Finite Element Exterior Calculus

DOUGLAS N. ARNOLD University of Minnesota Minneapolis, Minnesota

CBMS-NSF REGIONAL CONFERENCE SERIES IN APPLIED MATHEMATICS

SPONSORED BY CONFERENCE BOARD OF THE MATHEMATICAL SCIENCES SUPPORTED BY NATIONAL SCIENCE FOUNDATION

Finite Element Exterior Calculus

CBMS-NSF REGIONAL CONFERENCE SERIES IN APPLIED MATHEMATICS

A series of lectures on topics of current research interest in applied mathematics under the direction of the Conference Board of the Mathematical Sciences, supported by the National Science Foundation and published by SIAM.

GARRETT BIRKHOFF, The Numerical Solution of Elliptic Equations D. V. LINDLEY, Bayesian Statistics, A Review R. S. VARGA, Functional Analysis and Approximation Theory in Numerical Analysis R. R. BAHADUR. Some Limit Theorems in Statistics PATRICK BILLINGSLEY, Weak Convergence of Measures: Applications in Probability J. L. LIONS, Some Aspects of the Optimal Control of Distributed Parameter Systems ROGER PENROSE, Techniques of Differential Topology in Relativity HERMAN CHERNOFF, Sequential Analysis and Optimal Design J. DURBIN, Distribution Theory for Tests Based on the Sample Distribution Function SOL I. RUBINOW, Mathematical Problems in the Biological Sciences P. D. LAX, Hyperbolic Systems of Conservation Laws and the Mathematical Theory of Shock Waves I. J. SCHOENBERG, Cardinal Spline Interpolation IVAN SINGER, The Theory of Best Approximation and Functional Analysis WERNER C. RHEINBOLDT, Methods of Solving Systems of Nonlinear Equations HANS F. WEINBERGER, Variational Methods for Eigenvalue Approximation R. TYRRELL ROCKAFELLAR, Conjugate Duality and Optimization SIR JAMES LIGHTHILL, Mathematical Biofluiddynamics GERARD SALTON, Theory of Indexing CATHLEEN S. MORAWETZ, Notes on Time Decay and Scattering for Some Hyperbolic Problems F. HOPPENSTEADT, Mathematical Theories of Populations: Demographics, Genetics and Epidemics RICHARD ASKEY, Orthogonal Polynomials and Special Functions L. E. PAYNE, Improperly Posed Problems in Partial Differential Equations S. ROSEN, Lectures on the Measurement and Evaluation of the Performance of Computing Systems HERBERT B. KELLER, Numerical Solution of Two Point Boundary Value Problems J. P. LASALLE, The Stability of Dynamical Systems D. GOTTLIEB AND S. A. ORSZAG, Numerical Analysis of Spectral Methods: Theory and Applications PETER J. HUBER, Robust Statistical Procedures HERBERT SOLOMON, Geometric Probability FRED S. ROBERTS, Graph Theory and Its Applications to Problems of Society JURIS HARTMANIS, Feasible Computations and Provable Complexity Properties ZOHAR MANNA, Lectures on the Logic of Computer Programming ELLIS L. JOHNSON, Integer Programming: Facets, Subadditivity, and Duality for Group and Semi-group Problems SHMUEL WINOGRAD, Arithmetic Complexity of Computations J. F. C. KINGMAN, Mathematics of Genetic Diversity MORTON E. GURTIN, Topics in Finite Elasticity THOMAS G. KURTZ, Approximation of Population Processes JERROLD E. MARSDEN, Lectures on Geometric Methods in Mathematical Physics BRADLEY EFRON, The Jackknife, the Bootstrap, and Other Resampling Plans M. WOODROOFE, Nonlinear Renewal Theory in Sequential Analysis D. H. SATTINGER, Branching in the Presence of Symmetry R. TEMAM, Navier-Stokes Equations and Nonlinear Functional Analysis MIKLÓS CSÖRG, Quantile Processes with Statistical Applications J. D. BUCKMASTER AND G. S. S. LUDFORD, Lectures on Mathematical Combustion R. E. TARJAN, Data Structures and Network Algorithms

PAUL WALTMAN, Competition Models in Population Biology

S. R. S. VARADHAN, Large Deviations and Applications KIYOSI ITÔ, Foundations of Stochastic Differential Equations in Infinite Dimensional Spaces ALAN C. NEWELL, Solitons in Mathematics and Physics PRANAB KUMAR SEN, Theory and Applications of Sequential Nonparametrics László Lovász, An Algorithmic Theory of Numbers, Graphs and Convexity E. W. CHENEY, Multivariate Approximation Theory: Selected Topics JOEL SPENCER, Ten Lectures on the Probabilistic Method PAUL C. FIFE, Dynamics of Internal Layers and Diffusive Interfaces CHARLES K. CHUI, Multivariate Splines HERBERT S. WILF, Combinatorial Algorithms: An Update HENRY C. TUCKWELL, Stochastic Processes in the Neurosciences FRANK H. CLARKE, Methods of Dynamic and Nonsmooth Optimization **ROBERT B. GARDNER**, The Method of Equivalence and Its Applications GRACE WAHBA, Spline Models for Observational Data RICHARD S. VARGA, Scientific Computation on Mathematical Problems and Conjectures INGRID DAUBECHIES, Ten Lectures on Wavelets STEPHEN F. McCormick, Multilevel Projection Methods for Partial Differential Equations HARALD NIEDERREITER, Random Number Generation and Quasi-Monte Carlo Methods JOEL SPENCER, Ten Lectures on the Probabilistic Method, Second Edition CHARLES A. MICCHELLI, Mathematical Aspects of Geometric Modeling ROGER TEMAM, Navier-Stokes Equations and Nonlinear Functional Analysis, Second Edition GLENN SHAFER, Probabilistic Expert Systems PETER J. HUBER, Robust Statistical Procedures, Second Edition J. MICHAEL STEELE, Probability Theory and Combinatorial Optimization WERNER C. RHEINBOLDT, Methods for Solving Systems of Nonlinear Equations, Second Edition J. M. CUSHING, An Introduction to Structured Population Dynamics TAI-PING LIU, Hyperbolic and Viscous Conservation Laws MICHAEL RENARDY, Mathematical Analysis of Viscoelastic Flows GÉRARD CORNUÉJOLS, Combinatorial Optimization: Packing and Covering IRENA LASIECKA, Mathematical Control Theory of Coupled PDEs J. K. SHAW, Mathematical Principles of Optical Fiber Communications ZHANGXIN CHEN, Reservoir Simulation: Mathematical Techniques in Oil Recovery ATHANASSIOS S. FOKAS, A Unified Approach to Boundary Value Problems MARGARET CHENEY AND BRETT BORDEN, Fundamentals of Radar Imaging FIORALBA CAKONI, DAVID COLTON, AND PETER MONK, The Linear Sampling Method in Inverse Electromagnetic Scattering ADRIAN CONSTANTIN, Nonlinear Water Waves with Applications to Wave-Current Interactions and Tsunamis WEI-MING NI, The Mathematics of Diffusion ARNULF JENTZEN AND PETER E. KLOEDEN, Taylor Approximations for Stochastic Partial Differential Equations FRED BRAUER AND CARLOS CASTILLO-CHAVEZ, Mathematical Models for Communicable Diseases PETER KUCHMENT, The Radon Transform and Medical Imaging ROLAND GLOWINSKI, Variational Methods for the Numerical Solution of Nonlinear Elliptic Problems BENGT FORNBERG AND NATASHA FLYER, A Primer on Radial Basis Functions with Applications to the Geosciences FIORALBA CAKONI, DAVID COLTON, AND HOUSSEM HADDAR, Inverse Scattering Theory and Transmission Eigenvalues MIKE STEEL, Phylogeny: Discrete and Random Processes in Evolution PETER CONSTANTIN, Analysis of Hydrodynamic Models DONALD G. SAARI, Mathematics Motivated by the Social and Behavioral Sciences YUJI KODAMA. Solitons in Two-Dimensional Shallow Water DOUGLAS N. ARNOLD, Finite Element Exterior Calculus

Douglas N. Arnold

University of Minnesota Minneapolis, Minnesota

Finite Element Exterior Calculus



Copyright © 2018 by the Society for Industrial and Applied Mathematics

10987654321

All rights reserved. Printed in the United States of America. No part of this book may be reproduced, stored, or transmitted in any manner without the written permission of the publisher. For information, write to the Society for Industrial and Applied Mathematics, 3600 Market Street, 6th Floor, Philadelphia, PA 19104-2688 USA.

Trademarked names may be used in this book without the inclusion of a trademark symbol. These names are used in an editorial context only; no infringement of trademark is intended.

Publications Director	Kivmars H. Bowling
Acquisitions Editor	Paula Callaghan
Developmental Editor	Gina Rinelli Harris
Managing Editor	Kelly Thomas
Production Editor	Lisa Briggeman
Copy Editor	Susan Fleshman
Production Manager	Donna Witzleben
Production Coordinator	Cally A. Shrader
Compositor	Cheryl Hufnagle
Graphic Designer	Doug Smock

Library of Congress Cataloging-in-Publication Data

Names: Arnold, Douglas N., 1954- author.
Title: Finite element exterior calculus / Douglas N. Arnold (University of Minnesota, Minneapolis, Minnesota).
Description: Philadelphia : Society for Industrial and Applied Mathematics, [2018] | Series: CBMS-NSF regional conference series in applied mathematics ; 93 | Includes bibliographical references and index.
Identifiers: LCCN 2018045527 (print) | LCCN 2018047418 (ebook) | ISBN 9781611975543 | ISBN 9781611975536
Subjects: LCSH: Differential equations--Numerical solutions. | Finite element method. | Calculus. | Calculus of variations.
Classification: LCC QA372 (ebook) | LCC QA372 .A6845 2018 (print) | DDC 515/.35--dc23

LC record available at https://lccn.loc.gov/2018045527



Contents

Pre	face		ix
1	Intro	oduction	1
2	Basic	e notions of homological algebra	11
	2.1	Graded vector spaces, chain complexes, and differentials	11
	2.2	Cycles, boundaries, and homology	12
	2.3	Example: The simplicial chain complex	12
	2.4	Chain maps	15
	2.5	Cochain complexes	16
	2.6	Example: The de Rham complex	16
3	Basic	e notions of unbounded operators on Hilbert spaces	21
	3.1	Unbounded operators	21
	3.2	Closed operators	22
	3.3	The adjoint of an unbounded operator	23
	3.4	Example: grad, curl, and div and their adjoints	27
4	Hilbe	ert complexes	33
	4.1	Hilbert complexes and their duals	33
	4.2	Harmonic forms, the Hodge decomposition, and the Poincaré	
		inequality	35
	4.3	Example: The L^2 de Rham complex $\ldots \ldots \ldots \ldots \ldots$	37
	4.4	The abstract Hodge Laplacian	39
	4.5	Example: The Hodge Laplacian in \mathbb{R}^3	44
5	Appr	oximation of Hilbert complexes	51
	5.1	Galerkin discretization of the primal formulation	51
	5.2	Mixed methods of discretization	54
6	Basic	e notions of exterior calculus	63
	6.1	Exterior algebra	63
	6.2	Exterior calculus	68
7	Finit	e element differential forms	81
	7.1	The complete polynomial spaces of differential forms	82
	7.2	The Koszul complex	83
	7.3	The trimmed polynomial spaces of differential forms	85
	7.4	Finite element differential forms	86

	7.5	Properties of the finite element spaces	88
	7.6	The Whitney forms $\mathcal{P}_1^- \Lambda^k$ and de Rham's theorem	90
	7.7	Cubical elements	91
	7.8	Historical antecedents	93
8	Furthe	r directions and applications	97
	8.1	Eigenvalue problems	97
	8.2	Variable coefficients	98
	8.3	Lower order terms	99
	8.4	Differential equations on manifolds	101
	8.5	Parabolic and hyperbolic problems	101
	8.6	Maxwell's equations	103
	8.7	The Stokes equations	104
	8.8	Elasticity	107
Bibl	liography	7	113
Inde	ex		119

Preface

Computational methods to approximate the solution of differential equations play a crucial role in science, engineering, mathematics, and technology. Indeed, the key processes which govern the physical world-wave propagation, thermodynamics, fluid flow, solid deformation, gas dynamics, electricity and magnetism, quantum mechanics, general relativity, and many more-are described by differential equations, and we depend on numerical methods for the ability to simulate, explore, predict, and control systems involving these processes. The variety of differential equation problems that arise in these applications is vast, and much research has gone into developing numerical methods which can solve different problems accurately and efficiently. Mathematical analysis of these algorithms plays an essential role, furnishing rigorous validation to particular methods in clearly delineated circumstances, supplying quantitative error bounds, and enabling comparison among competing methods. In this book we will focus on finite element methods, a vast class of numerical methods for differential equations which is of wide applicability and great utility, and also, not coincidentally, one for which there is an extensive body of mathematical analysis.

The finite element method is a mature tool, in both practice and theory, in many areas of computational science. Nonetheless, the variety of partial differential equations (PDEs) which arise is vast, and there are still many important problems for which the known numerical approaches fail, and good numerical methods are yet to be devised. Consequently, research aimed at devising and analyzing new methods is flourishing. Traditionally, the key mathematical tools for the study of numerical PDEs, and especially of finite element methods, have come from functional analysis: Hilbert and Banach spaces, the Hilbert projection theorem, the Lax-Milgram lemma, the Bramble-Hilbert lemma, duality, Sobolev spaces, etc. The finite element exterior calculus (FEEC), presented in this book, also depends essentially on functional analysis, especially the theory of closed unbounded operators on Hilbert space. But FEEC's mathematical arsenal goes well beyond functional analysis, bringing in tools from geometry and topology to develop and analyze numerical methods for classes of PDEs resistant to more traditional approaches. Methods derived from FEEC are prime examples of structure-preserving numerical methods, in that they are designed to preserve key geometric, topological, and algebraic structures of the PDE at the discrete level. This turns out to be crucial to the development of successful finite element methods for a variety of problems for which standard methods fail. Specifically, FEEC focuses on PDEs which relate to complexes of differential operators acting on Hilbert function spaces and uses finite element spaces which form subcomplexes of these complexes, and which can be related to them via commuting projections.

While FEEC's antecedents go back decades, to the early days of the finite element method and even before, it first began to be defined as a distinct theory in my presentation to the International Congress of Mathematicians in 2002 [5] and was formalized in two long papers I coauthored with Richard Falk and Ragnar Winther in 2006 [11] and 2010 [13]. The first paper emphasized a particular complex of differential operators, namely, the de Rham complex of differential forms on a domain in \mathbb{R}^3 (or a Riemannian manifold). It was here that the name *finite element exterior calculus* first appeared, referring to the calculus of differential forms. In the 2010 paper, more emphasis was put on the abstract structure of a Hilbert complex, of which the L^2 de Rham complex is a special case, allowing FEEC to deal with other complexes that arise in other applications.

By June 2012 the basic outlines of FEEC theory were in place, and I was fortunate to be offered the opportunity to present an intensive short course on FEEC to an audience of nearly 70 faculty members, graduate students, and other researchers from around the world. The course was generously supported by the National Science Foundation and the Conference Board of the Mathematical Sciences as part of the NSF-CBMS conference series and expertly hosted at the Institute of Computational and Experimental Research in Mathematics (ICERM) at Brown University. This book grew out of that course. It shares with the course the goal of helping numerical analysts to master the fundamentals of FEEC, including the geometrical and functional analysis preliminaries, quickly and in one place. But the book has a broader audience in mind than the course, aiming to be accessible as well to mathematicians and students of mathematics from areas other than numerical analysis who are interested in understanding how techniques from geometry and topology come to play a role in numerical PDE. FEEC has been vigorously developing in the time since the course, and so the book contains much more material than was taught in the course, some of which was not even developed at that time.

The first portion of the book, Chapters 1–5, quickly develops the prerequisite material from homological algebra, algebraic topology, and functional analysis. These ingredients are combined in the basic structure of a Hilbert complex studied in Chapter 4. Remaining in this general abstract framework, the approximation of problems related to Hilbert complexes is developed in Chapter 5. The second portion of the book consists of Chapters 6 and 7, where we apply the general theory to the most canonical example of a Hilbert complex, the L^2 de Rham complex on a domain in \mathbb{R}^n . Finally, in the closing chapter we briefly survey some other examples and applications.

I am grateful to NSF and CBMS for their support of the FEEC course in 2012 and of this volume and for the support I received during the period I was developing FEEC and writing the book from NSF grants DMS-1115291, DMS-1418805, and DMS-1719694. Ron Rosier and David Bressoud, the former and current directors of CBMS, are to be particularly thanked for their patience and understanding. I am also grateful to ICERM for hosting the course and especially to Alan Demlow, Johnny Guzmán, and Dmitriy Leykekhman, who conceived and organized it. The audience for the course, many of whom have gone on to make important contributions to FEEC, was also a great source of stimulation and inspiration. Several people have proofread all or part of the manuscript and made countless valuable suggestions: thanks to Richard Falk, Ragnar Winther, Shawn Walker, Espen Sande, and Kaibo Hu. Johnny Guzmán, Anil Hirani, and Ragnar Winther have even used an early version of the book as a text for a course they taught, which was particularly helpful.



Participants in the NSF-CBMS Conference on Finite Element Exterior Calculus, held at ICERM, Brown University, in June 2012.

Chapter 1 Introduction

The finite element exterior calculus, or FEEC, is a theoretical approach to the design and understanding of finite element methods for the numerical solution of a variety of partial differential equations. The finite element method itself is one of the most important technologies of computational science and engineering. A major contributor to its success is the mathematical framework which has been developed over the past 50 years, enabling rigorous error analysis and understanding of the properties of different finite element methods and computations. Traditionally, this framework is primarily based upon the language and tools of functional analysis. FEEC brings in additional tools from topology and homological algebra. More specifically, FEEC captures essential structures of de Rham cohomology and Hodge theory at the discrete level and relates the continuous and discrete structures in order to obtain stable and convergent finite element discretizations.

In the next chapters, we will introduce the necessary homological and functional analytic background material, and then use it to develop FEEC and apply it in various contexts. But first, the reader may appreciate some motivation. Why do we need an elaborate new theory for finite elements? After all, the finite element method performs brilliantly in many applications. In this chapter, we will exhibit some simple numerical examples where standard finite element methods fail utterly, but for which more sophisticated finite elements—carefully constructed to preserve key structures of the underlying PDE problem—come to the rescue.

First we set the stage by recalling some of the classical theory of the finite element method. The simplest setting for finite elements is the numerical solution of Dirichlet's problem for the Poisson equation. Given a real-valued function f on a domain Ω in \mathbb{R}^n , we seek a function u satisfying

$$-\Delta u = f \text{ in } \Omega, \quad u = 0 \text{ on } \partial \Omega.$$

The finite element method proceeds not from this strong formulation but from a *weak formulation* of the boundary value problem. The solution u is sought in the Sobolev space $\mathring{H}^1(\Omega)$ consisting of $L^2(\Omega)$ functions which have their first partial derivatives also in $L^2(\Omega)$ and which vanish on $\partial\Omega$. It is determined by requiring that the equation

$$\int_{\Omega} \operatorname{grad} u \cdot \operatorname{grad} v \, dx = \int_{\Omega} f v \, dx \tag{1.1}$$

hold for all functions v in the test space $\mathring{H}^1(\Omega)$ (which, in this example, is the same as the trial space in which u is sought). This equation is derived by multiplying the original Poisson equation by the test function v and integrating over the domain Ω by parts. The weak formulation of the problem is well-posed, meaning that for any f (in $L^2(\Omega)$, or even in the dual Sobolev space $H^{-1}(\Omega) = \mathring{H}^1(\Omega)'$) there is a unique weak solution u and, moreover, it depends continuously in H^1 on f in L^2 or H^{-1} .

