.1c) and a corresponding expression for a_n , Eq. (A.1b) can be rewritten as

$$v_{n+1} = v_n - \frac{\Delta t K}{2M} (d_{n+1} + d_n) + \frac{\Delta t}{2M} (\theta_{n+1} + \theta_n).$$
 (A.2)

To assess the change in energy between time levels, it is convenient to rewrite Eqs. (A.1a) and (A.2) concisely in the form

$$Aq_{n+1} = Bq_n + p \tag{A.3a}$$

with

$$A := \begin{pmatrix} 1 & \frac{\Delta tK}{2M} \\ -\frac{\Delta t}{2} & 1 \end{pmatrix}, \qquad B := \begin{pmatrix} 1 & -\frac{\Delta tK}{2M} \\ \frac{\Delta t}{2} & 1 \end{pmatrix},$$
$$q := \begin{bmatrix} v \\ d \end{bmatrix}, \qquad p := \begin{bmatrix} \frac{\theta_{n+1}+\theta_n}{2}\frac{\Delta t}{M} \\ 0 \end{bmatrix}. \quad (A.3b)$$

The energy of the structure at time level t_{n+1} is given by

$$E_{n+1} := q_{n+1}^T \Lambda q_{n+1} \tag{A.4}$$

with $\Lambda := \text{diag}(M/2, K/2)$. Upon substituting Eq. (A.3) into (A.4) and denoting $A^{-1}B := C$, we obtain

$$E_{n+1} = q_n^T C^T \Lambda C q_n + q_n^T C^T \Lambda A^{-1} p + p^T A^{-T} \Lambda q_{n+1}.$$
(A.5)

Since $C^T \Lambda C = \Lambda$, the first term in the right-hand-side of Eq. (A.5) simplifies to $q_n^T \Lambda q_n =: E_n$. By manipulating the expressions of the second and third term and invoking Eq. (A.1a), Eq. (A.5) simplifies to

$$E_{n+1} = E_n + \Delta t \left(\frac{v_{n+1} + v_n}{2}\right) \left(\frac{\theta_{n+1} + \theta_n}{2}\right).$$
(A.6)

The second term on the right-hand side in Eq. (A.6) expresses the change in energy between time levels that is induced by the external forces. Based on Eq. (A.6), we conclude in Section 3.4.5 that the trapezoidal method does generally not conserve energy for forced structures.

A.2 Assessment of momentum conservation

To assess momentum conservation in the trapezoidal method, we determine the change in momentum in the generic time interval $[t_n, t_{n+1}]$ for the harmonic oscillator given by Eq. (3.3). Integration of Eq. (3.3) over time yields

$$Mv\Big|_{t_n}^{t_{n+1}} = -\int_{t_n}^{t_{n+1}} Kd \, \mathrm{d}t + \int_{t_n}^{t_{n+1}} \theta \, \mathrm{d}t \,. \tag{A.7}$$

The change in momentum that is implied by the trapezoidal discretization can be derived from Eq. (A.3). Upon multiplication with A^{-1} and some straightforward manipulations, we obtain for the change in velocity

$$v_{n+1} - v_n = \frac{1}{M + \frac{K}{4}\Delta t^2} \left(-\Delta t K d_n - \frac{1}{2}\Delta t^2 K v_n + \Delta t \frac{\theta_{n+1} + \theta_n}{2} \right).$$
(A.8)

Let us consider Eq. (A.8) for $\Delta t \to 0$ and invoke a Taylor series expansion for the first factor in the right member. Subsequent multiplication by M then yields

$$M(v_{n+1} - v_n) = (1 + O(\Delta t^2)) \left(-\Delta t K d_n - \frac{1}{2} \Delta t^2 K v_n + \Delta t \frac{\theta_{n+1} + \theta_n}{2} \right).$$
(A.9)

A comparison of Eqs. (A.7) and (A.9) conveys the approximation to the integrals that is implied by the trapezoidal method. In particular, the integral pertaining to the applied force is approximated by the midpoint rule modulo a higher-order term. We infer that the trapezoidal method does not generally conserve momentum for forced structures.