To compute an approximate solution u_h , we first consider *Galerkin methods*, a class that includes the finite element methods. A Galerkin method requires that we specify a finite dimensional subspace V_h of $\mathring{H}^1(\Omega)$ to be used on the discrete level as both trial and test space. Thus we seek the approximate solution $u_h \in V_h$ by requiring that the weak equation (1.1) be satisfied for test functions v belonging to V_h :

$$\int_{\Omega} \operatorname{grad} u_h \cdot \operatorname{grad} v \, dx = \int_{\Omega} f v \, dx, \quad v \in V_h.$$
(1.2)

This is a finite dimensional linear system of equations with dimension equal to dim V_h . To compute its solution, we need to specify a basis ϕ_i of V_h and to compute the *stiffness matrix* with entries $\int_{\Omega} \operatorname{grad} \phi_j \cdot \operatorname{grad} \phi_i \, dx$ and the *load vector* with entries $\int_{\Omega} f \phi_i \, dx$, and then to solve the resulting matrix equation, whose solution gives the coefficients of u_h in the selected basis. The stiffness matrix is, in this simple case, symmetric and positive definite, and thus the approximate solution consists of the cost of computing the stiffness matrix and load vector entries for each basis element, and the cost of solving the resulting linear system of equations.

The solution of the Galerkin method not only is uniquely determined and computable, but is also easy to derive an error estimate for it. In the present case, we may simply invoke Poincaré's inequality

$$||w||_{H^1} \le c ||\operatorname{grad} w||_{L^2}, \quad w \in \check{H}^1,$$

where the constant c depends only on the domain Ω . It follows that the bilinear form on the left-hand side of (1.1) defines an inner product on \mathring{H}^1 which is norm equivalent to the H^1 inner product and so we may apply the Hilbert projection theorem to prove the solution is quasi-optimal in the H^1 norm:

$$\|u - u_h\|_{H^1} \le c \inf_{v \in V_h} \|u - v\|_{H^1}, \tag{1.3}$$

with the same constant c, independent of f, u, and the choice of subspace V_h .

A finite element method is a Galerkin method for which the Galerkin subspace V_h is a piecewise polynomial space which is efficiently computable (more precisely, it can be constructed through the assembly procedure recalled at the beginning of Section 7.1). The simplest finite elements are the Lagrange piecewise linears. In two dimensions the finite element space is obtained by triangulating the domain Ω (so the domain is assumed to be a polygon, or approximated by one) and taking V_h to consist of all continuous functions which are polynomial of degree at most 1 on each triangle of the triangulation and which vanish on $\partial\Omega$. Figure 1.1 shows the domain $\Omega = (0,3) \times (0,3) \setminus [2/3,2] \times [3/4,2]$ partitioned by a rather coarse triangulation of 256 triangles. In the center of the figure is an example of a piecewise linear function satisfying the Dirichlet boundary conditions.



Figure 1.1. *Left: A triangulation of an annular polygonal domain. Center: A piecewise linear function. Right: A Lagrange basis function.*

This space has dimension equal to the number of interior vertices of the triangulation, and it admits a convenient basis called the Lagrange basis, with each basis function equal to unity at a unique interior vertex and zero at all the others, as illustrated in the right image in the figure. This is a local basis in the sense that each basis function is supported in only a small number of triangles. As a result, the integrals needed for the stiffness matrix and load vector can be computed very quickly, and the stiffness matrix is highly sparse (having only seven nonzero elements per row on average, asymptotically as the mesh is refined). Thus the finite element solution is efficiently computable. By (1.3) its accuracy in $H^1(\Omega)$ is comparable to the best continuous piecewise linear approximation on the mesh. To improve the accuracy of the finite element solution, we can enrich the finite element space by refining the mesh. Alternatively, we can use Lagrange finite elements of higher degree, i.e., increase the degree of the piecewise polynomials to some d > 1. Combining the estimate (1.3) with approximation theory we ensure that the error goes to 0 as either the mesh size tends to 0 or the element degree tends to ∞ .

We have illustrated the basic aspects of the finite element method on a simple model problem. Such a treatment would generally be continued by establishing rates of convergence and their dependence on the triangulation and the regularity of the solution, proving error estimates in norms other than the H^1 norm, developing algorithms for a posteriori error estimation and adaptive meshing and analyzing their performance, and, of course, extensions to other problems. See one of the many available books for such treatments, e.g., [21, 22, 31, 40].

We will soon discuss a problem for which the standard finite element approach fails, in order to motivate the study of FEEC. However, first, let us continue with the standard finite element method, not for the source problem just discussed but rather for the Laplacian *eigenvalue* problem. Again, this is a problem for which standard finite elements are very successful and for which the theoretical justification, which relies heavily on spectral theory in Hilbert space, is extensive [16, 18]. Retaining the domain of Figure 1.1, we seek the fundamental mode and frequency (or perhaps several modes and frequencies) of an annular polygonal drum. That is, we want to find the least $\lambda \in \mathbb{R}$ for which there exists nonzero $u : \Omega \to \mathbb{R}$ such that

$$-\Delta u = \lambda u \text{ in } \Omega, \quad u = 0 \text{ on } \partial \Omega. \tag{1.4}$$

Passing to a weak formulation and applying Galerkin's method brings us to the discrete eigenvalue problem of finding $\lambda_h \in \mathbb{R}$ and nonzero $u_h \in V_h$ such that

$$\int_{\Omega} \operatorname{grad} u_h \cdot \operatorname{grad} v \, dx = \lambda_h \int_{\Omega} u_h v \, dx, \quad v \in V_h.$$

	Deg	Degree 1		ree 3
# Elements	λ_1	λ_2	λ_1	λ_2
256	9.986	12.003	9.211	11.179
1,024	9.428	11.404	9.198	11.171
4,096	9.279	11.245	9.193	11.168
16,384	9.216	11.189	9.191	11.167
65,536	9.199	11.174	9.190	11.167
262,144	9.193	11.169	9.190	11.166

Table 1.1. First and second eigenvalues for the Laplacian on the domain of Figure 1.1, computed on uniform refinements of the given mesh, using Lagrange elements of degree 1 and 3. The correct values are 9.190 and 11.166 to the nearest 0.001.

Via a basis, this reduces to a symmetric sparse generalized eigenvalue problem, with one matrix the stiffness matrix defined above and the other the mass matrix defined in a similar way but using the L^2 inner product occuring on the righthand side of the equation. This matrix generalized eigenvalue problem can be efficiently solved by an iterative method, such as the inverse power method to find the fundamental eigenpair, or Krylov–Schur iteration to calculate the portion of the spectrum within some interval. As to theory, one can prove, in this case, that the first eigenfunction u_h satisfies a quasi-optimal error estimate like (1.3) and that the eigenvalue error is bounded by a multiple of the square of eigenfunction error. See, e.g., [18] for a lucid exposition, including the case of more than one eigenvalue and the complications arising from eigenvalues which are not simple.

Next we illustrate the performance with some numerical results. For the coarse mesh of 256 triangles shown in Figure 1.1, and Lagrange linear finite elements, the computed value of the fundamental eigenvalue is 9.986 and the computed fundamental eigenfunction is the one plotted in the center of Figure 1.1. It is no surprise that these are not very accurate (the eigenvalue error turns out to be about 9%), given the coarseness of the mesh. However, we can improve the accuracy to the extent desired by refining the mesh and/or increasing the polynomial degree. In Table 1.1, we show the results for polynomials of degree 1 and 3, both for the fundamental eigenvalue and the second eigenvalue, and for a sequence of meshes obtained from the original one by uniform refinement of the initial mesh. With piecewise linears and 65,536 triangles or piecewise cubics and 1.024 triangles we obtain the fundamental eigenvalue to within about 0.1%. Figure 1.2 shows the fundamental eigenfunction computed with linear Lagrange elements on the mesh with 65,536 triangles. It is worth remarking that the same accuracy as is achieved with these uniformly refined meshes could be obtained with far fewer triangles using adaptive refinement, but that is not a direction we shall discuss in this book.

Now we consider a problem for which the classical approach is insufficient and FEEC comes to the rescue. For this, we move slightly away from the scalar Laplacian. Instead we consider the vector Laplacian, the operator $-\Delta = \text{curl curl}$ $- \text{grad div acting on a vector field } u : \Omega \to \mathbb{R}^2$. (Since we are in two dimensions, curl u is a scalar function, while the second application of curl maps the scalar function back to a vector field.) Thus we consider the eigenvalue problem



Figure 1.2. The fundamental eigenfunction computed with piecewise linears on 65,536 triangles.

These boundary conditions, sometimes called magnetic for their application in electromagnetism, allow for a simple weak formulation. Define the space $\mathring{H}(\operatorname{div})$ consisting of vector fields in L^2 with divergence in L^2 and for which $u \cdot n$ vanishes on the boundary, and also the space $H(\operatorname{curl})$ of L^2 vector fields with curl in L^2 (with no boundary conditions imposed). If we take an arbitrary test vector field $v \in \mathring{H}(\operatorname{div}) \cap H(\operatorname{curl})$ and compute its L^2 inner product with the vector fields on both sides of (1.5), we get, after integration by parts, that

$$\int_{\Omega} \left[\operatorname{curl} u \cdot \operatorname{curl} v + (\operatorname{div} u)(\operatorname{div} v)\right] dx = \lambda \int_{\Omega} u \cdot v \, dx, \tag{1.6}$$

where no boundary integral term arises from the integration by parts of the curl since $\operatorname{curl} u = 0$ on the boundary, and none arises in the integration by parts of the gradient since $v \cdot n = 0$. Thus, we obtain a weak formulation of (1.5) by seeking $u \in \mathring{H}(\operatorname{div}) \cap H(\operatorname{curl})$ satisfying (1.6) for all v belonging to the same space. Note that the boundary condition on $u \cdot n$ is essential in this formulation—it is imposed on the space in which the solution is sought—while the boundary condition on $\operatorname{curl} u$ is natural, arising from the weak formulation.

If we naively attempt to compute the eigenvalues using a method analogous to that which we used for the scalar Laplacian, we are in for a big disappointment. To illustrate this, we discretize the weak formulation (1.6) using Galerkin's method, taking as the Galerkin subspace the space of vector fields both of whose components are continuous piecewise linear polynomials (or piecewise polynomials of higher degree: vector Lagrange elements of degree d). The results are reported in Table 1.2, which is altogether analogous to Table 1.1, except that it pertains to the vector Laplacian eigenvalue problem (1.5) instead of to the scalar Laplacian eigenvalue problem (1.4). The computed eigenvalues tabulated in Table 1.2 appear to be converging. For example, the computations strongly suggest that the fundamental eigenvalue is about 1.8. But this is entirely wrong! As we shall see later in the book, the vector Laplacian on a plane domain with the given boundary conditions has a kernel whose dimension is the first Betti number of the domain: the number of holes. Therefore, for our problem the Laplacian has a one-dimensional kernel and the first eigenvalue is precisely 0. The convergence of the finite element computation with the Lagrange elements to 1.8 is entirely

	Deg	ree 1		Degree 3			
# Elements	λ_1	λ_2	-	λ_1	λ_2		
256	2.270	2.360		1.896	1.970		
1,024	2.050	2.132		1.854	1.925		
4,096	1.940	2.016		1.828	1.897		
16,384	1.879	1.952		1.812	1.880		
65,536	1.843	1.914		1.802	1.870		
262,144	1.821	1.890		1.796	1.863		

Table 1.2. First and second eigenvalues for the vector Laplacian computed using vector Lagrange elements of degree 1 and 3. The results are completely wrong, the correct values being 0 and 0.617.

Table 1.3. First and second eigenvalues for the vector Laplacian computed using FEEC, exhibiting convergence to the correct values of 0 and 0.617.

	Deg	ree 1	Degree 3		
# Elements	λ_1	λ_2		λ_1	λ_2
256	0.000	0.638		0.000	0.619
1,024	0.000	0.625		0.000	0.618
4,096	0.000	0.620		0.000	0.617
16,384	0.000	0.618		0.000	0.617
65,536	0.000	0.618		0.000	0.617
262,144	0.000	0.617		0.000	0.617

misleading. The situation is no better for the second eigenvalue either. The Lagrange finite element computation converges to a value of around 1.86, but in fact the correct value is totally different, namely, 0.617.

The abject failure of the basic finite element method for the vector Laplacian eigenvalue problem can be remedied with FEEC. A FEEC approach to this problem uses a weak formulation which is different from (1.6) and discretizes it not with Lagrange finite elements but with specially designed finite elements. The results of the FEEC approach are shown in Table 1.3. Notice that the FEEC methods compute the fundamental eigenvalue of 0 *exactly* (except for roundoff error). As mentioned, this eigenvalue reflects a fundamental topological structure of the problem—the hole in the domain. This is a first example of FEEC methods being *structure-preserving*. In topological terms, they preserve the homology of the domain. The FEEC methods are not exact for the next eigenvalue, of course, but they do converge to it nicely. In Figure 1.3, we show the first two eigenfunctions computed with the lowest order FEEC method on the mesh with 4,096 triangles. Notice that the first eigenfunction is a vector field which swirls around the hole in the domain. Both eigenfunctions have strong singularities at the reentrant corners.

The reader may have noticed several features of this example. It is an eigenvalue problem rather than a source problem. The domain has both a hole, which is tied to the existence of the zero eigenvalue, and a reentrant corner, which gives rise to a singular solution. The singularity is in fact quite strong: the magnitude of the eigenfunction itself, and not just of its derivatives, grows unboundedly as we approach one of the reentrant corners. The last point is the most crucial for the failure of the standard finite element method based on the weak formulation (1.6).



Figure 1.3. Top: The first two eigenfunctions for (1.5) computed with the FEEC approach. The color gives the magnitude of the vector and the arrow length is proportional to it. Bottom: The magnitude of the first eigenfunction on the line y = 3/4 tangent to the bottom side of the rectangular hole. (Vertical scale left unspecified as eigenfunctions are determined only up to a constant multiple.)

We shall return to this point in Section 5.1 and fully explain the failure. There we will also give another example, a source problem on a simply connected domain with a reentrant corner, and see that the standard finite element method fails for it as well. (The interested reader may peek ahead to Figure 5.1 to see.) Then, in Section 5.2 we will show that these problems fit within the FEEC framework of analysis developed in the first part of the book, and so the convergence of the FEEC approach for them is guaranteed.

At the conclusion of Section 5.1 we also give a third, quite different, example of failure of the standard finite element method for a vector Laplace source problem. In that case the domain is a circular annulus and so has a smooth boundary and a smooth solution up to the boundary. The failure must therefore have a different origin, not dependent on singularity. Indeed, it turns out to be related to the zero eigenvalue, as we shall explain there.

Another popular class of examples for demonstrating the pitfalls of naive finite element methods and the need for methods based on the formulations and finite element spaces we develop in this book involves neither holes in the domain nor singular behavior. A simple example is the Maxwell eigenvalue problem, which in its simplest form is the eigenvalue problem for the curl curl operator, together with, say, the electric boundary condition $u \times n = 0$. The eigenvalues λ for this problem are all real and nonnegative. However, for a simply connected domain, the only eigenpairs which are of interest are those for which $\lambda > 0$, since for these the equation implies that div u = 0, which is a requirement of Maxwell's equations. There exists also an infinite dimensional space of eigenvectors for the eigenvalue $\lambda = 0$, comprising gradients of smooth functions which vanish on the boundary. However, we do not consider 0 to be a Maxwell eigenvalue, since these gradient eigenfunctions are not divergence-free. (In case the domain is not simply connected there is also be a finite dimensional space of divergence-free 0 eigenvectors.)

We will consider the Maxwell eigenvalue problem on a square domain $\Omega = (0, \pi) \times (0, \pi)$. In this case, the exact eigenpairs have a simple analytic expression which can be computed by separation of variables. The eigenvectors are

$$u(x,y) = \operatorname{curl}(\sin mx \, \sin ny) = (n \sin mx \, \cos ny, -m \cos mx \, \sin ny),$$

where m and n are nonnegative integers, not both zero, and the corresponding eigenvalues $\lambda = m^2 + n^2$, i.e.,

$$\lambda = 1, 1, 2, 4, 4, 5, 5, 6, 7, 7, 10, 10, 13, 13, \dots$$

A weak formulation of the eigenvalue problem seeks nonzero $u \in H(curl)$ and $\lambda > 0$ such that

$$\int_{\Omega} \operatorname{curl} u \, \operatorname{curl} v \, dx = \lambda \int_{\Omega} u \cdot v \, dx, \quad v \in \mathring{H}(\operatorname{curl}),$$

where the space $\dot{H}(\text{curl})$ incorporates the boundary condition $u \times n = 0$. We now show the results of discretizing this formulation using Galerkin's method. For the Galerkin finite element method we investigate two different sorts of uniform triangulations and, for each of these, two different finite element spaces. The first mesh we use is of diagonal type, obtained by dividing the square into $40 \times$ 40 subsquares and dividing each of them into two triangles using the positively sloped diagonal, and the second is of crisscross type, with each subsquare divided into four triangles using both diagonals. Meshes of the two types are shown in Figure 1.4, but with only 10×10 subsquares for improved visibility.

For the finite element spaces, we compare two possibilities on each mesh: the simple vector Lagrange elements of degree 1 (continuous piecewise linear



Figure 1.4. Uniform meshes, diagonal and crisscross. The meshes used for computation were finer, with 16 times as many elements.

Exact	1	1	2	4	4	5	5	8	9	9	10	10
Diagonal me	Diagonal mesh											
Lagrange FEEC	5.16 1.00	5.26 1.00	5.26 2.00	5.30 4.00	5.39 4.00	5.45 5.00	5.53 5.00	5.61 8.01	5.61 8.98	5.62 8.99	5.71 9.99	5.73 9.99
Crisscross mesh												
Lagrange FEEC	1.00 1.00	1.00 1.00	2.00 2.00	4.00 4.00	4.00 4.00	5.00 5.00	5.00 5.00	6.00 7.99	8.01 9.00	9.01 9.00	9.01 10.00	10.02 10.00

 Table 1.4. First 12 Maxwell eigenvalues and Galerkin approximations of them.

vector fields fulfilling the boundary conditions) and elements fitting the FEEC framework, the so-called Nédélec edge elements of lowest degree. In each case, we had the Krylov–Schur eigenvalue solver report the 12 eigenvalues nearest 5.5. The results are shown in Table 1.4. Comparing the exact eigenvalues to the ones computed with Lagrange elements on the diagonal mesh, we see *complete failure:* the eigenvalues are nowhere near the correct ones. On the crisscross mesh, we have failure again, but it is more subtle (and arguably more dangerous). In fact, 11 of the first 12 eigenvalues are accurately computed, but the eighth computed eigenvalue is 6.00. This value is *completely spurious:* the correct value is 8. Had we computed more than 12 eigenvalues, we would have found additional spurious eigenvalues, the next being a spurious double eigenvalue around 15. Moreover, when a finer crisscross mesh is used, these spurious values persist (and apparently converge to the precise, but spurious, values 6 and 15). By contrast, the FEEC method converges nicely for both meshes, a fact that can be backed by theory. We will return to this example in Section 8.1.

Having surveyed some motivating examples, we conclude the introduction with a description of the remainder of the book. In the first part of this exposition, in Chapters 2–5, we develop the basic abstract framework and results of FEEC. The key structure of this framework is a Hilbert complex. A Hilbert complex arises from the marriage of the algebraic structure of a chain complex, as studied in homological algebra, with the analytic setting of unbounded operators on Hilbert spaces from functional analysis. In Chapter 2 we provide an elementary self-contained introduction to the few results we need from homological algebra, and in Chapter 3 we introduce the basic results we need concerning unbounded operators on Hilbert space. These two chapters establish material which is prerequisite to what follows. Both of these, and subsequent chapters, end with a box summarizing the main results obtained in the chapter. Readers familiar with the subject of the chapter can skip forward to the boxed summary at the end and decide whether it is necessary to read the body of the chapter. After the first two prerequisite chapters, we develop Hilbert complexes in Chapter 4. Throughout, we illustrate the abstract concepts with the important and canonical example of the L^2 de Rham complex on a domain in three-dimensional Euclidean space. Each Hilbert complex gives rise to a sequence of operator equations, which are the abstract generalizations of scalar and vector Poisson equations in the case of the three-dimensional de Rham complex, as well as other elliptic partial differential equations. The accurate numerical solution of these problems is the goal of FEEC. The approximation of these so-called abstract Hodge Laplacian problems is studied within the Hilbert complex framework in Chapter 5, where the main

theoretical results of FEEC are established. In particular, we show that consistent, stable, and so convergent approximations are obtained by Galerkin's method applied to a mixed formulation of the equations as long as the Galerkin subspaces satisfy a few important hypotheses.