Subject Index

aggregated variational problem, 22 approximation space, 23 arbitrary Lagrangian-Eulerian formulation, 13 augmented problem, 37 auxiliary conditions strong enforcement of —, 18 weak enforcement of —, 18

balance of linear momentum, 19 boundary condition farfield —, 121 no-slip —, 17 slip —, 121

coincidence condition, 37 compatibility condition, 36 conservation law, 16 consistency, 24 convergence behaviour asymptotic —, 96 transient —, 96 convergence bounds GMRES —, 107 subiteration —, 97

defect correction method, 56 Discrete Geometric Conservation Law, 24discretization error, 57

equation of state, 16 error discretization —, 42

linearization —, 73

numerical evaluation —, 57 error amplification — of Interface-GMRES(R), 105 — of subiteration, 94 Euler equations, 16, 29, 109, 120 Eulerian description, 10 existence, 23 finite-element discretization, 23

flutter, 1, 123 free-boundary problem, 3, 26, 30 Fréchet derivative, 53

Gauss-Seidel method, 56 Gram-Schmidt orthonormalization, 73

harmonic oscillator, 29 hemodynamics, 58 homeomorphic transformation, 18

impermeability, 17, 30, 121
incompatible discretization, 37
inner-iteration, 56
instability

numerical —, 57
physical —, 123

interface condition

dynamic —, 22, 30, 121
kinematic —, 21, 30, 121

Interface-GMRES, 72

algorithm, 73, 74
Interface-GMRESR, 75
interfield iteration, 59

Krylov space, 72, 97

Lagrangian description, 12 least-squares problem, 73, 102 mass ratio, 33, 122 material description, see Lagrangian description method of spines, 54 monolithic methods, 53 moving domain, 10, 15 multi-physics, 2 Navier-Stokes equations, 16 nested preconditioning, 105 Newton's method, 53, 72, 91 Newton-Krylov method, 71 non-dimensionalization, 33, 122 nonnormality, 96, 107 panel problem, 119, 120 partitioned methods loosely-coupled -, 56, 58, 60, 127strongly-coupled —, 55, 58, 60, 127partitioning, 51 Picard iteration, 56 piston problem, 28, 32, 119 prediction techniques, 57 pseudospectrum, 115 residual estimate, 73 residual-sensitivity space, 72, 101 reuse of Krylov vectors, 75, 104, 107 Schur complement, 99, 100, 105 search space, 73, 101 search space injection, 104 security set, 18 shape derivative, 54, 92 software modularity, 51 space/time - domain, 15, 24 finite-element discretization, 23, 24- slab, 23

— variational formulation, 15, 17 spatial description, see Eulerian description spectral radius, 96, 107 state variables, 15 steepest-ascent method, 112 subcycling, 42 subiteration, 55, 71, 93 - algorithm, 55, 93 — operator, 71, 94, 96 — preconditioning, 97 substructuring schemes, 75 successive approximation, 56 time-discontinuous Galerkin method, 23, 40, 41 trapezoidal method, 38, 149 unconditional stability, 59 underrelaxation, 73 van-der-Pol oscillator, 109

Summary

Efficient Numerical Methods for Fluid-Structure Interaction

The interaction between a fluid and a structure occurs in many physical systems and engineering applications. The accurate prediction of the behaviour of such fluid-structure interactions is therefore of great importance, for instance, in the design of aircraft, bridges and artificial heart valves. Numerical methods play an important role in the solution of fluid-structure-interaction problems. The computation of fluid-structure interactions is challenging on account of their multi-physics and free-boundary character. In addition, maintaining the conservation properties at the fluid-structure interface under discretization is in general non-trivial. Moreover, the customary solution methods often lack robustness and efficiency. This motivates research into conservative discretization techniques and robust and efficient iterative solution methods.