In the second part of the book, Chapters 6 and 7, we specialize from a general Hilbert complex to the L^2 de Rham complex on an *n*-dimensional domain or Riemannian manifold. In this case, the spaces are not abstract Hilbert spaces but rather Hilbert spaces of differential forms. To treat this case, we introduce the language and basic results of exterior calculus in Chapter 6, which allows for a unified and more insightful treatment of all the spaces and operators involved. Our next task is then to define, in the concrete setting of the de Rham complex, Galerkin subspaces of differential forms. This is the main task of Chapter 7.

Finally, in the concluding chapter of the book we give a glimpse of extensions and applications that take us beyond the de Rham complex and to additional applications.

Chapter 2 Basic notions of homological algebra

The basic structures of homological algebra, such as chain complexes and homology groups, were introduced in the 19th century with the aim of defining topological invariants such as the Betti numbers via the counting of discrete structures. However, their antecedents go back much further, to Euler and even Descartes. With the development of differential forms by E. Cartan at the start of the 20th century and the study of their cohomology by de Rham in the 1930s these same algebraic structures appeared in the context of spaces of functions acted on by partial differential operators. Later yet they became objects of study from an abstract algebraic point of view. Thus homological algebra has long played a fundamental role in algebraic topology, differential geometry, and algebra. As useful as it is, however, it is often unfamiliar to applied mathematicians and numerical analysts. Therefore, in this chapter, we include a short self-contained introduction to the basic aspects of homological algebra we shall need. The key points are summarized in Box 2.1 at the end of the chapter.

2.1 • Graded vector spaces, chain complexes, and differentials

A graded vector space is a vector space V expressed as a direct sum of subspaces V_k indexed by the integers:

$$V = \bigoplus_{k=-\infty}^{\infty} V_k$$

For example, the space $\mathcal{P}(\mathbb{R}^n)$ of polynomial functions in n variables is graded with the *k*th summand equal to the space of homogeneous polynomials of degree k if $k \ge 0$ and 0 otherwise. A graded vector space is called *nonnegative* if, as in this case, $V_k = 0$ for k < 0. Unlike the full space of polynomials, most of the graded spaces we encounter will be nonnegative and, in addition, *finite*, meaning that only finitely many of the summands are nonzero. For example, the space $\mathcal{P}_r(\mathbb{R}^n)$ of polynomials of degree at most r is graded with only r + 1 nonzero summands.

A linear map $f: V \to W$ between graded vector spaces is called a *graded* map of degree p if $f(V_k) \subset W_{k+p}$ for k. The Laplace operator, for example, acts

as a graded linear map of degree -2 on the polynomials, while multiplication by some fixed homogeneous polynomial of degree $p \ge 0$ is a graded map of degree p.

A *chain complex* is simply a graded vector space V furnished with a graded linear map $\partial : V \to V$ of degree -1 which satisfies $\partial^2 = 0$. The operator ∂ , or, equivalently, the sequence of the operators $\partial_k = \partial|_{V_k}$, is called the *differential* of the chain complex, or, sometimes, the *boundary operator*. In other words, a chain complex consists of a sequence of vector spaces V_k and linear maps $\partial_k : V_k \to$ $V_{k-1}, k \in \mathbb{Z}$, with the property that $\partial_k \circ \partial_{k+1} = 0$. We will often write a chain complex in the form

$$\cdots \to V_{k+1} \xrightarrow{\partial_{k+1}} V_k \xrightarrow{\partial_k} V_{k-1} \to \cdots$$

but when we wish to refer to it briefly, we will simply write (V, ∂) or even just V. The elements of V_k are sometimes called k-chains.

A graded subspace of a graded vector space is, naturally, a subspace S which decomposes as the direct sum of subspaces $S_k \subset V_k$. If S is a graded subspace of a complex (V, ∂) such that $\partial S_k \subset S_{k-1}$, then $(S_k, \partial_k|_{S_k})$ is itself a complex, called a *subcomplex* of (V, ∂) .

2.2 - Cycles, boundaries, and homology

Given a chain complex (V, ∂) , the null space \mathfrak{Z} and the range \mathfrak{B} of ∂ are each graded subspaces of V. (They are subcomplexes, but with pretty boring differentials, since the restriction of ∂ to each is zero.) The elements of \mathfrak{Z}_k are called the *k*-cycles and the elements of the range \mathfrak{B}_k of ∂_{k+1} are called the *k*-boundaries. The differential property $\partial^2 = 0$ means that $\mathfrak{B}_k \subset \mathfrak{Z}_k$, and the *k*th homology space is defined to be the quotient space $\mathcal{H}_k = \mathfrak{Z}_k/\mathfrak{B}_k$. Thus the elements of the homology space are equivalence classes of cycles, with two *k*-cycles considered equivalent if their difference is the boundary of a (k + 1)-chain. If the boundary spaces \mathfrak{B}_k are exactly equal to the cycle spaces \mathfrak{Z}_k , rather than proper subspaces, then the homology spaces \mathcal{H}_k all vanish. In this case the complex is said to be *exact*.

2.3 • Example: The simplicial chain complex

A classical example of a chain complex, the one which accounts for much of the terminology, is the *simplicial chain complex* associated to a simplicial complex. By a k-dimensional *simplex* in some Euclidean space \mathbb{R}^n we mean the convex hull f of k + 1 points x_0, \ldots, x_k , called the vertices of f, which we assume to be in general position in \mathbb{R}^n . Each nonempty subset of the vertex set of f determines a simplex of some dimension $\leq k$, which is called a face of f. A *simplicial complex* is a finite set S of simplices in \mathbb{R}^n such that (1) whenever a simplex belongs to S, each of its faces also belongs to S, and (2) if f and g are two simplices in S, then their intersection $f \cap g$ is either empty or a face of f and of g. The underlying space of the simplicial complex, that is, the subset of \mathbb{R}^n obtained as the union of the simplices, may be a polyhedral domain, a piecewise flat manifold, or a more complicated set. The dimension of the simplicial complex is the highest dimension of any of its simplices, and a simplicial complex of dimension k is said to be *pure* if every simplex belonging to it is contained in a simplex of dimension k.



Figure 2.1. Oriented simplices of dimensions 1, 2, and 3.

The notion of a pure simplicial complex is essentially the same as that of a *tri-angulation* of the underlying space. More precisely, a finite set of n-dimensional simplices constitute a triangulation of their union if and only if they, together with all their faces of all dimensions, form a pure n-dimensional simplicial complex.

A simplex may be *oriented* by selecting an ordering of its vertices, with the understanding that two orderings are equivalent if they differ by an even permutation. Thus every simplex of dimension k > 0 has two possible orientations. See Figure 2.1. We shall write $[x_0, \ldots, x_k]$ for the simplex with vertices x_0, \ldots, x_k endowed with the orientation determined by the given ordering of the vertices. We may (and shall) designate a preferred orientation for all the simplices in the simplicial complex, e.g., lexicographic ordering. The oriented simplex $[x_0, \ldots, x_k]$ is then said to be of *positive orientation* if the vertex sequence x_0, \ldots, x_k , or some even permutation of it, is in increasing order.

Let S be a simplicial complex in \mathbb{R}^n and denote by $\Delta_k(S)$ its set of simplices of dimension k. A k-chain is a formal linear combination of k-dimensional oriented simplices,

$$c = \sum_{f \in \Delta_k(\mathcal{S})} c_f f,$$

with -1 times an oriented simplex being understood as the same simplex with the opposite orientation. We shall mostly be concerned with chains with real coefficients, meaning that $c_f \in \mathbb{R}$. The space of all k-chains is thus a real vector space of dimension $\#\Delta_k(S)$, denoted by $C_k(S)$, or simply C_k . Given any oriented k-simplex $f = [x_0, x_1, \ldots, x_k]$ in S, we define its *boundary*, ∂f , as the (k-1)-chain

$$\partial[x_0, x_1, \dots, x_k] := \sum_{i=0}^k (-1)^i [x_0, \dots, \hat{x}_i, \dots, x_k],$$

where the hat over x_i means that that vertex is omitted. Thus the boundary of an oriented k-simplex, ∂f , is the formal sum of the k-1 simplices contained in the geometrical boundary of the simplex, each endowed with the proper orientation (through the signs in the above sum). See Figure 2.2. The boundary operator extends to general chains by linearity: $\partial c = \sum c_f \partial f$. As a linear operator ∂_k can be expressed as a matrix which is the *oriented incidence matrix* for the simplicial complex. It has a column for each k-chain, a row for each (k-1)-chain, and its entries are 1, -1, or 0 according to whether the (k-1)-chain belongs to the boundary of the k-chain with the same orientation, with the opposite orientation, or not at all.



Figure 2.2. Boundaries of chains.



Figure 2.3. Two cycles which belong to the same nonzero homology class.

In this context, the k-boundaries \mathfrak{B}_k are truly the k-chains which arise as the boundaries of (k + 1)-dimensional chains, and the cycles are those chains which have no boundary, like the final chain in Figure 2.2, which describes a closed curve, i.e., one without end points. The 1-cycles which are not boundaries signal the existence of holes going through the domain. For example, Figure 2.3 shows two such 1-cycles. However, they both signal the same hole since their difference is a 1-boundary. In other words, they represent the same coset in the homology group $\mathcal{H}_k = \mathfrak{Z}_k/\mathfrak{B}_k$. The dimension of the kth homology group is called the kth *Betti number* of the simplicial complex \mathcal{S} : $b_k = \dim \mathcal{H}_k$. The Betti numbers are an invariant of the underlying space: they do not depend on the particular triangulation. They are, in fact, one of the simplest topological invariants and so one of the motivating examples of algebraic topology. Their meanings for domains in two and three dimensions are as follows:

- b_0 is the number of connected components of the domain.
- b_1 is the number of holes through the domain ("tunnels" in the three-dimensional case). For example, b_1 is equal to 1 for a solid torus.

- b_2 is zero for any bounded domain in \mathbb{R}^2 (because such a domain always has a boundary and so cannot be a cycle). For a three-dimensional domain, b_2 is the number of voids enclosed by the domain, e.g., $b_2 = 1$ for a hollow ball.
- b₃ is zero for any domain in two or three dimensions.

The Betti numbers of some three-dimensional domains are given in Figure 2.4.



Figure 2.4. *Betti numbers* b_0, b_1, b_2, b_3 *for some domains in* \mathbb{R}^3 .

2.4 • Chain maps

A chain map from a chain complex (V, ∂) to a second chain complex (W, ξ) is a sequence of linear maps $f_k : V_k \to W_k$ such that the following diagram commutes:

$$\cdots \longrightarrow V_{k+1} \xrightarrow{\partial_{k+1}} V_k \xrightarrow{\partial_k} V_{k-1} \longrightarrow \cdots$$

$$f_{k+1} \downarrow \qquad f_k \downarrow \qquad f_{k-1} \downarrow \qquad f_{k-1} \downarrow \qquad \cdots$$

$$\cdots \longrightarrow W_{k+1} \xrightarrow{\xi_{k+1}} W_k \xrightarrow{\xi_k} W_{k-1} \longrightarrow \cdots$$

In other words, a chain map is a graded linear map of degree 0 which commutes with the differentials $f \circ \partial = \xi \circ f$. Since $f(\partial c) = \xi f(c)$ for $c \in V_{k+1}$, f maps the k-boundaries of V to k-boundaries of W, and the same equation applied to $c \in \mathfrak{Z}_k$ shows that f maps k-cycles to k-cycles. It follows that the chain map f induces a linear map \overline{f} on homology, i.e., the map $\overline{f}_k : \mathcal{H}_k(V) \to \mathcal{H}_k(W)$ which takes the homology class \overline{v} of some $v \in V_k$ to the homology class $\overline{f}_k v$ of $f_k v \in W_k$. This map is functorial. That is, if (X, ζ) is a third chain complex and g is a chain map from W to X, then $g \circ f$ is a chain map from V to X and $\overline{g \circ f} = \overline{g} \circ \overline{f}$.

A case we shall find of great interest later is when the complex W is a subcomplex of V and the maps f_k are projections of V_k onto W_k , i.e., $f_k v = v$ for $v \in W_k$. In this case we say that f is a *chain projection*.

Proposition 2.1. A chain projection from a complex to a subcomplex induces a surjection on homology.

Proof. The subcomplex inclusion maps $W_k \hookrightarrow V_k$ certainly form a chain map, *i*. When *f* is a projection, $f \circ i$ is the identity map from *W* to itself, and so $\overline{f} \circ \overline{i} = \overline{f \circ i}$ is the identity on homology and is, in particular, surjective. Therefore \overline{f} must be surjective. \Box

2.5 - Cochain complexes

A cochain complex is the same thing as a chain complex except with the indexing reversed: cochain differentials have degree +1 rather than -1. One typically writes the index of a cochain complex as a superscript rather than a subscript and often uses d rather than ∂ to denote the differential:

$$\cdots \to V^{k-1} \xrightarrow{d^{k-1}} V^k \xrightarrow{d^k} V^{k+1} \to \cdots$$

Of course, all the definitions and results above carry over in an obvious way. For example, the *k*th cohomology space \mathcal{H}^k is defined as the quotient of the cocycles \mathfrak{Z}^k by the coboundaries \mathfrak{B}^k .

Since any linear map $L: V \to W$ between vector spaces induces a dual map $L': W' \to V'$ between the dual spaces, any chain complex gives rise to a dual cochain complex. For example, from the simplicial chain complex associated to a simplicial complex, we obtain the simplicial cochain complex. Since the k-dimensional simplices form a basis for C_k , a k-cochain is an element of $C^k := C'_k$, i.e., a function ω which assigns to each k-dimensional simplex a real number $\omega(f)$ (extended to chains by linearity). The differential $\partial^k : C^k \to C^{k+1}$ is defined by $\partial^k = \partial'_{k+1}$, i.e.,

$$(\partial^k \omega)g = \omega(\partial_{k+1}g), \quad \omega \in C^k, \ g \in C_{k+1}.$$

This means that the matrix of the operator ∂^k is the transpose of the oriented incidence matrix which represents ∂_{k+1} . For example, ∂^0 applied to a vertex is the sum of the edges incident on that vertex, each multiplied by 1 or -1 according to whether the edge points into or out of the vertex. Of course, the information contained in the simplicial cochain complex

$$0 \to C^0 \xrightarrow{\partial^0} C^1 \xrightarrow{\partial^1} \cdots \xrightarrow{\partial^{n-1}} C^n \to 0 \tag{2.1}$$

is essentially the same as that contained in the simplicial chain complex. In particular, it can be shown that the cohomology space \mathcal{H}^k is naturally isomorphic to \mathcal{H}'_k , the dual of the corresponding homology space, and so its dimension is the *k*th Betti number.

2.6 • Example: The de Rham complex

We will have a lot more to say about the de Rham complex

$$0 \to \Lambda^0(\Omega) \xrightarrow{d^0} \Lambda^1(\Omega) \xrightarrow{d^1} \cdots \xrightarrow{d^{n-1}} \Lambda^n(\Omega) \to 0$$
 (2.2)

associated to any *n*-dimensional domain or smooth manifold Ω . Here $\Lambda^k(\Omega)$ is the space of (smooth) differential *k*-forms on Ω and the operator *d* is the exterior derivative, as introduced by Élie Cartan [27]. We shall put off the introduction of the language of differential forms until Chapter 6. For now we just describe the case of Ω a domain in \mathbb{R}^3 , for which the de Rham complex can be described in the terminology of vector calculus.

The complex is then

$$0 \to C^{\infty}(\Omega) \xrightarrow{\text{grad}} C^{\infty}(\Omega; \mathbb{R}^3) \xrightarrow{\text{curl}} C^{\infty}(\Omega; \mathbb{R}^3) \xrightarrow{\text{div}} C^{\infty}(\Omega) \to 0.$$
 (2.3)

The calculus identities $\operatorname{curl}\operatorname{grad}\phi = 0$ and $\operatorname{div}\operatorname{curl} u = 0$ ensure that (2.3) is indeed a complex. Its cohomology spaces \mathcal{H}_{dR}^k comprise the de Rham cohomology of the domain.

If the domain Ω is triangulated by some simplicial complex S, there is an interesting cochain map from the de Rham complex to the simplicial cochain complex called the *de Rham map*. Namely, we associate to a function $\phi \in C^{\infty}(\Omega)$ the 0-cochain $F^{0}\phi$ determined by $F^{0}\phi(x) = \phi(x)$ for each $x \in \Delta_{0}(S)$. As we shall see in Chapter 6 a differential k-form may be integrated on an oriented k-dimensional manifold, and this is the simplest case of k = 0. This is extended by linearity to arbitrary 0-chains, so $F^{0}\phi(\sum c_{i}x_{i}) = \sum c_{i}\phi(x_{i})$. Similarly we associate to a vector field $\psi \in C^{\infty}(\Omega; \mathbb{R}^{3})$ the 1-cochain whose value on any oriented edge $e \in \Delta_{1}(S)$ is

$$F^{1}\psi(e) = \int_{e} \psi \cdot t \, ds_{e}, \quad e \in \Delta_{1}(\mathcal{S}),$$

where the t denotes the unit vector tangent to the edge e directed according to the orientation. The vector field ψ may be viewed as a differential 1-form, and it has a trace on e which is a differential 1-form on the one-dimensional manifold e and so may be integrated over it. The line integral above is the expression of that integral in vector calculus terminology. A vector field in \mathbb{R}^3 may alternately be viewed as a differential 2-form, in which case the integral of its trace over a surface is defined and may be realized as

$$F^2\psi(f) = \int_f \psi \cdot n \, ds_f, \quad f \in \Delta_2(\mathcal{S}),$$

where now n is the unit normal to the face f, its direction determined by the orientation of f according to the right-hand rule. Finally, a scalar function ϕ may be thought of as a 3-form, leading to

$$F^3\phi(T) = \int_T \phi \, dx$$

for a positively oriented tetrahedron $T \in \Delta_3(S)$. We have thus defined the de Rham map connecting the de Rham complex (2.3) to the simplicial cochain complex (2.1):

The de Rham map is a cochain map, which is to say that this diagram commutes. For example, let us verify that $F^1 \circ \text{grad} = \partial^0 \circ F^0$. We must show that, for an arbitrary smooth function ϕ , the 1-cochains $F^1 \operatorname{grad} \phi$ and $\partial^0 F^0 \phi$ coincide, i.e., that

$$F^1 \operatorname{grad} \phi([x_0, x_1]) = \partial^0 F^0 \phi([x_0, x_1])$$

for any oriented edge $[x_0, x_1]$ in the complex. Letting t denote the unit vector

 $(x_1 - x_0)/|x_1 - x_0|$, we have

$$F^{1} \operatorname{grad} \phi([x_{0}, x_{1}]) = \int_{[x_{0}, x_{1}]} \operatorname{grad} \phi \cdot t \qquad (\text{definition of } F^{1})$$

$$= \phi(x_{1}) - \phi(x_{0}) \qquad (\text{fundamental theorem of calculus})$$

$$= F^{0}\phi([x_{1}] - [x_{0}]) \qquad (\text{definition of } F^{0})$$

$$= F^{0}\phi(\partial_{1}[x_{0}, x_{1}]) \qquad (\text{definition of } \partial_{1})$$

$$= \partial^{0}F^{0}\phi([x_{0}, x_{1}]) \qquad (\partial^{0} = \partial'_{1}).$$

The commutativity of the other two cells in the diagram is proven in the same way, using the Stokes theorem and the divergence theorem from vector calculus:

$$\int_{f} \operatorname{curl} \psi \cdot n = \int_{\partial f} \psi \cdot t, \quad \int_{T} \operatorname{div} \psi = \int_{\partial T} \psi \cdot n.$$

The de Rham complex which constitutes the first row of the diagram (2.4) is very different from the simplicial cochain complex which is its second row. The former is composed of infinite dimensional function spaces connected by differential operators, while the latter consists of finite dimensional spaces with operators that are essentially combinatorial (multiplication by matrices whose entries are all 1, -1, or 0). However, the *cohomology* of the two complexes is identical. This statement is a form of de Rham's theorem.

Theorem 2.2 (de Rham). The map on cohomology induced by the de Rham map is an isomorphism from the de Rham cohomology spaces onto the simplicial cohomology spaces.