We present a formulation of fluid-structure-interaction problems in classical and in generic variational form. The generic form serves to identify the characteristic features of fluid-structure interaction and, moreover, to formulate solution algorithms such that they are essentially applicable to *any* fluid-structure-interaction problem.

On the basis of a prototypical model problem, we investigate the conservation properties of the discrete numerical model at the fluid-structure interface. Energy conservation at the interface is only trivially maintained under restrictive compatibility conditions on the approximation spaces for the fluid and the structure, i.e., matching meshes and identical orders of approximation at the interface. These conditions are prohibitive for practical use. We then consider an approach based on coincidence and orthogonality conditions which enables conservation also for incompatible discretizations. Our numerical results demonstrate that a method which maintains conservation at the interface yields a much more accurate solution than a non-conservative method, at the same computational expense. The results also illustrate that violation of the conservation properties can induce numerical instability.

Next, we compare loosely-coupled and strongly-coupled solution methods in terms of stability, accuracy and efficiency. Our numerical results demonstrate that the numerical evaluation error incurred by loose coupling can compromise the aforementioned properties. The results moreover convey that strongly-coupled methods are superior to loosely-coupled methods provided that the underlying discretization maintains the conservation properties. In particular, our results indicate that the higher computational cost of strongly-coupled methods is only justified by a greater accuracy if the underlying discretization is conservative. To overcome the deficiencies of the customary subiteration solution method, viz., only conditional stability, potential convergence difficulties due to nonnormality and the inability to reuse information from previously solved similar problems, we propose to combine subiteration with GMRES acceleration. Because the Krylov vectors need only contain the discrete representation associated with the interface position, the corresponding *Interface-GMRES acceleration* requires only negligible computational resources. Moreover, the method allows for optional reuse of Krylov vectors in subsequent invocations of GMRES, which we refer to as GMRESR. Such reuse can considerably enhance the efficiency of the method. The Interface-GMRES(R) method retains the modularity of the underlying subiteration method. We illustrate the effectiveness of the proposed method through numerical experiments on the prototypical piston model problem. Detailed convergence studies and a comparison to standard subiteration show that the Interface-GMRES(R) method is much more robust and efficient, and that it converges even if subiteration itself diverges.

We present an error-amplification analysis of subiteration and of the Interface-GMRES(R) method. The adopted linear-algebra setting enables a clear explanation of the relation between the *local* GMRES acceleration (i.e., on the interface degrees-of-freedom), and the *global* error-amplification properties (i.e., for the aggregated fluid-structure system). We show that subiteration condenses errors into a low-dimensional subspace associated with the interface degrees-of-freedom. The rank of the corresponding error-amplification matrix is at most equal to the dimension N of the approximation space of the interface variables. This implies that any Krylov method converges in at most N steps, independent of the choice of the acceleration space, e.g., aggregated variables, structure variables, or interface variables. However, the acceleration on the interface variables is the most efficient, because the computational cost and the storage required by the Krylov acceleration itself increase with the dimension of the acceleration space. Numerical experiments are conducted to investigate the effect of the reuse of Krylov vectors on the error-amplification properties of the Interface-GMRESR method and to assess the limitations on the reuse option.

To demonstrate the versatility of the Interface-GMRES(R) method, we assess the solution method on an alternative model problem, viz., the panel problem. The panel problem is distinctly different from the piston problem. In contrast to the one-dimensional piston problem, for the two-dimensional panel problem the interface degrees-of-freedom pertain to both space and time. Furthermore, the panel problem can exhibit parameter-dependent stability behaviour. In particular, it admits instabilities such as flutter and divergence. Our numerical results confirm the findings on the piston problem. In addition, the results indicate that flutter instability can affect the effectiveness of the reuse of Krylov vectors. However, this effect appears to be minor and reuse remains beneficial.