Remark 2.3. De Rham's theorem as just stated equates de Rham cohomology, which is constructed analytically using differential operators, with the more geometric simplicial cohomology, constructed from chains of simplices. Although simplicial homology is in some sense the simplest and most easily computed homology theory, topologists often prefer singular homology for its greater generality and theoretical flexibility. In particular, simplicial homology is defined only for spaces which can be triangulated (or are homeomorphic images of such spaces), while singular homology, being built from chains of singular simplices (arbitrary continuous images of simplices) is defined for any topological space. An early, quite elementary, result in most expositions of homology is that the two approaches give identical homology (or cohomology) spaces when both apply, that is, for triangulated spaces. De Rham's theorem, which is a deeper result, is then usually stated as giving an isomorphism between de Rham cohomology and singular cohomology on a smooth manifold.

De Rham's theorem is a deep result. However, we shall see in Section 7.6 that a version of it can be proven quite easily once we have assembled the tools of FEEC. Box 2.1. Summary of Chapter 2 on homological algebra.

A chain complex is a sequence of vector spaces and operators

$$\cdots \to V_{k+1} \xrightarrow{\partial_{k+1}} V_k \xrightarrow{\partial_k} V_{k-1} \to \cdots$$

with successive operators composing to zero. It gives rise to three vector spaces at each level k:

- \mathfrak{Z}_k , the null space of the boundary operator ∂_k , consisting of the *cycles*;
- \mathfrak{B}_k , the range of ∂_{k+1} , the *boundaries*;
- $\mathcal{H}_k = \mathfrak{Z}_k/\mathfrak{B}_k$ the homology space.

An important example is the *simplicial chain complex* associated to a triangulation or more general simplicial complex. The dimensions of its homology spaces are the Betti numbers of the underlying space.

A cochain complex has the indexing reversed,

$$\cdots \to V^{k-1} \xrightarrow{d^{k-1}} V^k \xrightarrow{d^k} V^{k+1} \to \cdots,$$

giving rise to cocyles, coboundaries, and cohomology.

The dual spaces and dual maps of a chain complex form a cochain complex.

Another cochain complex is the *de Rham complex* which, on a threedimensional domain, can be written

$$0 \to C^{\infty}(\Omega) \xrightarrow{\text{grad}} C^{\infty}(\Omega; \mathbb{R}^3) \xrightarrow{\text{curl}} C^{\infty}(\Omega; \mathbb{R}^3) \xrightarrow{\text{div}} C^{\infty}(\Omega) \to 0.$$

The de Rham map is a *cochain map* from the de Rham complex to the dual of the simplicial chain complex and, as such, induces a map on cohomology. De Rham's theorem states that the induced map is an isomorphism.

Chapter 3 Basic notions of unbounded operators on Hilbert spaces

Much of the modern theory of partial differential equations and of their numerical discretization is most elegantly studied in the framework of Hilbert spaces and operators acting on them. We assume that the reader is familiar with the basic definitions and results concerning Hilbert spaces and bounded linear operators on them. Here we focus on closed unbounded linear operators with an emphasis on adjoints and duality. This is also a mature and well-developed theory, even if less studied. Most of the results here can be found in numerous references, such as [23, Chapter 2.6], [56, Chapter III, Section 5, and Chapter IV, Section 5.2], or [71, Chapter II, Section 6, and Chapter VII], often in the greater generality of Banach spaces or even more general topological vector spaces. However, for our applications, the Hilbert space context suffices and, moreover, it allows for a brief, but fairly self-contained, exposition. At the end of the chapter we return to the differential operators which arose in the last chapter as differentials of the de Rham complex and interpret them as closed unbounded operators on Hilbert space and compute their adjoints. A summary of the key points is given in Box 3.1 at the end of this chapter.

3.1 • Unbounded operators

When considering unbounded operators, one must allow for operators whose domain may not be the entire Hilbert space but only a subspace. Thus, by an *unbounded operator* T from a Hilbert space X to a Hilbert space Y, we mean a linear operator mapping a *subspace* D(T) of X, called the domain of T, into Y. The operator T is not necessarily bounded and the domain D(T) is not necessarily a closed subspace. Such operators are called *unbounded operators* on Hilbert space, even though the case of a bounded linear operator defined on all of X is included. The more accurate term *not-necessarily-everywhere-defined*, *not-necessarily-bounded linear operators* would be a bit unwieldy! We shall almost always be concerned with *densely defined* operators, meaning that the domain D(T) is a dense subspace of X. A typical example to keep in mind is $X = L^2(\Omega)$ for some domain $\Omega \subset \mathbb{R}^n$, $Y = L^2(\Omega; \mathbb{R}^n)$, the domain D(T) is the Sobolev space $H^1(\Omega)$, which is a dense subspace of $L^2(\Omega)$, and the operator T is the gradient. That is, $T\phi = \operatorname{grad} \phi \in L^2(\Omega; \mathbb{R}^n)$ for all $\phi \in H^1(\Omega)$. A variation is the operator \tilde{T} with $D(\tilde{T}) = \mathring{H}^1(\Omega)$, which is also dense in L^2 , again with $\tilde{T}\phi = \operatorname{grad} \phi$, but with $\tilde{T}\phi$ only defined for $\phi \in \mathring{H}^1$. Other possibilities would be to define a gradient operator whose domain consists of C^1 functions on $\bar{\Omega}$, or C^∞ functions, or C^∞ functions with compact support inside Ω . Although we might refer to any one of these operators as the gradient, they are different, and it is important to always be clear about the domain. If we say that two unbounded operators S and T are equal, we mean both that D(S) = D(T) and that the actions of the operators coincide.

Two unbounded operators S and T from X to Y can be added to give an operator S + T, with domain $D(S + T) = D(S) \cap D(T)$. If S is an unbounded operator $X \to Y$ and T an unbounded operator $Y \to Z$, then we can similarly define the composition $T \circ S$ with domain $\{v \in D(S) | Sv \in D(T)\}$. One must be careful, since the sum or composition of densely defined operators need not be densely defined.

The *null space* of an unbounded operator T, the *range* of T, and the *graph* of T are each defined in the obvious way:

$$\mathcal{N}(T) = \{ x \in D(T) \mid Tx = 0 \} \subset X,$$

$$\mathcal{R}(T) = \{ Tx \mid x \in D(T) \} \subset Y,$$

$$\Gamma(T) = \{ (x, Tx) \mid x \in D(T) \} \subset X \times Y.$$

The function $v \mapsto (v, Tv)$ is a one-to-one mapping from the domain of T onto its graph. Applying the norm in $X \times Y$ to (v, Tv) defines a norm on D(T). This norm, called the *graph norm* on D(T), is given by

$$||v||_{D(T)}^2 := ||v||_X^2 + ||Tv||_Y^2, \quad v \in D(T),$$

and is associated to the graph inner product

$$\langle v, w \rangle_{D(T)} := \langle v, w \rangle_X + \langle Tv, Tw \rangle_Y, \quad v, w \in D(T).$$

Obviously the graph norm is stronger than the restriction to T(D) of the X norm, and commonly they are not equivalent. For example, in the example given above of the gradient operator, the graph norm is the usual H^1 norm, which is strictly stronger than the L^2 norm. We always have

$$||Tv||_Y \le ||v||_{D(T)}, \quad v \in D(T),$$

so the operator T is a bounded operator from D(T) to Y when the former is endowed with the graph norm. The graph norm makes D(T) an inner product space, or pre-Hilbert space, but it may not be complete.¹

3.2 • Closed operators

Many important properties of bounded linear operators extend to a much larger class of unbounded operators, called closed. Banach's closed graph theorem states than an everywhere defined linear operator between Hilbert (or Banach) spaces is bounded if and only if its graph is closed. This inspires the definition of a *closed*

¹We could have instead defined the graph norm by $||v||_{D(T)} = ||v||_X + ||Tv||_Y$ without the squares. This definition gives an equivalent norm, but it would no longer be associated to an inner product.

operator, which is any unbounded operator whose graph is closed in $X \times Y$. In other words, an unbounded operator T from X to Y is a closed linear operator if and only if the following condition holds:

Whenever $v_1, v_2, \ldots \in D(T)$ satisfy $v_n \xrightarrow{X} x$ and $Tv_n \xrightarrow{Y} y$ for some $x \in X$ and $y \in Y$, then $x \in D(T)$ and Tx = y.

The following proposition gives another, equivalent definition of a closed operator.

Proposition 3.1. An unbounded operator T is closed if and only if its domain, endowed with the graph norm, is a Hilbert space.

Proof. We must show that D(T) is complete if and only if T is closed. Let $\{v_n\}_{n\in\mathbb{N}}$ be a sequence of elements of D(T). To say the sequence is Cauchy with respect to the graph norm means that it is Cauchy with respect to the X norm and $\{Tv_n\}$ is Cauchy with respect to the Y norm. Since X and Y are complete, this implies that $v_n \xrightarrow{X} x$ and $Tv_n \xrightarrow{Y} y$ for some $x \in X, y \in Y$. If T is closed, then $x \in D(T)$ and Tx = y, so $v_n \xrightarrow{D(T)} x$. Thus D(T) is complete.

Conversely, suppose D(T) is complete with respect to the graph norm and that $v_n \xrightarrow{X} x$ and $Tv_n \xrightarrow{Y} y$. Then the sequence v_n is Cauchy with respect to the graph norm on D(T), so $v_n \xrightarrow{D(T)} v$ for some $v \in D(T)$, i.e., $v_n \xrightarrow{X} v$ and $Tv_n \xrightarrow{Y} Tv$. Thus $x = v \in D(T)$ and y = Tv = Tx, so T is closed.

3.3 • The adjoint of an unbounded operator

Now suppose that T is any densely defined linear operator from X to Y. Then we may define an unbounded linear operator from Y to X called the *adjoint* of T and denoted by T^* . To this end we consider, for any $w \in Y$, the linear mapping $D(T) \to \mathbb{R}$ given by

$$v \mapsto \langle w, Tv \rangle_Y, \quad v \in D(T).$$
 (3.1)

The domain $D(T^*)$ of the adjoint is defined as the set of $w \in Y$ such that the linear functional (3.1) is bounded in X norm, i.e., for which there exists a constant c_w such that

$$|\langle w, Tv \rangle_Y| \le c_w ||v||_X, \quad v \in D(T).$$

For any such $w \in D(T^*)$ there is a unique extension of the map (3.1) to a bounded linear functional from all of X to \mathbb{R} . Moreover, by the Riesz representation theorem, there then exists a unique element $x \in X$ such that

$$\langle w, Tv \rangle_Y = \langle x, v \rangle_X, \quad v \in D(T),$$
(3.2)

and we set $T^*w = x$. In short,

$$\langle T^*w, v \rangle_X = \langle w, Tv \rangle_Y, \quad v \in D(T),$$

and $w \in D(T^*)$ exactly when there exists an element of X (necessarily unique) which, when taken for T^*w , fulfills this equation.

Many properties of the adjoint follow from a simple relation between the graph of T and the graph of its adjoint T^* . Since the latter lives in $Y \times X$ rather than $X \times Y$, we define a variant of the graph of T^* , the rotated graph, by

$$\tilde{\Gamma}(T^*) = \{ (-x, y) \, | \, (y, x) \in \Gamma(T^*) \, \} = \{ (-T^*w, w) \, | \, w \in D(T^*) \, \} \subset X \times Y,$$

where the minus sign has been inserted to make the relationship more evident. In fact, the rotated graph of T^* is simply $\Gamma(T)^{\perp}$, the orthogonal complement in $X \times Y$ of the graph $\Gamma(T)$.

Proposition 3.2. Let T be a densely defined unbounded operator. Then

$$\Gamma(T)^{\perp} = \tilde{\Gamma}(T^*), \qquad (3.3)$$

$$\overline{\Gamma(T)} = \tilde{\Gamma}(T^*)^{\perp}.$$
(3.4)

Proof. An element (x, w) of $X \times Y$ belongs to $\Gamma(T)^{\perp}$ if and only if

$$\langle x, v \rangle_X + \langle w, Tv \rangle_Y = 0, \quad v \in D(T),$$

which is exactly the condition (3.2) for w to belong to $D(T^*)$ with $T^*w = -x$, i.e., for (x, w) to belong to $\tilde{\Gamma}(T^*)$. This proves (3.3). Taking orthogonal complements then gives (3.4).

As an immediate application, we see that the adjoint T^* of a densely defined unbounded operator T is always closed, whether T is closed or not. Moreover, the original operator T is closed if and only if $\Gamma(T) = \tilde{\Gamma}(T^*)^{\perp}$, since $\overline{\Gamma(T)} = \Gamma(T)$ exactly when T is closed.

As another application of Proposition 3.2, we show that the adjoint operator of a *closed* densely defined operator is itself densely defined.

Proposition 3.3. If T is a closed densely defined unbounded operator from X to Y, then T^* is a closed densely defined operator from Y to X.

Proof. To show that T^* is densely defined, we must prove that any $y \in Y$ which is orthogonal to $D(T^*)$ must vanish. Now

$$y \perp D(T^*) \iff (0, y) \in \Gamma(T^*)^{\perp} = \Gamma(T),$$

with the last equality coming from Proposition 3.2 and the assumption that T is closed. Thus, if $y \perp D(T^*)$, then y = T0 and so indeed vanishes.

In view of this result, when T is a closed densely defined operator, then so is T^* , and we may define T^{**} , which is again a closed densely defined operator $X \to Y$. In fact, T^{**} coincides with T.

Proposition 3.4. If T is a closed densely defined unbounded operator from X to Y, then $T^{**} = T$.

Proof. Since T is closed, $\Gamma(T) = \tilde{\Gamma}(T^*)^{\perp}$ and, similarly, $\tilde{\Gamma}(T^*) = \Gamma(T^{**})^{\perp}$. Combining these we get $\Gamma(T) = \Gamma(T^{**})^{\perp \perp}$, which equals $\Gamma(T^{**})$, since T^{**} is closed. Since T and T^{**} have the same graphs, they are equal. If T is a closed operator, then its null space $\mathcal{N}(T)$ is a closed subspace of X. This follows directly from the definition of a closed operator. However, the range of a closed linear operator may well not be closed, even in the case of a bounded linear operator (consider, for example, the inclusion of H^1 into L^2). The following theorem summarizes the relationship between the null spaces and the ranges of T and T^* . Its proof is another application of Proposition 3.2.

Theorem 3.5. Let T be a closed densely defined operator X to Y. Then

 $\mathcal{R}(T)^{\perp} = \mathcal{N}(T^*), \ \mathcal{N}(T)^{\perp} = \overline{\mathcal{R}(T^*)}, \ \mathcal{R}(T^*)^{\perp} = \mathcal{N}(T), \ \mathcal{N}(T^*)^{\perp} = \overline{\mathcal{R}(T)}.$

Proof. We verify the third equation first:

$$z \in \mathcal{N}(T) \iff (z,0) \in \Gamma(T) = \tilde{\Gamma}(T^*)^{\perp}$$
$$\iff \langle z, T^* w \rangle_X = 0 \; \forall w \in D(T^*) \iff z \in \mathcal{R}(T^*)^{\perp}.$$

The third equation then implies the first equation by replacing T by T^* . Taking orthogonal complements of the first and third equations then gives the fourth and second equations, respectively.

Although the range of a closed linear operator T may not be closed, if it is, then so also is the range of T^* (and vice versa, since we can apply this result with T replaced by T^*). This is a consequence of Banach's closed range theorem. A proof may be found in the more general setting of closed operators on Banach spaces, for example, in [71, Chapter VII.5]. For the convenience of the reader we include a proof in the Hilbert space case here. First we state a simple lemma.

Lemma 3.6. Let T be a closed operator between Hilbert spaces X and Y and suppose that there exists $\gamma > 0$ such that

$$||Tv||_Y \ge \gamma ||v||_X, \quad v \in D(T).$$

Then T has a closed range.

Proof. If $\{v_n\}$ is a sequence in D(T) such that Tv_n converges to some y in Y, we must show that $y \in \mathcal{R}(T)$. The hypothesis implies that the sequence $\{v_n\}$ is Cauchy and so converges to some x in X. Since T is closed, $x \in D(T)$ and $y = Tx \in \mathcal{R}(T)$. \Box

Theorem 3.7 (closed range theorem). *If the range of a closed densely defined operator* T *from* X *to* Y *is closed in* Y*, then the range of* T^* *is closed in* X*.*

Proof. Let $W = \mathcal{R}(T)$, a closed subspace of Y, and let \overline{T} be the range restriction of T, i.e., the unbounded operator from X to W with the same domain as T with $\overline{T}v = Tv$, $v \in D(T)$. Obviously \overline{T} is also a closed densely defined operator between Hilbert spaces, and moreover it is a surjection. We shall show that $\mathcal{R}(\overline{T}^*) = \mathcal{R}(T^*)$. This follows from two observations. First, note that $D(\overline{T}^*) = D(T^*) \cap W$ and \overline{T}^* is the restriction of T^* to this space. This follows
directly from the definition of the adjoint operator and the fact that $\overline{T}v = Tv$ for all v in $D(T) = D(\overline{T})$. Thus $\mathcal{R}(\overline{T}^*) \subset \mathcal{R}(T^*)$. But $W^{\perp} = \mathcal{N}(T^*)$. Therefore, if $y \in D(T^*)$ and w is its orthogonal projection onto W, then $w \in D(\overline{T}^*)$ and $\overline{T}^*w = T^*y$, showing that $\mathcal{R}(T^*) \subset \mathcal{R}(\overline{T}^*)$.

In view of this, it suffices to show that if T is a closed densely defined *surjec*tion, then $\mathcal{R}(T^*)$ is closed. We know that D(T) is a Hilbert space when endowed with the graph norm. Let Z be the orthogonal complement of $\mathcal{N}(T)$ in D(T) and let $\tilde{T}: Z \to Y$ be the restriction of T. Then \tilde{T} is a bounded linear operator which is one-to-one and onto. It follows that it is invertible, so there exists a positive constant c such that for any $y \in Y$ there exists a unique $z \in Z$ with Tz = y and $\|z\|_X \leq c\|y\|_Y$. Applying this for $y \in D(T^*)$ we have

$$||y||_Y^2 = \langle y, Tz \rangle_Y = \langle T^*y, z \rangle_X \le ||T^*y||_X ||z||_X \le c ||y||_Y ||T^*y||_X \le c ||T^*y$$

Thus

$$||T^*y||_X \ge c^{-1} ||y||_Y, \quad y \in D(T^*).$$

Now we apply the preceding lemma to T^* to conclude that $\mathcal{R}(T^*)$ is closed. \Box

We conclude this section with some useful results for verifying that the range of an operator is closed. First we show that if the range of a closed linear operator is finite codimensional, then it is closed. A proof in the case of bounded operators is given in [52, Lemma 19.1.1], and it extends easily to closed operators.

Theorem 3.8. Let T be a closed linear operator between Hilbert spaces X and Y. If the quotient space $Y/\mathcal{R}(T)$ is finite dimensional, then the range of T is closed in Y.

Proof. First consider the case where T is a bounded linear operator from X to Y. Restricting T to the orthogonal complement of its null space, we may assume that T is injective. Let $y_1, \ldots, y_n \in Y$ be such that their cosets modulo $\mathcal{R}(T)$ are a basis for $Y/\mathcal{R}(T)$. Then every element of Y may be written as $r+c_1y_1+\cdots c_ny_n$ with $r \in \mathcal{R}(T)$ and $c_i \in \mathbb{R}$ in a unique way. Thus if we define $\tilde{T} : X \times \mathbb{R}^n \to Y$ by

$$T(x,c_1,\ldots,c_n) = Tx + c_1y_1 + \cdots + c_ny_n,$$

then T is a bounded linear operator which is one-to-one and onto and hence, by Banach's theorem, a homeomorphism. Since $X \times \{0\}$ is closed in $X \times \mathbb{R}^n$, it follows that $\mathcal{R}(T) = \tilde{T}(X \times \{0\})$ is closed in Y.

In the general case of a closed linear operator T from X to Y, we may view T as a bounded linear operator from D(T) to Y, and so, from the previous case, if its range is finite codimensional, it is closed.

Our second criterion for closed range is based on compactness.

Theorem 3.9. Let T be a closed linear operator between Hilbert spaces X and Y, and suppose that the inclusion of its domain into X is compact (where the domain is given the graph norm). Then $\mathcal{R}(T)$ is closed in Y.

Proof. Given a sequence of elements $x_n \in D(T)$ and $y \in Y$ such that $Tx_n \to y$ in Y, we must show that $y \in \mathcal{R}(T)$. Without loss of generality we may take

 $x_n \in \mathcal{N}(T)^{\perp}$. In this case, the sequence x_n is bounded in X, as we now show by contradiction. Indeed, if it were not, we could pass to a subsequence, which we continue to denote x_n , with $||x_n||_X \to \infty$. Defining $w_n = x_n/||x_n||_X \in D(T) \cap \mathcal{N}(T)^{\perp}$ we would thus have a sequence with $||w_n||_X = 1$ and $Tw_n \to 0$. Such a sequence is bounded in D(T), and so, by compactness, we could pass to a subsequence which converges in the norm of X to some $w \in X$. Since T is closed, the resulting w would belong to D(T) and satisfy Tw = 0, i.e., would belong to $\mathcal{N}(T)$. However, as the limit of the w_n , w would also belong to $\mathcal{N}(T)^{\perp}$ and be of norm 1. These conditions are in clear contradiction.

Thus it suffices to show that for a bounded sequence of elements $x_n \in D(T)$ with $\lim Tx_n = y$ in Y, y belongs to $\mathcal{R}(T)$. We again invoke the compactness hypothesis to replace the sequence with a subsequence for which, in addition, x_n converges to some x in X. Since T is closed we have $x \in D(T)$ and Tx = y, as desired. \Box

3.4 • Example: grad, curl, and div and their adjoints

In this section we consider the three vector calculus operators which enter the de Rham complex, showing how to view them as closed unbounded operators on Hilbert space, and deriving their basic properties in that framework.

Let Ω denote a bounded domain in \mathbb{R}^n . We start by considering the unbounded operator from $L^2(\Omega; \mathbb{R}^n)$ to $L^2(\Omega)$ with domain $C_0^{\infty}(\Omega; \mathbb{R}^n)$ given by $\phi \mapsto -\operatorname{div} \phi$ for $\phi \in C_0^{\infty}(\Omega; \mathbb{R}^n)$. We shall write this operator as $(-\operatorname{div}, C_0^{\infty})$ to make clear its domain. Although it is not a closed operator, it is clearly densely defined, so its adjoint $(-\operatorname{div}, C_0^{\infty})^*$ is a closed operator $L^2(\Omega) \to L^2(\Omega; \mathbb{R}^n)$. By definition, a real-valued function u on Ω belongs to the domain of this adjoint if it belongs to $L^2(\Omega)$ and there exists a vector field v in $L^2(\Omega; \mathbb{R}^n)$ for which

$$-\int_{\Omega} u \operatorname{div} \phi \, dx = \int_{\Omega} v \cdot \phi \, dx, \quad \phi \in C_0^{\infty}(\Omega; \mathbb{R}^n).$$

But this equation is exactly the statement that the distributional gradient of u, grad u, belongs to L^2 (namely, grad $u = v = (-\operatorname{div}, C_0^{\infty})^* u$). Thus the domain of $(-\operatorname{div}, C_0^{\infty})^*$ is the set of all L^2 functions with distributional gradient in L^2 , which is precisely the definition of the Sobolev space $H^1(\Omega)$. We have shown that

$$(-\operatorname{div}, C_0^\infty)^* = (\operatorname{grad}, H^1).$$

The operator (grad, H^1) not only is closed, but it is evidently also densely defined (its domain contains C_0^{∞} , for example, which is already dense in L^2). The fact that it is closed establishes the completeness of H^1 .

We now reverse the roles of div and grad and consider the unbounded operator $(-\operatorname{grad}, C_0^{\infty})$ from $L^2(\Omega)$ to $L^2(\Omega; \mathbb{R}^n)$. Reasoning just as above,

$$(-\operatorname{grad}, C_0^\infty)^* = (\operatorname{div}, H(\operatorname{div}))$$

where the operator on the right-hand side is closed and densely defined. Its domain is the Hilbert space $H(\text{div}, \Omega)$, consisting of L^2 vector fields with distributional divergence in L^2 . Our next goal is to compute the adjoints of each of the two operators (grad, H^1) and (div, H(div)) just derived. The key to this is the integration-by-parts formula

$$\int u \operatorname{div} w \, dx + \int \operatorname{grad} u \cdot w \, dx = \int_{\partial \Omega} u|_{\partial \Omega} w|_{\partial \Omega} \cdot n \, ds.$$
(3.5)

This certainly holds for $u \in C^1(\overline{\Omega})$, $w \in C^1(\overline{\Omega}; \mathbb{R}^n)$ and a smooth boundary, but it is not clear that quantities like the trace of w on $\partial\Omega$ make sense for w only belonging to H(div). So first we investigate that issue. The key result is given in Theorem 3.12.

For this, and the remainder of the section, we suppose that Ω has a Lipschitz boundary. We begin by recalling a few elementary facts about the Sobolev space H^1 . First, the space $C^{\infty}(\bar{\Omega})$ is dense in $H^1(\Omega)$ [41, Section 5.3, Theorem 2]. Second, the restriction operator $C^1(\bar{\Omega}) \to L^2(\partial\Omega)$ extends to a bounded trace operator tr : $H^1(\Omega) \to L^2(\partial\Omega)$ [41, Section 5.5, Theorem 1]. The range of the trace operator is $H^{1/2}(\partial\Omega)$ and, in fact, this is one possible way to define that space. In this approach, the norm in $H^{1/2}(\partial\Omega)$ is given by

$$\|g\|_{1/2,\partial\Omega} = \inf_{\text{tr}\, u=g} \|u\|_{1,\Omega}.$$
(3.6)

Third, the kernel of the trace operator is the space $\mathring{H}^1(\Omega)$, defined as the closure of $C_0^{\infty}(\Omega)$ in $H^1(\Omega)$. See [41, Section 5.5, Theorem 2]. Knowing both the range and the kernel of the trace operator acting on $H^1(\Omega)$, we can use duality to characterize the functionals on $H^1(\Omega)$ which vanish on the kernel $\mathring{H}^1(\Omega)$. The following result is elementary.

Lemma 3.10. If $\gamma : X \to Y$ is a bounded linear surjection between normed linear spaces with kernel Z, the dual map $\gamma' : Y' \to X'$ is a bounded injection with range equal to the annihilator of Z (defined as $\{L \in X' \mid L|_Z \equiv 0\}$). Thus, for L in the annihilator there exists a unique $g \in Y'$ such that

$$Lu = \langle g, \gamma u \rangle, \quad u \in X.$$

Applying this lemma with γ equal to the trace operator tr : $H^1(\Omega) \rightarrow H^{1/2}(\partial\Omega)$ and defining $H^{-1/2}(\partial\Omega)$ as the dual space of $H^{1/2}(\partial\Omega)$, we obtain the following characterization.

Proposition 3.11. If $L : H^1(\Omega) \to \mathbb{R}$ is a bounded linear functional which annihilates $\mathring{H}^1(\Omega)$, then there exists a unique $g \in H^{-1/2}(\partial\Omega)$ for which

$$Lu = \langle g, \operatorname{tr} u \rangle, \quad u \in H^1(\Omega).$$
(3.7)

Moreover, $||g||_{H^{-1/2}(\partial\Omega)} = ||L||_{H^1(\Omega)'}$.

Proof. The existence of a unique g satisfying (3.7) comes directly from the lemma. From (3.7) and (3.6), we have

$$|Lu| \le ||g||_{H^{-1/2}(\partial\Omega)} ||\operatorname{tr} u||_{H^{1/2}(\partial\Omega)} \le ||g||_{H^{-1/2}(\partial\Omega)} ||u||_{H^{1}(\Omega)}$$

so $||L|| \leq ||g||_{H^{-1/2}(\partial\Omega)}$. For the reverse inequality, let $f \in H^{1/2}(\partial\Omega)$ be arbitrary and choose $u \in H^1(\Omega)$ such that $\operatorname{tr} u = f$ and $||u||_1 \leq ||f||_{H^{1/2}(\partial\Omega)}$. Then

$$\langle g, f \rangle = Lu \le \|L\| \, \|u\|_1 \le \|L\| \, \|f\|_{H^{1/2}(\partial\Omega)}.$$

Next we apply this result to define the trace of $w \in H(\operatorname{div}, \Omega)$. The functional

$$L_w(u) = \int u \operatorname{div} w \, dx + \int \operatorname{grad} u \cdot w \, dx, \quad u \in H^1(\Omega),$$

is a bounded linear functional on H^1 and, by the definition of the distributional derivative, $L_w(u)$ vanishes for $u \in C_0^{\infty}(\Omega)$ and hence, by density, for $u \in \mathring{H}^1(\Omega)$. We may thus apply the proposition to obtain a unique $\gamma_n w \in H^{-1/2}(\partial\Omega)$ for which

$$\int u \operatorname{div} w \, dx + \int \operatorname{grad} u \cdot w \, dx = \langle \gamma_n w, \operatorname{tr} u \rangle, \quad u \in H^1(\Omega), \ w \in H(\operatorname{div}).$$
(3.8)

Moreover,

$$\|\gamma_n w\|_{-1/2,\partial\Omega} = \sup_u \frac{\int u \,\mathrm{div} \, w \,dx + \int \operatorname{grad} u \cdot w \,dx}{\|u\|_{1,\Omega}} \le \|w\|_{H(\mathrm{div})}.$$

To understand the meaning of $\gamma_n w$, we compare (3.8) with the integration-byparts formula (3.5). Fixing $w \in C^1(\overline{\Omega}; \mathbb{R}^3)$, note that, as a function of $u \in C^1(\overline{\Omega})$, each of the three integrals in (3.5) is bounded with respect to the H^1 norm. The density of smooth functions in H^1 then implies (3.5) for all $u \in H^1(\Omega)$. However, $\gamma_n w$ is uniquely determined by (3.8), so we have

$$\gamma_n w = w|_{\partial\Omega} \cdot n$$

for $w \in C^1(\overline{\Omega}; \mathbb{R}^3)$. We have established the following theorem.

Theorem 3.12. The operator $w \mapsto w|_{\partial\Omega} \cdot n$ on $C^1(\overline{\Omega}; \mathbb{R}^3)$ extends to a bounded linear operator $\gamma_n : H(\operatorname{div}) \to H^{-1/2}(\partial\Omega)$. The extension satisfies formula (3.8) for integration by parts and the bound $\|\gamma_n w\|_{-1/2,\partial\Omega} \leq \|w\|_{H(\operatorname{div})}$.

We are now able to compute the adjoint of the gradient operator with domain H^1 . For this, define

$$\dot{H}(\operatorname{div}) = \{ w \in H(\operatorname{div}) \mid \gamma_n w = 0 \}.$$

Theorem 3.13. The adjoint of the unbounded operator (grad, H^1) is the operator $(-\operatorname{div}, \mathring{H}(\operatorname{div}))$.

Proof. A vector field $w \in L^2(\Omega; \mathbb{R}^n)$ is in the domain of the adjoint if and only if there exists $v \in L^2(\Omega)$ such that

$$\int w \cdot \operatorname{grad} u \, dx = \int v u \, dx, \quad u \in H^1(\Omega).$$

In that case, clearly we must have $w \in H(\operatorname{div})$ and $v = -\operatorname{div} w$ (by considering $u \in C_0^\infty$). Thus w belongs to the domain of the adjoint if and only if $w \in H(\operatorname{div})$ and

$$\int w \cdot \operatorname{grad} u \, dx = -\int (\operatorname{div} w) u \, dx, \quad u \in H^1(\Omega).$$

Comparing with the definition (3.8) of $\gamma_n w$, we see that w belongs to the domain of the adjoint if and only if $\gamma_n w = 0$.

We next show that the normal trace operator maps H(div) not only into $H^{-1/2}(\partial\Omega)$ but onto.

Theorem 3.14. The operator $\gamma_n : H(\operatorname{div}) \to H^{-1/2}(\partial\Omega)$ is surjective.

Proof. Given $g \in H^{-1/2}(\partial \Omega)$, define $u \in H^1$ by the Neumann problem

$$\int \operatorname{grad} u \cdot \operatorname{grad} v \, dx + \int uv \, dx = \langle g, \operatorname{tr} v \rangle, \quad v \in H^1.$$

Then $-\Delta u + u = 0$. Take $w = \operatorname{grad} u$. Then $\operatorname{div} w = \Delta u = u$, and so $w \in H(\operatorname{div})$. We have

$$\int w \cdot \operatorname{grad} v \, dx + \int (\operatorname{div} w) v \, dx = \langle g, \operatorname{tr} v \rangle, \quad v \in H^1,$$

which shows that $\gamma_n w = g$.

Using this, we can argue exactly as in the proof of Theorem 3.13, to prove the following result, which establishes the adjoint of the divergence operator.

Theorem 3.15. *The adjoint of the unbounded operator* (div, H(div)) *is the operator* $(-\operatorname{grad}, \mathring{H}^1)$.

Since $T^{**} = T$ for a closed densely defined operator, from Theorems 3.13 and 3.15 we can read off two more adjoints.

Theorem 3.16.

$$(\operatorname{div}, \check{H}(\operatorname{div}))^* = (-\operatorname{grad}, H^1), \qquad (\operatorname{grad}, \check{H}^1)^* = (-\operatorname{div}, H(\operatorname{div}))$$

Next we verify that all of these operators have closed range. For $(\operatorname{div}, H(\operatorname{div}))$ this is clear, since in fact $\operatorname{div} H(\operatorname{div}) = L^2$, as can be seen by writing $f \in L^2$ as $\Delta u = \operatorname{div} \operatorname{grad} u$ for some $u \in H^1$. If $\int f \, dx = 0$, then we can choose u satisfying the Neumann condition $\operatorname{grad} u \cdot n = 0$, and so $f \in \operatorname{div} \mathring{H}(\operatorname{div})$. Thus the range $\operatorname{div} \mathring{H}(\operatorname{div})$ has codimension 1 in L^2 , and so this is closed as well, by Theorem 3.8. The closed range theorem then implies that the adjoints of these two maps, namely, $(\operatorname{grad}, \mathring{H}^1)$ and $(\operatorname{grad}, H^1)$, both have closed range as well.

Finally, we turn to the curl operator, assuming the domain $\Omega \subset \mathbb{R}^3$. Most of the arguments we applied above to grad and div apply to curl as well, but not all, because neither the kernel nor the range of curl is finite codimensional. For the curl, the relevant integration-by-parts formula is

$$\int v \cdot \operatorname{curl} u \, dx - \int \operatorname{curl} v \cdot u \, dx = \int_{\partial \Omega} v|_{\partial \Omega} \times n \cdot u|_{\partial \Omega} \, ds, \quad u, v \in C^1(\bar{\Omega}; \mathbb{R}^3),$$
(3.9)

where *n* is the unit outward normal (which is defined almost everywhere since the boundary is Lipschitz). The space H(curl) is defined to be the space of vector fields in $L^2(\Omega; \mathbb{R}^3)$ whose curl taken in the sense of distributions also belongs to $L^2(\Omega; \mathbb{R}^3)$. Thus, a vector field $u \in L^2(\Omega; \mathbb{R}^3)$ belongs to H(curl) if and only if there exists a vector field $v \in L^2(\Omega; \mathbb{R}^3)$ such that

$$\int_{\Omega} u \cdot \operatorname{curl} w \, dx = \int_{\Omega} v \cdot w \, dx, \quad w \in C_0^{\infty}(\Omega; \mathbb{R}^3),$$
(3.10)

and, in this case, $v = \operatorname{curl} u$. Thus, just as above,

$$(\operatorname{curl}, C_0^{\infty})^* = (\operatorname{curl}, H(\operatorname{curl})).$$

Being an adjoint, this last is a closed operator, which means that H(curl) is a Hilbert space when endowed with the graph norm

$$||u||_{H(\operatorname{curl})}^2 = ||u||_{L^2}^2 + ||\operatorname{curl} u||_{L^2}^2.$$

Now we develop a trace for $v \in H(\text{curl})$. The map

$$u \mapsto \int v \cdot \operatorname{curl} u \, dx - \int \operatorname{curl} v \cdot u \, dx, \quad u \in H^1(\Omega; \mathbb{R}^3),$$

is a bounded linear functional on $H^1(\Omega; \mathbb{R}^3)$ which vanishes on $C_0^{\infty}(\Omega; \mathbb{R}^3)$ and so, by density, on $\mathring{H}^1(\Omega; \mathbb{R}^3)$. Therefore we can apply Proposition 3.11 to infer the existence of a unique $\gamma_\tau v \in H^{-1/2}(\partial\Omega; \mathbb{R}^3)$ for which

$$\int v \cdot \operatorname{curl} u \, dx - \int \operatorname{curl} v \cdot u \, dx = \langle \gamma_{\tau} v, \operatorname{tr} u \rangle, \quad u \in H^1(\Omega; \mathbb{R}^3), \, v \in H(\operatorname{curl}),$$
(3.11)

and to obtain the norm bound

$$\|\gamma_{\tau}v\|_{-1/2,\partial\Omega} \le \|v\|_{H(\operatorname{curl})}.$$
 (3.12)

Comparing with (3.9) we see that we have extended the *tangential trace* operator $\gamma_{\tau} v := v|_{\partial\Omega} \times n$ to a bounded linear map on H(curl) with range belonging to the Sobolev space of $H^{-1/2}(\partial\Omega; \mathbb{R}^3)$ of vector fields on $\partial\Omega$, and so we have established the analogue of Theorem 3.12 for H(curl).

Theorem 3.17. The operator $v \mapsto v|_{\partial\Omega} \times n$ on $C^1(\overline{\Omega}; \mathbb{R}^3)$ extends to a bounded linear operator $\gamma_{\tau} : H(\text{curl}) \to H^{-1/2}(\partial\Omega; \mathbb{R}^3)$. The extension satisfies the integration-by-parts formula (3.11) and the bound (3.12).

Note, however, that the trace operator γ_{τ} is *not* onto $H^{-1/2}(\partial\Omega; \mathbb{R}^3)$, i.e., the analogue of Theorem 3.14 does not hold. For one thing, $\gamma_{\tau}w$ is always a tangential vector field. This is clear at least if v and $\partial\Omega$ are smooth, while the precise way to define this in the case of a merely Lipschitz boundary is subtle. There are further restrictions on the range of γ_{τ} as well. See [25, 26] for details.

Finally, we compute the adjoint of (curl, H(curl)). If w belongs to the domain of the adjoint, then there exists $v = (\text{curl}, H(\text{curl}))^* w \in L^2(\Omega; \mathbb{R}^3)$ with

$$\int w \cdot \operatorname{curl} u \, dx = \int v \cdot u \, dx, \quad u \in H(\operatorname{curl})$$

Such a w must belong to H(curl) and have curl w = v. Thus w belongs to the domain of $(\text{curl}, H(\text{curl}))^*$ if and only if $w \in H(\text{curl})$ and

$$\int w \cdot \operatorname{curl} u \, dx = \int \operatorname{curl} w \cdot u \, dx, \quad u \in H(\operatorname{curl})$$

It can be shown (by a mollification argument similar to that for H^1) that smooth functions are dense in H(curl). Thus we need only require that the last equation

Box 3.1. Summary of Chapter 3 on operators on Hilbert spaces.

An *(unbounded) operator* from one Hilbert space X to another Y is a linear operator mapping its domain D(T), a subspace of X, into Y. The domain need not be closed in X and the operator need not be bounded, but usually we require that the operator be *densely defined*, i.e., that D(T) is dense in X.

An operator is *closed* if its graph is closed in $X \times Y$ or, equivalently, if D(T) is complete when furnished with the *graph norm*. An everywhere defined operator is closed if and only if it is bounded.

A densely defined operator T from X to Y gives rise to an adjoint operator T^* from Y to X. The adjoint T^* is always closed. If T is closed, then T^* is also densely defined, and, in this case, $T^{**} = T$.

The range of a closed, densely defined linear operator $X \to Y$ need not be a closed subspace, but the range of T is closed in Y if and only if the range of T^* is closed in X. If dim $Y/\mathcal{R}(T) < \infty$, then T has closed range. If the inclusion $D(T) \subset X$ is compact, then T has closed range.