Samenvatting

Efficiënte Numerieke Methoden voor Vloeistof-Constructie Interactie

De interactie tussen een vloeistof en een constructie komt voor in vele fysische systemen en praktische toepassingen. De nauwkeurige voorspelling van het gedrag van dergelijke vloeistof-constructie interactie is dus van groot belang, bijvoorbeeld in het ontwerp van vliegtuigen, bruggen en kunstmatige hartkleppen. Numerieke methoden spelen een belangrijke rol in het oplossen van vloeistof-constructie interactie problemen. De berekening van de interactie tussen een vloeistof en een constructie is uitdagend op grond van het multifysica en vrije-rand karakter. Daarnaast is het bewaren van de behoudseigenschappen aan de rand tussen de vloeistof en de constructie (de 'interface') onder discretisatie in het algemeen niet triviaal. Bovendien missen de conventionele oplosmethoden voor dit type problemen vaak robuustheid en efficiëntie. Dit is de motivatie voor onderzoek naar behoudende discretisatie-technieken en robuste en efficiënte iteratieve oplosmethoden.

We presenteren een formulering van vloeistof-constructie interactie problemen in een klassieke en in een generieke variationele vorm. De generieke vorm dient om de karakteristieke eigenschappen van vloeistof-constructie interactie te identificeren en, daarnaast, om oplosalgorithmen te formuleren zodat deze in principe toepasbaar zijn voor elk vloeistof-constructie interactie probleem.

Door middel van een prototypisch model-probleem onderzoeken we de behoudseigenschappen van het discrete numerieke model aan de interface. Aan energiebehoud aan de interface wordt slechts triviaal voldaan onder zeer beperkende compatibiliteitsvoorwaarden met betrekking tot de benaderingsruimtes voor de vloeistof en de constructie, d.w.z., aansluitende roosters en gelijke orde van benadering aan de interface. Deze voorwaarden zijn te restrictief om het in de praktijk toe te passen. Daarnaast beschouwen we een methodiek die op 'coincidence'en orthogonaliteitsvoorwaarden gebaseerd is. Deze methodiek maakt het mogelijk om ook voor incompatibele discretisaties aan de behoudseigenschappen te voldoen. Onze numerieke resultaten laten zien dat een methode die aan de behoudseigenschappen aan de interface voldoet een veel nauwkeurigere oplossing oplevert dan een methode die niet behoudend is, bij gelijke rekenkosten. De resultaten illustreren ook dat het niet voldoen aan de behoudseigenschappen kan leiden tot numerieke instabiliteit.

Vervolgens vergelijken we zwak-gekoppelde en sterk-gekoppelde oplosmethoden qua stabiliteit, nauwkeurigheid en efficiëntie. Onze numerieke resultaten laten zien dat de numerieke fout die door zwak-gekoppelde methoden geïntroduceerd wordt de stabiliteit, nauwkeurigheid en efficiëntie in gevaar kan brengen. De resultaten tonen daarnaast aan dat sterk-gekoppelde methoden superieur zijn ten opzichte van zwak-gekoppelde methoden, onder de voorwaarde dat de onderliggende discretisatie aan de behoudseigenschappen voldoet. Onze resultaten tonen aan dat de hogere rekenkosten van sterk-gekoppelde methoden alleen maar gerechtvaardigd zijn door een hogere nauwkeurigheid als de onderliggende discretisatie behoudend is.

Om de nadelen van de gebruikelijke subiteratie oplosmethode, namelijk slechts voorwaardelijke stabiliteit, potentiële convergentie moeilijkheden vanwege nietnormaliteit en de onbekwaamheid om informatie van reeds opgeloste vergelijkbare problemen te hergebruiken, te overwinnen stellen we voor subiteratie met GMRES versnelling te combineren. Omdat de Krylov vektoren slechts de diskrete representatie van de interface positie hoeven te bevatten, heeft de bijbehorende Interface-GMRES versnelling minimale computer capaciteit nodig. Bovendien staat de methode het optionele hergebruik ('reuse') van Krylov vektoren in opvolgende toepassingen van GMRES toe, hetgeen we GMRESR noemen. Dit hergebruik kan de efficiëntie van de methode aanzienlijk verbeteren. De Interface-GMRES(R) methode behoudt de modulariteit van de onderliggende subiteratie methode. We illustreren de effectiviteit van de voorgestelde methode aan de hand van numerieke experimenten met het prototypisch 'piston model-probleem'. Gedetailleerde convergentie studies en een vergelijking met de standaard subiteratie methode tonen aan dat de Interface-GMRES(R) methode veel robuuster en efficiënter is, en dat de methode convergeert zelfs als subiteratie divergeert.