On a bounded domain in \mathbb{R}^3 with Lipschitz boundary, the operators grad, curl, and div may be viewed as unbounded operators on L^2 (scalars or vectors), with domains H^1 , H(curl), H(div), respectively. They are densely defined, closed, and have closed ranges. Their adjoints are - div, curl, and - grad with domains $\mathring{H}(\text{div})$, $\mathring{H}(\text{curl})$, and \mathring{H}^1 , respectively.

hold for all $u \in H^1(\Omega; \mathbb{R}^3)$. But then (3.11) implies that this holds if and only if $\gamma_{\tau} w = 0$, and so we have established that the adjoint of (curl, H(curl)) is (curl, \dot{H} (curl)). Finally, we can take adjoints and get the reverse relation. These results are summarized in the following theorem.

Theorem 3.18.

 $(\operatorname{curl}, H(\operatorname{curl}))^* = (\operatorname{curl}, \mathring{H}(\operatorname{curl})), \qquad (\operatorname{curl}, \mathring{H}(\operatorname{curl}))^* = (\operatorname{curl}, H(\operatorname{curl})),$

where

$$\check{H}(\operatorname{curl}) := \{ w \in H(\operatorname{curl}) \, | \, \gamma_{\tau} w = 0 \}.$$

In Section 4.3 we show that the range $\operatorname{curl} H(\operatorname{curl})$ is closed.

Chapter 4 Hilbert complexes

We have just seen that the differential operators entering the de Rham complex can be viewed as closed unbounded operators between L^2 spaces. In this way we obtain a structure which combines the homological algebraic features of a cochain complex, discussed in Chapter 2, with the functional analytic features of closed unbounded operators on Hilbert space, discussed in Chapter 3. This structure is called a *Hilbert complex* and was first promulgated by Brüning and Lesch [24]. It is the structure underlying Hodge theory of Riemannian manifolds,² and, despite its simplicity, the main results of Hodge theory, for example, the identification of cohomology with harmonic forms and the Hodge decomposition, follow easily in this framework. Associated to each Hilbert complex is a graded map, or sequence of operators, called the abstract Hodge Laplacian of the complex. We explore several different but equivalent formulations of the Hodge Laplacian and demonstrate their well-posedness up to harmonic forms. The case of the de Rham complex on a three-dimensional domain is treated in detail, and the results of the chapter are summarized in Box 4.1 at the end of the chapter.

4.1 • Hilbert complexes and their duals

4.1.1 • Hilbert complexes

Definition 4.1. A Hilbert complex is a sequence of Hilbert spaces W^k and a sequence of closed densely defined linear operators d^k from W^k to W^{k+1} such that $\mathcal{R}(d^k) \subset \mathcal{N}(d^{k+1})$.

Note that a Hilbert complex is more than just a complex in which the spaces are Hilbert. At each level, there are two Hilbert spaces: the base space W^k and the domain space, which we shall denote by V^k . Since we assume that the differentials d^k are closed, the domain spaces V^k are Hilbert spaces with the graph norm

$$\|v\|_{V^k}^2 = \|v\|_{W^k}^2 + \|d^k v\|_{W^{k+1}}^2$$

 $^{^{2}}$ An indication of how natural is the notion of a Hilbert complex is that in the early drafts of [13] my coauthors and I introduced this definition, before we discovered that it already occured in the literature [24] with the very same name and definition.

The operators d^k may be viewed as *bounded* linear maps of V^k into V^{k+1} , and so they form a bounded Hilbert complex—a cochain complex in which the spaces are Hilbert spaces and the differentials are bounded linear operators—called the *domain complex*:

$$0 \to V^0 \xrightarrow{d} V^1 \xrightarrow{d} \cdots \xrightarrow{d} V^n \to 0. \tag{4.1}$$

However, the Hilbert complex (W, d) contains more information than its domain complex (V, d). From the former you can construct the latter, but knowing only the latter you cannot infer the spaces W^k or the W^k norm. Besides Hilbert cochain complexes, just defined, we can, of course, also talk about Hilbert chain complexes which differ only in that the indices decrease.

As we did previously for complexes, we denote the range and null space of the differential by $\mathfrak{B}^k = d^{k-1}V^{k-1}$ and $\mathfrak{Z}^k = \mathcal{N}(d^k)$, the spaces of coboundaries and cocycles. We have $\mathfrak{B}^k \subset \mathfrak{Z}^k$ and the quotients are the cohomology spaces $\mathcal{H}^k = \mathfrak{Z}^k/\mathfrak{B}^k$. The cycle spaces are always closed, but the boundary spaces may or may not be. A Hilbert complex is called *closed* if the range \mathfrak{B} is closed in W for all k. This holds if and only if \mathfrak{B} is complete in the W norm. Since the V norm on \mathfrak{Z} coincides with the W norm there, it is equivalent to require that \mathfrak{B} is complete with respect to the V norm, or that \mathfrak{B} is closed in V. A Hilbert complex is called *Fredholm* if the cohomology spaces are finite dimensional, i.e., the boundary spaces are finite codimensional inside the cycle spaces. Most of the complexes we encounter will be Fredholm. In view of Theorem 3.8, a Fredholm complex is closed.

For the few occasions when we work with Hilbert complexes which are not closed, or not known to be closed, we also introduce the *reduced cohomology* space $\mathfrak{Z}^k/\bar{\mathfrak{B}}^k$, which is a Hilbert space. Of course, for a closed Hilbert complex it coincides with the cohomology space \mathcal{H}^k .

4.1.2 • The dual complex

If (W, d) is a Hilbert complex, then d^{k-1} is a closed densely defined operator from W^{k-1} to W^k , so its adjoint, which we denote by d_k^* , is a closed densely defined operator from W^k to W^{k-1} . We denote its domain by V_k^* , a dense subset of W^k which is, in general, quite different from V^k . We set $\mathfrak{Z}_k^* = \mathcal{N}(d_k^*)$ and $\mathfrak{B}_k^* = \mathcal{R}(d_{k+1}^*)$. Thus Theorem 3.5 gives

$$\mathfrak{B}^{\perp}=\mathfrak{Z}^{*}, \quad \mathfrak{Z}^{\perp}=ar{\mathfrak{B}}^{*}, \quad \mathfrak{B}^{*\perp}=\mathfrak{Z}, \quad \mathfrak{Z}^{*\perp}=ar{\mathfrak{B}}.$$

Consequently,

$$\mathfrak{B}^* \subset \overline{\mathfrak{B}}^* = \mathfrak{Z}^\perp \subset \mathfrak{B}^\perp = \mathfrak{Z}^*,$$

so

$$0 \to V_n^* \xrightarrow{d_n^*} V_{n-1}^* \xrightarrow{d_{n-1}^*} \cdots \xrightarrow{d_1^*} V_0^* \to 0$$

is a chain complex, and (W, d^*) is a Hilbert complex. If the original Hilbert complex (W, d) is closed, then the dual complex (W, d^*) is closed as well, by the closed range theorem.

4.2 - Harmonic forms, the Hodge decomposition, and the Poincaré inequality

Next we derive several simple, but crucial, consequences of the Hilbert space structure. In this section, we assume that (W, d) is a Hilbert complex, which is not necessarily closed.

4.2.1 - Harmonic forms

Definition 4.2. An element $v \in W^k$ is called harmonic or a harmonic k-form if both dv = 0 and $d^*v = 0$.

We denote the space of harmonic k-forms by \mathfrak{H}^k , so $\mathfrak{H}^k = \mathfrak{Z}^k \cap \mathfrak{Z}_k^*$. By the first assertion of Theorem 3.5, we may as well write

$$\mathfrak{H}^k = \mathfrak{Z}^k \cap \mathfrak{B}^{k,\perp},\tag{4.2}$$

meaning that the harmonic forms are the cocycles which are orthogonal to coboundaries. If the coboundary space is not closed, we may close it and equivalently say that \mathfrak{H}^k is the orthogonal complement of the closed subspace $\overline{\mathfrak{B}}^k$ inside the Hilbert space \mathfrak{Z}^k . We may then canonically identify the quotient space, which is the reduced cohomology space $\mathfrak{Z}^k/\overline{\mathfrak{B}}^k$, with \mathfrak{H}^k , which is a subspace of W^k . Of course, for a closed Hilbert complex, we do not need to specify "reduced," and the result is simply

$$\mathfrak{H}^k \cong \mathcal{H}^k. \tag{4.3}$$

Thus the harmonic forms give a concrete realization of the cohomology of the complex as a closed subspace of the Hilbert space W^k . This is illustrated in Figure 4.1. The elements of the cohomology space \mathcal{H}^k are the equivalence classes



Figure 4.1. The space \mathfrak{B}^k of coboundaries is a subspace of the space \mathfrak{Z}^k of cocycles. The cohomology space consists of equivalence classes of cocycles, with the class [z] of a cocycle z being the affine subspace through z parallel to \mathfrak{B}^k . Each cohomology class [z] is represented by a unique element z_0 belonging to \mathfrak{H}^k , which is its element of its smallest norm.

of elements of \mathfrak{Z}^k modulo \mathfrak{B}^k . These provide a decomposition of \mathfrak{Z}^k into closed affine subspaces obtained from \mathfrak{B}^k by translation. Each such subspace contains a distinguished element, namely, the unique element orthogonal to \mathfrak{B}^k , which is also characterized as the unique element of the smallest norm. These are the harmonic *k*-forms.

We may use harmonic forms to obtain a useful criterion for a Hilbert complex to be Fredholm, i.e., for its homology spaces to be finite. The space $V^k \cap V_k^* \subset W^k$ is a Hilbert space when endowed with the norm

$$\|v\|_{V\cap V^*}^2 = \|v\|_V^2 + \|v\|_{V^*}^2 = 2\|v\|^2 + \|dv\|^2 + \|d^*v\|^2$$

Here and below, we write $\|\cdot\|_V$ for the norm in each of the spaces V^k , $\|\cdot\|_{V^*}$ for the norm in the spaces V_k^* , and simply $\|\cdot\|$ for the norm in the spaces W^k .

Definition 4.3. A Hilbert complex is said to have the compactness property if the inclusion $V^k \cap V_k^* \subset W^k$ is compact for each k.

Theorem 4.4. If a Hilbert complex has the compactness property, then it is closed and Fredholm.

Proof. We first check that d has closed range. Let T be the restriction of d to $D(T) := V \cap \mathfrak{Z}^{\perp}$, viewed as an unbounded operator $W \to W$. It is easy to check that the range of T coincides with the range of d and, since d is a closed operator, T is closed as well. Since $\mathfrak{Z}^{\perp} = \mathfrak{B}^* \subset \mathfrak{Z}^* \subset V^*$, $D(T) \subset V \cap V^*$ and the graph norm on D(T) is equivalent to the $V \cap V^*$ norm. By Theorem 3.9, T, and thus d, has closed range.

Next, we note that, on $\mathfrak{H} \subset V \cap V^*$, the $V \cap V^*$ norm is equivalent to the W norm. Thus the closed unit ball in \mathfrak{H} is compact, which implies that $\mathfrak{H} \cong \mathcal{H}$ is finite dimensional, i.e., the complex is Fredholm. \Box

4.2.2 • Hodge decomposition

The Hilbert complex structure leads directly to the Hodge decomposition, a generalization of the Helmholtz decomposition of vector calculus.

Since the space of harmonic k-forms is the orthogonal complement of \mathfrak{B}^k within \mathfrak{Z}^k , it follows immediately that

$$\mathfrak{Z}^k = \bar{\mathfrak{B}}^k \oplus \mathfrak{H}^k, \tag{4.4}$$

where the symbol \oplus indicates a sum which is orthogonal (note that in this context we need not distinguish between orthogonality in W and orthogonality in V). Since \mathfrak{Z}^k is a closed subspace of both W^k and V^k , we have

$$W^k = \mathfrak{Z}^k \oplus \mathfrak{Z}^{k\perp} = \mathfrak{Z}^k \oplus \bar{\mathfrak{B}}^*_k, \quad V^k = \mathfrak{Z}^k \oplus \mathfrak{Z}^{k\perp_V},$$

where $\mathfrak{Z}^{k\perp_V} = \mathfrak{Z}^{k\perp} \cap V = \overline{\mathfrak{B}}_k^* \cap V$. Combining with (4.4) we obtain the Hodge decomposition.

Theorem 4.5 (Hodge decomposition). Any Hilbert complex gives rise to the orthogonal decompositions

$$W^{k} = \underbrace{\widetilde{\mathfrak{B}}^{k} \bigoplus \mathfrak{H}^{k}}_{\mathfrak{Z}^{*\perp}_{k}} \underbrace{\mathfrak{S}^{k}_{k} \bigoplus \widetilde{\mathfrak{B}}^{k}_{k}}_{\mathfrak{Z}^{*\perp}_{k} = \mathfrak{B}^{k\perp}}$$
(4.5)

and

$$V^{k} = \bar{\mathfrak{B}}^{k} \oplus \mathfrak{H}^{k} \oplus \mathfrak{Z}^{k\perp_{V}}.$$

$$(4.6)$$

For a closed Hilbert complex we may, of course, drop the bars over the boundary spaces. In this case, the Hodge decomposition says that every element of Whas a unique decomposition into an element of the range of d, an element of the range of d^* , and a harmonic form, and this decomposition is orthogonal.

4.2.3 • Poincaré inequality

The final result we derive is the Poincaré inequality, for which it is necessary to assume that the Hilbert complex is closed.

Theorem 4.6 (Poincaré inequality). Given a closed Hilbert complex, for each k, there exists a constant c_P such that

$$||z||_V \le c_P ||dz||, \quad z \in \mathfrak{Z}^{k \perp V}.$$
 (4.7)

Proof. Since the Hilbert complex is closed, \mathfrak{B}^{k+1} is a Hilbert space. The operator d^k is thus a bounded linear isomorphism between the Hilbert spaces $\mathfrak{Z}^{k\perp_V}$ and \mathfrak{B}^{k+1} , and so it has a bounded inverse, which leads directly to (4.7). \Box

4.3 • Example: The L^2 de Rham complex

The prototypical example of a Hilbert complex is the de Rham complex associated to a Riemannian manifold. We continue to discuss the case of a domain Ω in \mathbb{R}^3 , which we shall assume to have a Lipschitz boundary. The base Hilbert spaces will be $W^0 = W^3 = L^2(\Omega)$, $W^1 = W^2 = L^2(\Omega; \mathbb{R}^3)$. The operators are the ones discussed in Section 3.4: $d^0 = \operatorname{grad}$, $d^1 = \operatorname{curl}$, $d^2 = \operatorname{div}$, with the domains taken to be H^1 , $H(\operatorname{curl})$, and $H(\operatorname{div})$, respectively. We saw there that these are closed operators, so we meet the criteria for a Hilbert complex. The domain complex, consisting of the spaces V^k , is

$$0 \to H^1 \xrightarrow{\text{grad}} H(\text{curl}) \xrightarrow{\text{curl}} H(\text{div}) \xrightarrow{\text{div}} L^2 \to 0.$$
(4.8)

In view of the adjoints computed in Section 3.4, the domain complex of the dual complex (consisting of the V_k^*) is

$$0 \leftarrow L^2 \xleftarrow{-\operatorname{div}} \mathring{H}(\operatorname{div}) \xleftarrow{\operatorname{curl}} \mathring{H}(\operatorname{curl}) \xleftarrow{-\operatorname{grad}} \mathring{H}^1 \leftarrow 0.$$
(4.9)

An important fact is that the L^2 de Rham complex satisfies the compactness property, i.e., that $V^k \cap V_k^*$ is compactly included in W^k . For k = 0, we have $V^0 \cap V_0^* = H^1 \cap L^2 = H^1$ which is indeed compact in $W^0 = L^2$ by Rellich's theorem. The result for k = 3 is similar. For k = 1 and 2 the desired results are that $H(\operatorname{curl}) \cap \mathring{H}(\operatorname{div})$ and $\mathring{H}(\operatorname{curl}) \cap H(\operatorname{div})$ are compact in $L^2(\Omega; \mathbb{R}^3)$. This was proved, in the greater generality of Riemannian manifolds with Lipschitz boundary, by Picard [62]. According to Theorem 4.4, the compactness property implies that the de Rham complex is Fredholm: the operators have closed range, and the harmonic form spaces are finite dimensional. In particular, we have established that $\operatorname{curl} H(\operatorname{curl})$ is closed, as claimed at the end of Chapter 3. We now calculate the spaces of harmonic forms and exhibit the Hodge decomposition for each degree.

The space of harmonic 0-forms coincides with \mathfrak{Z}^0 , the null space of the gradient. This is the space of constants, or, if we allow domains with more than the one connected component, the space of local constants. Its dimension is the number b_0 of connected components of Ω (the zeroth Betti number of the domain). Since the space $\mathfrak{B}^0 = 0$, the Hodge decomposition expresses L^2 as the sum div $\mathring{H}(\text{div})$ and the constant (or locally constant) functions.

The space of harmonic 1-forms is more complicated. It consists of the L^2 vector fields p for which $\operatorname{curl} p = 0$, $\operatorname{div} p = 0$, and the normal trace $\gamma_n p$ vanishes on $\partial\Omega$. An elementary example of such a function is pictured on the left of Figure 4.2, in which the domain is an annular cylinder. In cylindrical coordinates the domain is given by $a \leq r \leq b, c \leq z \leq d$, for some 0 < a < b, c < d, and the harmonic forms are simply the multiples of $p = \operatorname{grad} \theta$. On a general domain $\dim \mathfrak{H}^1 = b_1$, the first Betti number, which we defined in Section 2.3 and which counts the number of tunnels through the domain. For a simply connected domain, $b_1 = 0$, and the Hodge decomposition is the classical Helmholtz decomposition of a vector field into a gradient (which is irrotational, meaning curl-free) and a curl (solenoidal, meaning divergence-free). For a general domain, we must



Figure 4.2. Left: Harmonic 1-form on an annular cylinder in \mathbb{R}^3 . Right: Harmonic 2-form on a spherical shell (half the shell has been cut away for visibility).

also include the harmonic part:

$$L^{2}(\Omega; \mathbb{R}^{3}) = \operatorname{grad} H^{1}(\Omega) \oplus \operatorname{curl} \check{H}(\operatorname{curl}, \Omega) \oplus \mathfrak{H}^{1}.$$

The space of harmonic 2-forms consists of vector fields which are again curland divergence-free but, in this case, with vanishing tangential trace on the boundary. An elementary example is given in spherical coordinates on the closed spherical shell a < r < b by $p = \operatorname{grad}(1/r)$. See Figure 4.2. On a general domain, dim $\mathfrak{H}^2 = b_2$, the number of voids in the domain. The Hodge decomposition is again into gradients and curls, but the boundary conditions are different, as is the space of harmonic forms:

 $L^{2}(\Omega; \mathbb{R}^{3}) = \operatorname{grad} \mathring{H}^{1}(\Omega) \oplus \operatorname{curl} H(\operatorname{curl}, \Omega) \oplus \mathfrak{H}^{2}.$

Finally, the space of harmonic 3-forms is simple: $\mathfrak{H}^3 = 0$, since there are no nonzero functions in \mathring{H}^1 with vanishing gradient. Of course, \mathfrak{B}_3^* also vanishes, so the Hodge decomposition is trivial.

4.4 • The abstract Hodge Laplacian

Let (W, d) be a closed Hilbert complex. Associated to it is a graded operator from $W \to W$ of degree 0 given by $L = dd^* + d^*d$. In the case of the de Rham complex associated to a Riemannian manifold, this operator is called the Hodge Laplacian, and we refer to it generally as the (abstract) Hodge Laplacian associated to the complex. Specifically,

$$D(L^k) = \{ u \in V^k \cap V_k^* \, | \, du \in V_{k+1}^*, \, d^*u \in V^{k-1} \}, \tag{4.10}$$

and

$$L^{k}u = d_{k+1}^{*}d^{k}u + d^{k-1}d_{k}^{*}u = (d^{k})^{*}d^{k}u + d^{k-1}(d^{k-1})^{*}u, \quad u \in D(L^{k}).$$
(4.11)

A major subject of FEEC is the numerical solution of the Hodge Laplace equation $L^k u = f$ for given $f \in W^k$. Whether this problem has a unique solution depends on whether the space \mathfrak{H}^k of harmonic forms vanishes. Indeed, we now show that the null space of L^k is precisely \mathfrak{H}^k , explaining the terminology *harmonic*. From (4.2) we see that $\mathfrak{H}^k \subset D(L^k)$ and $L^k u = 0$ for all $u \in \mathfrak{H}^k$. Thus $\mathfrak{H}^k \subset \mathcal{N}(L^k)$. On the other hand, from the definition of L we have that

$$\langle L^k u, v \rangle = \langle du, dv \rangle + \langle d^* u, d^* v \rangle, \quad u \in D(L^k), \ v \in V^k \cap V_k^*.$$
(4.12)

Taking v = u, we get

$$\|du\|^2 + \|d^*u\|^2 = \langle L^k u, u \rangle, \quad u \in D(L^k),$$

and so,

$$||du||^2 + ||d^*u||^2 = 0,$$

if $u \in \mathcal{N}(L^k)$. Thus, if $u \in \mathcal{N}(L^k)$, then du and d^*u vanish, and so $u \in \mathfrak{H}^k$. This verifies that $\mathcal{N}(L^k) = \mathfrak{H}^k$, i.e., that solutions of the Hodge Laplacian are determined only up to addition of a harmonic form.

The harmonic forms also determine a necessary condition for the existence of solutions. Taking the inner product of the equation $L^k u = f$ with $q \in \mathfrak{H}^k$ and using (4.12), we see that $\langle f, q \rangle = 0$. Thus a necessary condition for the existence of a solution is the orthogonality condition $f \perp \mathfrak{H}^k$. In fact, in Theorem 4.8 below we shall show that this condition is not only necessary but also sufficient for existence of a solution.

4.4.1 • Three formulations of the Hodge Laplace problem

In order to formulate a problem which is solvable for any $f \in W^k$, we make use of the orthogonal projection $P_{\mathfrak{H}}$ of W^k onto \mathfrak{H}^k . Then we can formulate the following *strong form* of the Hodge Laplace equation, which we shall show is well-posed.

Strong formulation

Given $f \in W^k$, find $u \in D(L^k)$ such that

$$Lu = f - P_{\mathfrak{H}}f, \qquad u \perp \mathfrak{H}^k. \tag{4.13}$$

Note that if u is a solution to this problem, then

$$f = dd^*u + P_{5}f + d^*du. (4.14)$$

Clearly $dd^*u \in \mathfrak{B}$, $P_{\mathfrak{H}}f \in \mathfrak{H}$, and $d^*du \in \mathfrak{B}^*$, so (4.14) is the Hodge decomposition of f, i.e., $dd^*u = P_{\mathfrak{B}}f$ and $d^*du = P_{\mathfrak{B}^*}f$.

We will also consider two weak formulations of this problem. The primal formulation corresponds directly to the strong form.

Primal weak formulation

Given $f \in W^k$, find $u \in V^k \cap V_k^*$ such that $u \perp \mathfrak{H}^k$ and

$$\langle du, dv \rangle + \langle d^*u, d^*v \rangle = \langle f - P_{\mathfrak{H}}f, v \rangle, \quad v \in V^k \cap V_k^*.$$
 (4.15)

For the mixed weak formulation, we introduce two new variables, $\sigma = d^* u \in V^{k-1}$ and $p = P_{\mathfrak{H}} f \in \mathfrak{H}^k$.

Mixed weak formulation

Given $f \in W^k$, find $\sigma \in V^{k-1}$, $u \in V^k$, and $p \in \mathfrak{H}^k$ such that

$$\langle \sigma, \tau \rangle - \langle u, d\tau \rangle = 0, \qquad \tau \in V^{k-1},$$

$$\langle d\sigma, v \rangle + \langle du, dv \rangle + \langle p, v \rangle = \langle f, v \rangle, \quad v \in V^k,$$

$$\langle u, q \rangle = 0, \qquad q \in \mathfrak{H}^k.$$
 (4.16)

An important observation is that the mixed weak formulation, unlike the other two formulations, does not treat the operators d and d^* symmetrically. Only the d operator appears in this formulation, the terms involving d^* being treated weakly.

Despite their different appearance, the three formulations are completely equivalent.

Theorem 4.7 (equivalence of formulations). Let $f \in W^k$. An element $u \in W^k$ solves the strong formulation of the abstract Hodge Laplace problem if and only if it solves the primal weak formulation. Moreover, in this case, if we set $\sigma = d^*u$ and $p = P_{\mathfrak{H}}f$, then the triple (σ, u, p) solves the mixed weak formulation. Finally, if some (σ, u, p) solves the mixed weak formulation, then u solves the other two formulations of the problem, and $\sigma = d^*u$, $p = P_{\mathfrak{H}}f$.

Proof. First we show the equivalence of the strong formulation and the primal weak formulation. If u satisfies the strong formulation, then certainly $u \in V \cap V^*$ and satisfies the primal weak formulation. On the other hand, if u satisfies the primal weak formulation and *if* u *belongs to* D(L), then it satisfies the strong formulation. Therefore, we need to show that any u which satisfies the primal weak formulation, and so a priori only belongs to $V \cap V^*$, actually belongs to D(L) defined in (4.10), i.e., satisfies $du \in V^*$ and $d^*u \in V$. The statement that $du \in V^*$ means that there exists a constant c such that

$$\langle du, dv \rangle \le c \|v\|, \quad v \in V.$$
(4.17)

Using the Hodge decomposition, we split $v = v_1 + v_2$ with $v_1 \in \mathfrak{Z}, v_2 \in V \cap \mathfrak{B}^*$. Obviously,

$$\langle du, dv_1 \rangle = 0,$$

while

$$\langle du, dv_2 \rangle = \langle du, dv_2 \rangle + \langle d^*u, d^*v_2 \rangle = \langle f - P_{\mathfrak{H}}f, v_2 \rangle \le c \|v_2\| \le c \|v\|_{\mathcal{H}}$$

where $c = ||f - P_{55}f||$. Adding these results gives (4.17) and shows that $du \in V^*$. An analogous argument shows that $d^*u \in V$, and so we have proven that $u \in D(L)$ and so that a solution of the primal weak formulation is a strong solution.

Next, suppose u solves the strong formulation, and set $\sigma = d^*u$, $p = P_{55}f$. Since $u \in D(L)$, $\sigma \in V$. The first equation of the mixed formulation is a consequence of the definition of σ . We have, from the Hodge Laplace equation, that

$$d\sigma + d^*du + p = f,$$

and, taking the inner product with any $v \in V$, we obtain the second equation of the mixed method. The third equation follows from the assumption that $u \perp \mathfrak{H}^k$. This shows that (σ, u, p) solves the weak formulation.

Finally, if $(\sigma, u, p) \in V \times V \times \mathfrak{H}$ solves the mixed weak formulation, then the first equation implies $u \in V^*$ and $d^*u = \sigma$, and the last equation implies that $u \perp \mathfrak{H}$. Taking $v \in \mathfrak{H}$ in the second equation shows that $p = P_{\mathfrak{H}} f$. Finally, taking $v \in V \cap V^*$ in the second equation we conclude that u solves the primal weak formulation. \Box

4.4.2 • Well-posedness

Finally, we show that the abstract Hodge Laplace equation is well-posed once we account for any harmonic forms.

Theorem 4.8 (well-posedness). For each $f \in W^k$ there exists a unique solution $u \in D(L^k)$ such that $u \perp \mathfrak{H}^k$ and $Lu = f - P_{\mathfrak{H}}f$. Moreover

$$||u|| + ||du|| + ||d^*u|| + ||dd^*u|| + ||d^*du|| + ||p|| \le c||f||.$$

The constant c entering this estimate depends only on the constant c_P entering the Poincaré inequality (Theorem 4.6).

Proof. When a solution exists, the bounds on $||dd^*u||$, $||d^*du||$, and ||p|| are immediate from the Hodge decomposition $f = dd^*u + d^*du + p$, so it suffices to show that for any $f \in W$ a unique solution exists and satisfies

$$||u||_V + ||d^*u|| \le c||f||.$$

Our proof is based on the mixed weak formulation. Writing $X = V^{k-1} \times V^k \times \mathfrak{H}^k$ for brevity, define a bounded bilinear form $B : X \times X \to \mathbb{R}$ by combining the left-hand sides of the three equations defining the mixed formulation:

$$B(\sigma, u, p; \tau, v, q) = \langle \sigma, \tau \rangle - \langle u, d\tau \rangle - \langle d\sigma, v \rangle - \langle du, dv \rangle - \langle p, v \rangle - \langle u, q \rangle.$$
(4.18)

Note that B is symmetric. Also define a bounded linear functional $F: X \to \mathbb{R}$ by

$$F(\tau, v, q) = \langle f, v \rangle. \tag{4.19}$$

Then the mixed weak formulation has a standard form: find $x \in X$ such that

$$B(x,y) = F(y), \quad y \in X. \tag{4.20}$$

Nečas's generalization of the Lax-Milgram theorem [61, Theorem 3.1] tells us when such a problem is well-posed. Well-posedness (for arbitrary $F \in X'$) is equivalent to the *inf-sup condition*

$$\gamma := \inf_{0 \neq w \in X} \sup_{0 \neq y \in X} \frac{B(w, y)}{\|w\|_X \|y\|_X} > 0,$$
(4.21)

in which case the solution x satisfies $||x||_X \leq \gamma^{-1} ||F||_{X'}$. (A second inf-sup condition with the roles of w and y reversed is generally required but in our case is unnecessary, since B is symmetric.) The inf-sup condition is proven in the following theorem. \Box

Theorem 4.9. The bilinear form B defined by (4.18) satisfies the inf-sup condition (4.21) with γ bounded below by a constant depending only on the Poincaré constant c_P .

Proof. For any given $(\sigma, u, p) \in V^{k-1} \times V^k \times \mathfrak{H}$ we must exhibit $(\tau, v, q) \in V^{k-1} \times V^k \times \mathfrak{H}$, not identically zero, such that

$$B(\sigma, u, p; \tau, v, q) \ge \gamma(\|\sigma\|_V + \|u\|_V + \|p\|)(\|\tau\|_V + \|v\|_V + \|q\|),$$

where γ depends only on c_P . For this, we use the Hodge decomposition $u = u_{\mathfrak{B}} + u_{\mathfrak{H}} + u_{\mathfrak{B}^*}$, where $u_{\mathfrak{B}} = P_{\mathfrak{B}}u$, $u_{\mathfrak{H}} = P_{\mathfrak{H}}u$, and $u_{\mathfrak{B}^*} = P_{\mathfrak{B}^*}u$. Since $u_{\mathfrak{B}} \in \mathfrak{B}^k$, we may write $u_{\mathfrak{B}} = d\rho$ with $\rho \in \mathfrak{Z}^{k-1\perp}$. By the Poincaré inequality,

$$\|\rho\|_V \le c_P \|u_{\mathfrak{B}}\|. \tag{4.22}$$

Then we define

$$\tau = \sigma - \frac{1}{c_P^2} \rho \in V^{k-1}, \quad v = -u - d\sigma - p \in V^k, \quad q = -p + u_{\mathfrak{H}} \in \mathfrak{H}^k.$$
(4.23)

From (4.22) and the orthogonality of the Hodge decomposition, we have

$$\|\tau\|_{V} + \|v\|_{V} + \|q\| \le C(\|\sigma\|_{V} + \|u\|_{V} + \|p\|).$$
(4.24)

Plugging (4.23) into (4.18), we have

$$B(\sigma, u, p; \tau, v, q) = \|\sigma\|^2 + \|d\sigma\|^2 + \|du\|^2 + \|p\|^2 + \|u_{\mathfrak{H}}\|^2 + \frac{1}{c_P^2} \|u_{\mathfrak{H}}\|^2 - \frac{1}{c_P^2} \langle \sigma, \rho \rangle.$$

We use the Cauchy–Schwarz inequality, the algebraic-geometric mean inequality, and (4.22) to bound the final term:

$$\langle \sigma, \rho \rangle \le \|\sigma\| \|\rho\| \le \frac{c_P^2}{2} \|\sigma\|^2 + \frac{1}{2c_P^2} \|\rho\|^2 \le \frac{c_P^2}{2} \|\sigma\|^2 + \frac{1}{2} \|u_{\mathfrak{B}}\|^2.$$

Thus

$$B(\sigma, u, p; \tau, v, q) \ge \frac{1}{2} \|\sigma\|^2 + \|d\sigma\|^2 + \|du\|^2 + \|p\|^2 + \|u_{\mathfrak{H}}\|^2 + \frac{1}{2c_P^2} \|u_{\mathfrak{H}}\|^2.$$

Now we apply the Poincaré inequality to $u_{\mathfrak{B}^*}$, for which we have $du_{\mathfrak{B}^*} = du$:

$$\|u_{\mathfrak{B}^*}\|_V \le c_P \|du\|.$$

We obtain

$$B(\sigma, u, p; \tau, v, q) \\ \geq \frac{1}{2} \|\sigma\|^2 + \|d\sigma\|^2 + \frac{1}{2} \|du\|^2 + \|p\|^2 + \|u_{\mathfrak{H}}\|^2 + \frac{1}{2c_P^2} \|u_{\mathfrak{B}}\|^2 + \frac{1}{2c_P^2} \|u_{\mathfrak{B}^*}\|^2 \\ \geq \eta(\|\sigma\|^2 + \|d\sigma\|^2 + \|u_{\mathfrak{B}}\|^2 + \|u_{\mathfrak{H}}\|^2 + \|u_{\mathfrak{B}^*}\|^2 + \|du\|^2 + \|p\|^2) \\ = \eta(\|\sigma\|_V^2 + \|u\|_V^2 + \|p\|^2),$$

where $\eta = \min(1/2, 1/2c_P^2)$ depends only on c_P . The inf-sup condition easily follows from this bound and (4.24).

Remark 4.10. The inf-sup condition in Theorem 4.9 implies the well-posedness of the mixed problem (4.20) for general $F \in X'$, not only for F of the form (4.19).

Having established well-posedness of the Hodge Laplacian, we denote by K the solution operator. That is, for $f \in W$, we define $Kf \in D(L)$ by the equations $LKf = f - P_{\mathfrak{H}}f$, $Kf \perp \mathfrak{H}$. Written in terms of K, the solution to the mixed formulation is given by $\sigma = d^*Kf$, u = Kf, $p = P_{\mathfrak{H}}f$. Thus the Hodge projections are $P_{\mathfrak{B}} = dd^*K$, $P_{\mathfrak{B}^*} = d^*dK$, and

$$u = dd^*Ku + P_{\mathfrak{H}}u + d^*dKu, \quad u \in W.$$

Remark 4.11. The reader may wonder why we chose to establish well-posedness of the abstract Hodge Laplacian using the mixed formulation, rather than the primal weak formulation. After all, the bilinear form for the primal formulation is coercive (by the Poincaré inequality) and hence we can apply the Riesz representation theorem and dispense with the verification of the inf-sup condition. The reason is that we are anticipating the discretization of the Hodge Laplacian, which will be the subject of the next chapter. The discretization will be based on the mixed formulation, which, as we shall see, turns out to be far more appropriate for discretization than the primal formulation. The inf-sup condition for the mixed bilinear form (and not only well-posedness for F of the form (4.19)) will then be the key to the analysis of the discretization method.

4.4.3 \bullet The \mathfrak{B} and \mathfrak{B}^* problems

In this section, we examine the Hodge Laplace equation Lv = g when the data g belongs to \mathfrak{B} or \mathfrak{B}^* . In the former case, g = df for some $f \in V$ (if $g \in V^k$, then $f \in V^{k-1}$). From the solution u of the Hodge Laplace equation

$$Lu = f - P_{\mathfrak{H}}f \tag{4.25}$$

with data f we obtain the solution to Lv = g by setting v = du. This is reasonable since, formally,

$$Lv = (dd^* + d^*d)du = dd^*du = d(dd^* + d^*d)u = d(f - P_{\mathfrak{H}}f) = g.$$

This formal calculation is justified as long as $v = du \in D(L)$. To see this, note that, since $u \in D(L)$, $v = du \in V^*$ and $v \in V$. Moreover dv = 0 and $d^*v = d^*du = P_{\mathfrak{B}^*}f = f - P_{\mathfrak{Z}}f \in V$, since $f \in V$ and $\mathfrak{Z} \subset V$. This shows that $v \in D(L)$, as claimed.

Thus, when the data g for the Hodge Laplacian belongs to \mathfrak{B} , the corresponding solution v also belongs to \mathfrak{B} . But then the term d^*dv in the Hodge Laplacian vanishes, so v satisfies another problem.

The \mathfrak{B} problem

Given $g \in \mathfrak{B}^k$, find $v \in \mathfrak{B}^k$ such that $dd^*v = g$.

By definition, a solution to this problem must belong as well to $D(dd^*) = \{v \in V_k^* \mid d^*v \in V^{k-1}\}$. We have just seen that, for $g \in \mathfrak{B}^k$, if v solves the Hodge Laplace problem with data g, then v solves the \mathfrak{B} problem. Conversely, if v solves the \mathfrak{B} problem, then we have dv = 0, and so $Lv = dd^*v = g$, i.e., v solves the Hodge Laplace problem. From the well-posedness of the Hodge Laplace problem, we obtain well-posedness of the \mathfrak{B} problem (but only for $g \in \mathfrak{B}^k$; otherwise there is no solution).

Theorem 4.12. If $g \in \mathfrak{B}^k$, then there is a unique solution v to the \mathfrak{B} problem with data g and this solution coincides with the unique solution to the Hodge Laplace problem with the same data.

Of course, we have the analogous result for d^* as well: if $f \in V^*$ and u the corresponding solution to the Hodge Laplace problem, then $v = d^*u \in D(L)$ and is the solution to the Hodge Laplace problem with data $g = d^*f$. Moreover, if $g \in \mathfrak{B}^*$, then it is also the unique solution to the following problem.

The \mathfrak{B}^* problem

Given $g \in \mathfrak{B}_k^*$, find $v \in \mathfrak{B}_k^*$ such that $d^*dv = g$.

4.5 • Example: The Hodge Laplacian in \mathbb{R}^3

We now return to our running example, the de Rham complex on a domain with Lipschitz boundary in \mathbb{R}^3 . We saw in Section 4.3 that this forms a closed Hilbert complex. Therefore, for each k = 0, 1, 2, 3, we obtain a Hodge Laplace problem

associated to the complex, which is a well-posed boundary value problem in partial differential equations. To interpret these problems we just have to interpret the spaces V^k and V_k^* and the operators d^k and d_k^* , for which we will have frequent recourse to the domain complex (4.8) and the dual domain complex (4.9).

4.5.1 • k = 0

Assuming that the domain is connected, the space \mathfrak{H}^0 of harmonic forms consists only of the constants and $P_{\mathfrak{H}}f = \overline{f}$, the mean value of f on Ω . Since $d_0^* = 0$, $L^0 u = (dd^* + d^*d)u = d^*du = -\operatorname{div}\operatorname{grad} u = -\Delta u$, with domain

$$D(L) = \{ u \in H^1(\Omega) \mid \operatorname{grad} u \in \mathring{H}(\operatorname{div}, \Omega) \}$$

= $\{ u \in H^1(\Omega) \mid \Delta u \in L^2(\Omega), \ \partial u / \partial n = 0 \text{ on } \partial \Omega \},$

where Δ denotes the ordinary Laplacian $\Delta = \sum_{i=1}^{3} \partial^2 / \partial x_i^2$. Given $f \in L^2(\Omega)$, the strong formulation (4.13) then seeks $u \in H^1$ such that

$$-\Delta u = f - \bar{f} \text{ in } \Omega, \quad \frac{\partial u}{\partial n} = 0 \text{ on } \partial \Omega, \quad \int u = 0.$$
 (4.26)

Thus the k = 0 Hodge Laplace problem associated to the de Rham complex is just the usual Neumann problem for Poisson's equation, the adjustment by \bar{f} being required to make the data f compatible with the homogeneous Neumann boundary condition.