We presenteren een analyse van de fout-versterking van subiteratie en Interface-GMRES(R). De aangenomen lineaire algebra opzet maakt een duidelijke verklaring mogelijk van de relatie tussen de *locale* GMRES versnelling (d.w.z., op de interface vrijheidsgraden) en de *qlobale* fout-versterkingseigenschappen (d.w.z., voor het totale systeem van vloeistof- en constructie vergelijkingen). We laten zien dat subiteratie fouten in een laag-dimensionale ruimte projecteert die met de interface vrijheidsgraden geassocieerd kan worden. De rang van de bijbehorende foutversterkingsmatrix is hoogstens gelijk aan de dimensie N van de benaderingsruimte van de interface variabelen. Dit impliceert dat iedere Krylov methode in hooguit Nstappen convergeert, onafhankelijk van de keuze van de versnellingsruimte, bijvoorbeeld, alle variabelen van het systeem, de constructie variabelen, of de interface variabelen. Echter, de versnelling op de interface variabelen is het meest efficiënt omdat de voor de Krylov versnelling benodigde berekeningskosten en geheugen toenemen met de dimensie van de versnellingsruimte. Numerieke experimenten zijn uitgevoerd om het effect van het hergebruik van Krylov vektoren op de foutversterkingseigenschappen van de Interface-GMRESR methode te onderzoeken en om de grenzen van hergebruik vast te stellen.

Om de ruime toepasbaarheid van de Interface-GMRES(R) methode aan te tonen, testen we de oplosmethode op een ander model probleem, namelijk het 'panel' probleem. Het panel probleem is wezenlijk verschillend van het piston probleem. In tegenstelling tot het een-dimensionale piston probleem, hebben de interface vrijheidsgraden voor het twee-dimensionale panel probleem betrekking op de ruimte en de tijd. Bovendien kan het panel probleem parameter-afhankelijke stabiliteits-eigenschappen vertonen. Het systeem laat bijvoorbeeld instabiliteiten zoals 'flutter' en 'divergence' toe. Onze numerieke resultaten bevestigen de bevindingen met betrekking tot het piston probleem. Daarnaast tonen de resultaten aan dat flutter instabiliteit de effectiviteit van het hergebruik van Krylov vektoren kan beïnvloeden. Echter, dit effect blijkt van weinig belang te zijn en hergebruik blijft voordelig.

Curriculum Vitae

May 22, 1976	Born in Celle, Germany
1988 - 1995	Scharnhorst-Gymnasium, Hildesheim
Sept. 1995 - Sept. 1996	Civil service at the Landesbildungszentrum für Hörgeschädigte, Hildesheim
Oct. 1996 - Aug. 1999	Chemical Engineering studies at the Technical University of Clausthal, Germany
Aug. 1999	Diploma (M.Sc. degree) with distinction 'cum laude'
Oct. 1997 - Aug. 1999	Part-time employment at IDV-Systeme, Clausthal (engineering office for finite-element analysis of structures)
Sept. 1999 - June 2000	Post-graduate Diploma Course at the von Kármán Institute for Fluid Dynamics (Department of Aeronautics / Aerospace), Rhode-St-Genèse, Belgium
Since Oct. 2000	Research assistant at the Faculty of Aerospace Engineering, Delft University of Technology, The Netherlands