Next we turn to the primal and mixed weak formulations. Since $V^{-1} = 0$, for k = 0, σ and τ vanish in the mixed weak formulation (4.16). Thus there is no distinction between the two weak formulations for k = 0. They both seek $u \in H^1$ such that

$$\int \operatorname{grad} u \cdot \operatorname{grad} v = \int (f - \bar{f})v, \quad v \in H^1, \quad \int_{\Omega} u = 0.$$
(4.27)

We have seen in Theorem 4.7 that this problem is equivalent to the strongly formulated problem (4.26). For concreteness, we now check directly that (4.27) implies (4.26), the other direction being easier. Taking $v \in C_0^{\infty}(\Omega)$ in (4.27) we see that

$$-\operatorname{div}\operatorname{grad} u = f - \bar{f} \tag{4.28}$$

in the sense of distributions and so, since $f \in L^2$, grad u belongs to H(div). We may therefore apply the integration-by-parts formula (3.8) with w = grad u and then apply (4.27) and (4.28) to get

$$\langle \gamma_n(\operatorname{grad} u), \operatorname{tr} v \rangle = \int (\operatorname{div} \operatorname{grad} u)v + \int \operatorname{grad} u \cdot \operatorname{grad} v = 0, \quad v \in H^1.$$

Since tr v can be an arbitrary element of $H^{1/2}(\partial\Omega)$ we conclude that $\gamma_n(\operatorname{grad} u) = 0$, i.e., $\partial u/\partial n = 0$ on the boundary. Notice that, in the weak formulation, the boundary condition is imposed *naturally*, meaning that it is implied by the equations via integration by parts rather than being imposed in the space in which u is sought.

Turning to the \mathfrak{B} and \mathfrak{B}^* problems, we do not find anything new, since $\mathfrak{B}^0 = 0$. Therefore, the \mathfrak{B} problem does not arise for k = 0, while the \mathfrak{B}^* problem is essentially the full Hodge Laplacian.

Summarizing the situation for k = 0, the Hodge Laplacian for 0-forms is just the Neumann problem for the Poisson equation. The primal and mixed weak formulations coincide in this case, being just the standard weak formulation for the Neumann problem, in which the boundary condition is imposed *naturally*.

4.5.2 ■ *k* = 1

The case k = 1 is more interesting. The space of harmonic forms \mathfrak{H}^1 is a space of vector fields whose dimension is b_1 , the first Betti number. The strong formulation seeks $u \in H(\operatorname{curl}) \cap \mathring{H}(\operatorname{div})$ with $\operatorname{curl} u \in \mathring{H}(\operatorname{curl})$ and $\operatorname{div} u \in H^1$ satisfying the partial differential equation

$$\operatorname{curl}\operatorname{curl} u - \operatorname{grad}\operatorname{div} u = f - P_{\mathfrak{H}}f \text{ in }\Omega.$$
(4.29)

Note that the differential operator curl curl – grad div is the negative vector Laplacian $-\vec{\Delta}$ in that curl curl u – grad div $u = (-\Delta u_1, -\Delta u_2, -\Delta u_3)$ for a smooth vector field $u = (u_1, u_2, u_3)$.

The definition of the domain $D(L^1)$ implies the boundary conditions

$$u \cdot n = 0, \quad (\operatorname{curl} u) \times n = 0 \text{ on } \partial\Omega,$$
(4.30)

and, of course, we have the side condition $u \perp \mathfrak{H}^1$ to impose uniqueness. The conditions (4.30) are sometimes referred to as *magnetic boundary conditions*, for reasons that will become clearer when we discuss Maxwell's equations in Section 8.6.

The primal weak formulation seeks $u \in H(\text{curl}) \cap \mathring{H}(\text{div})$ with

$$\langle \operatorname{curl} u, \operatorname{curl} v \rangle + \langle \operatorname{div} u, \operatorname{div} v \rangle = \langle f - P_{\mathfrak{H}} f, v \rangle, v \in H(\operatorname{curl}) \cap \mathring{H}(\operatorname{div}).$$
 (4.31)

Notice that, in this formulation, the boundary condition $u \cdot n = 0$ is essential (imposed in the space where u is sought), while the condition $\operatorname{curl} u \times n = 0$ is natural.

By contrast, the mixed weak formulation seeks $\sigma \in H^1$, $u \in H(\text{curl})$, $p \in \mathfrak{H}^1$, without any imposed boundary conditions, satisfying

$$\int \sigma \, \tau - \int u \cdot \operatorname{grad} \tau = 0, \qquad \tau \in H^1,$$

$$\int \operatorname{grad} \sigma \cdot v + \int \operatorname{curl} u \cdot \operatorname{curl} v + \int p \cdot v = \int f \cdot v, \quad v \in H(\operatorname{curl}), \quad (4.32)$$

$$\int u \cdot q = 0, \qquad q \in \mathfrak{H}^1.$$

The boundary conditions (4.30) are both natural in this formulation, arising from the integration by parts in the first and second equations, respectively. Notice that, in addition to the unknown vector field u, this formulation takes as a second unknown the scalar function $\sigma = -\operatorname{div} u$ (and of course $p = P_{\mathfrak{H}} f$).

In the case $f \in \mathfrak{B}^1$, i.e., f is a gradient, then the solution u is a gradient as well, which we can express by saying that $\operatorname{curl} u = 0$ and $u \perp \mathfrak{H}^1$. Thus the \mathfrak{B}^1 problem is to find u such that

$$-\operatorname{grad}\operatorname{div} u = f$$
, $\operatorname{curl} u = 0$ in Ω , $u \cdot n = 0$ on $\partial\Omega$, $u \perp \mathfrak{H}^1$.

From the abstract theory, we know that this problem has a unique solution for $f \in \mathfrak{B}^1$ which, moreover, can be computed by solving the full Hodge Laplacian problem for that same f (e.g., via the mixed weak formulation). One application of such a problem is toward solving the div-curl problem

$$\operatorname{div} u = F, \quad \operatorname{curl} u = 0,$$

which can be done by taking $f = \operatorname{grad} F$ (with $\int F = 0$ required for compatibility with the boundary condition on u).

On the other hand, the \mathfrak{B}^* problem arises when $f \in \mathfrak{B}_1^*$, meaning $f \in \operatorname{curl} \mathring{H}(\operatorname{curl})$. Then u is determined by

$$\operatorname{curl}\operatorname{curl} u = f, \operatorname{div} u = 0 \operatorname{in} \Omega, \quad u \cdot n = 0, \operatorname{curl} u \times n = 0 \operatorname{on} \partial \Omega, \quad u \perp \mathfrak{H}^{1}.$$

Such curl-curl problems arise in electromagnetics. Again, the solution may be determined by solving the full Hodge Laplace problem.

4.5.3 • *k* = 2

This case again leads to the vector Laplace equation (4.29), but with a different set of boundary conditions (and different harmonic forms with dim $\mathfrak{H}^2 = b_2$). The domain of L^2 is the set of $u \in H(\text{div}) \cap \mathring{H}(\text{curl})$ with div $u \in \mathring{H}^1$ and curl $u \in H(\text{curl})$. Thus we obtain *electric boundary conditions*:

$$u \times n = 0$$
, div $u = 0$ on $\partial \Omega$.

The first of these conditions is essential in the primal variational formulation, while the second is natural. Both are, of course, natural in the mixed weak formulation, which seeks $\sigma \in H(\text{curl}), u \in H(\text{div}), p \in \mathfrak{H}^2$ such that

$$\int \sigma \cdot \tau - \int u \cdot \operatorname{curl} \tau = 0, \qquad \tau \in H(\operatorname{curl}),$$
$$\int \operatorname{curl} \sigma \cdot v + \int \operatorname{div} u \operatorname{div} v + \int p \cdot v = \int f \cdot v, \quad v \in H(\operatorname{div}), \qquad (4.33)$$
$$\int u \cdot q = 0, \qquad q \in \mathfrak{H}^2.$$

Notice that this formulation is very different from the mixed weak formulation (4.32) for k = 1. Here the additional field σ is a vector field, while it was a scalar field for k = 1.

The \mathfrak{B} problem for k = 2 is

curl curl
$$u = f$$
, div $u = 0$ in Ω , $u \times n = 0$, on $\partial \Omega$, $u \perp \mathfrak{H}^2$,

which has a unique solution if $f \in \operatorname{curl} H(\operatorname{curl})$. The \mathfrak{B}_2^* problem is

$$-\operatorname{grad}\operatorname{div} u = f, \operatorname{curl} u = 0 \text{ in } \Omega, \quad u \times n = 0, \operatorname{div} u = 0 \text{ on } \partial\Omega, \quad u \perp \mathfrak{H}^2,$$

solvable for $f \in \operatorname{grad} \mathring{H}^1$. Notice that the differential equations, but not the boundary conditions, for the \mathfrak{B}^2 problem coincide with those of the \mathfrak{B}^*_1 problem and those for the \mathfrak{B}^*_2 problem coincide with those of the \mathfrak{B}^1 problem.

Table 4.1. The Hodge Laplace problem for the de Rham complex. σ is the extra variable in mixed weak formulation. Underlined boundary conditions (the ones not involving derivatives) are essential in the primal weak formulation, while the others are natural. All the boundary conditions are natural in the mixed weak formulation.

k	L^k	σ	Boundary conditions		Harmonic forms
0	$-\Delta$	0	Neumann:	$\partial u/\partial n=0$	\mathfrak{H}^0 (dim b_0)
1	$-ec\Delta$	$-\operatorname{div} u$	magnetic:	$\underline{u \cdot n = 0}$, $\operatorname{curl} u \times n = 0$	\mathfrak{H}^1 (dim b_1)
2	$-ec\Delta$	$\operatorname{curl} u$	electric:	$\underline{u \times n = 0}$, div $u = 0$	\mathfrak{H}^2 (dim b_2)
3	$-\Delta$	$-\operatorname{grad} u$	Dirichlet:	$\underline{u=0}$	$\mathfrak{H}^3=0$

4.5.4 • k = 3

Both \mathfrak{B}_3^* and \mathfrak{H}^3 vanish. The domain of the Laplacian $D(L^3) = \{ u \in \mathring{H}^1 | \Delta u \in L^2 \}$, so the boundary value problem is just the Dirichlet problem for Poisson's equation:

$$-\Delta u = f \text{ in } \Omega, \quad u = 0 \text{ on } \partial \Omega$$

The primal weak formulation is the obvious one: find $u \in \mathring{H}^1$ such that

$$\int \operatorname{grad} u \cdot \operatorname{grad} v = \int f v, \quad v \in \mathring{H}^1$$

The mixed weak formulation seeks $\sigma \in H(\text{div}), u \in L^2$ such that

$$\int \sigma \cdot \tau - \int u \operatorname{div} \tau = 0, \quad \tau \in H(\operatorname{div}), \qquad \int \operatorname{div} \sigma v = \int fv, \quad v \in L^2.$$

The Dirichlet condition is natural in the mixed formulation.

Table 4.1 collects key features of the Hodge Laplacian problems for the de Rham complex.

4.5.5 • The de Rham complex with boundary conditions and duality

We have just considered the application of the abstract theory to the L^2 de Rham complex, whose domain complex is given by (4.8). Instead, we may apply the theory to the L^2 de Rham complex with boundary conditions, by which we mean the Hilbert complex whose base spaces remain L^2 but whose domain complex is

$$0 \to \mathring{H}^1 \xrightarrow{\text{grad}} \mathring{H}(\text{curl}) \xrightarrow{\text{curl}} \mathring{H}(\text{div}) \xrightarrow{\text{div}} L^2 \to 0.$$

This Hilbert complex is essentially the dual complex to (4.8), i.e., (4.9), except for some signs and for a change of indexing. Specifically, the index k space of the de Rham complex with boundary conditions is the index 3 - k space of the dual complex to the original de Rham complex. The dual complex to the de Rham complex with boundary conditions is then

$$0 \leftarrow L^2 \xleftarrow{-\operatorname{div}} H(\operatorname{div}) \xleftarrow{\operatorname{curl}} H(\operatorname{curl}) \xleftarrow{-\operatorname{grad}} H^1 \leftarrow 0$$

i.e., the original de Rham complex, except for signs and indexing.

Table 4.2. The Hodge Laplace problem for the de Rham complex with boundary conditions. All boundary conditions are essential for the mixed formulation and underlined boundary conditions are essential for the primal formulation. See Table 4.1.

k	L^k	σ	Boundary conditions		Harmonic forms
0	$-\Delta$	0	Dirichlet:	$\underline{u=0}$	$\mathring{\mathfrak{H}}^0 = 0$
1	$-\vec{\Delta}$	$-\operatorname{div} u$	electric:	$\underline{u \times n = 0}$, div $u = 0$	$\mathring{\mathfrak{H}}^1 = \mathfrak{H}^2 \ (\dim b_2)$
2	$-\vec{\Delta}$	$\operatorname{curl} u$	magnetic:	$\underline{u \cdot n = 0}$, $\operatorname{curl} u \times n = 0$	$\mathring{\mathfrak{H}}^2 = \mathfrak{H}^1$ (dim b_1)
3	$-\Delta$	$-\operatorname{grad} u$	Neumann:	$\partial u/\partial n=0$	$\mathring{\mathfrak{H}}^3 = \mathfrak{H}^0$ (dim b_0)

If we write $\mathring{\mathfrak{B}}^k$, $\mathring{\mathfrak{Z}}^k$, and $\mathring{\mathfrak{H}}^k$ for the spaces of boundaries, cycles, and harmonic forms of the complex with boundary conditions, then $\mathring{\mathfrak{B}}^k = \mathfrak{B}^{3-k}$ and similarly for \mathfrak{Z} and \mathfrak{H} . (The relation $\mathfrak{H}^{n-k} = \mathring{\mathfrak{H}}^k$ on an *n*-dimensional manifold is known as *Poincaré–Lefschetz duality*. For manifolds without boundary it becomes $\mathfrak{H}^{n-k} = \mathfrak{H}^k$, which is Poincaré duality.) It is easy to check that the strong formulation and the primal weak formulation for the complex with boundary conditions coincide with those for the ordinary de Rham complex, except for an index change, so we obtain the same boundary value problems as above. The k = 0 problem, for instance, is the Poisson equation with Dirichlet boundary conditions. The mixed weak formulation for the problems with k > 0, however, brings in something new. For k = 1, for example, we seek $\sigma \in \mathring{H}^1$, $u \in \mathring{H}(\operatorname{curl})$, $p \in \mathring{\mathfrak{H}}^1 = \mathfrak{H}^2$ such that

$$\int \sigma \, \tau - \int u \cdot \operatorname{grad} \tau = 0, \qquad \tau \in \mathring{H}^1,$$
$$\int \operatorname{grad} \sigma \cdot v + \int \operatorname{curl} u \cdot \operatorname{curl} v + \int p \cdot v = \int f \cdot v, \quad v \in \mathring{H}(\operatorname{curl}).$$
$$\int u \cdot q = 0, \qquad q \in \mathfrak{H}^2.$$

This is a mixed formulation for the vector Laplacian with boundary conditions div u = 0, $u \times n = 0$. These are the same boundary conditions as for the k = 2 problem for the ordinary de Rham complex. However, the mixed method (4.33) for that problem was altogether different, using a vector field for the auxiliary variable σ rather than a scalar field, as here, and imposing the boundary conditions naturally, rather than essentially as here.

A similar situation holds for k = 2 and k = 3, which we leave to the reader to work out in detail. A summary is presented in Table 4.2.

Box 4.1. Summary of Chapter 4 on Hilbert complexes.

A *Hilbert complex* is a sequence of Hilbert spaces W^k and closed densely defined operators d^k from one to the next such that d^k maps its domain into the null space of d^{k+1} . The Hilbert complex is called *closed* if the d^k all have closed range.

Setting $V^k = D(d^k)$, we get the domain complex

$$0 \to V^0 \xrightarrow{d} V^1 \xrightarrow{d} \cdots \xrightarrow{d} V^n \to 0,$$

a cochain complex of Hilbert spaces and bounded linear operators.

We denote the cocycles and coboundaries as \mathfrak{Z}^k and \mathfrak{B}^k and, for the dual complex, by \mathfrak{Z}^*_k and \mathfrak{B}^*_k .

The Hodge Laplacian is the unbounded operator L^k from $W^k \to W^k$ given by $d^*d + dd^*$ with the natural domain. Its kernel is the space $\mathfrak{H}^k = \mathfrak{Z}^k \cap \mathfrak{Z}^*_k$ of harmonic forms. For a closed Hilbert complex, \mathfrak{H}^k is isomorphic to the cohomology space \mathcal{H}^k .

Any closed Hilbert complex gives rise to the orthogonal Hodge decomposition

$$W^k = \mathfrak{B}^k \oplus \mathfrak{H}^k \oplus \mathfrak{Z}^k$$

and satisfies the Poincaré inequality

$$||z||_V \le c_P ||dz||, \quad z \in \mathfrak{Z}^{k \perp V}.$$

The Hodge Laplace problem

$$L^k u = f - P_{\mathfrak{H}} f, \quad u \perp \mathfrak{H}^k,$$

has a unique solution in $D(L^k)$ for any $f \in W^k$. The solution may be characterized as well by the primal weak formulation (4.15) or the mixed weak formulation (4.16).

Chapter 5 Approximation of Hilbert complexes

Our goal now is to design and analyze numerical methods for the approximate solution of a Hodge Laplacian problem associated to some closed Hilbert complex. We shall use Galerkin methods, i.e., we shall replace the trial and test spaces occuring in the primal or mixed weak formulation by finite dimensional subspaces. In practice, when the Hilbert complex is the de Rham complex or something similar, the subspaces we consider will be finite element spaces.

To define the Hodge Laplace problem at some level k of the Hilbert complex, we require only the operators d^{k-1} from W^{k-1} to W^k with domain V^{k-1} and the d^k from W^k to W^{k+1} with domain V^k ; cf. (4.11). Therefore, in the interest of carrying no more baggage than necessary, in this chapter we do not require the whole Hilbert complex (4.1) but only a short segment of it:

$$V^{k-1} \xrightarrow{d} V^k \xrightarrow{d} V^{k+1}.$$
(5.1)

We shall see that Galerkin discretization of the primal formulation is problematic and shall focus on Galerkin discretization of the mixed formulation. This requires two compatible Galerkin subspaces, one for V^{k-1} and one for V^k . The main result of this chapter is to establish assumptions on the choice of subspaces, which ensures consistency and stability of the Galerkin method, and to deduce consequences from it. Box 5.1 summarizes the main results of the chapter.

5.1 - Galerkin discretization of the primal formulation

A simple approach to discretization of the Hodge Laplacian would be the application of Galerkin's method to the primal weak formulation (4.15), i.e., restricting the trial and test functions to some finite dimensional subspace $V_h \subset V^k \cap V_k^* \cap \mathfrak{H}^{k\perp}$. However, this approach suffers from two serious deficiencies, as we shall now show. First, we can anticipate some difficulty working with the condition $V_h \perp \mathfrak{H}^k$, since in many problems, we do not know the space \mathfrak{H}^k explicitly and may not even know its dimension. We will return to this issue in a moment. Second, *the primal approach may have difficulties even when there are no non-vanishing harmonic forms*. The difficulty here is that the space $V^k \cap V_k^*$ is not a

standard Sobolev space, and it may be difficult to design a finite element subspace of it with good approximation properties. More precisely, the space may contain vector fields which are too singular to belong to H^1 or even to be approximated arbitrarily closely by H^1 vector fields. Such singular fields arise as solutions of the Hodge Laplacian on domains with nonconvex corners and in other situations, so this is a serious defect.

As a simple example, suppose that Ω is a nonconvex simply connected polygon in the plane and consider the two-dimensional L^2 de Rham complex, for which the domain complex is

$$0 \to H^1(\Omega) \xrightarrow{\text{grad}} H(\text{curl}, \Omega) \xrightarrow{\text{curl}} L^2(\Omega) \to 0$$

(here curl $u = \partial u_2 / \partial x_1 - \partial u_1 / \partial x_2$). Since the domain is simply connected there are no harmonic 1-forms, and so the primal weak formulation (4.31) of the k = 1 Hodge Laplacian for this complex seeks $u \in Y := H(\text{curl}) \cap \mathring{H}(\text{div})$ such that

$$\int_{\Omega} \left[(\operatorname{curl} u)(\operatorname{curl} v) + (\operatorname{div} u)(\operatorname{div} v) \right] dx = \int_{\Omega} f \cdot v \, dx, \quad v \in Y.$$
(5.2)

Let $H_n^1 = H^1(\Omega; \mathbb{R}^2) \cap \mathring{H}(\operatorname{div})$ denote the space of H^1 vector fields with vanishing normal component on the boundary, which is a closed subspace of $H^1(\Omega; \mathbb{R}^2)$. Clearly $H_n^1 \subset Y$. Moreover,³

$$\|\operatorname{curl} u\|^2 + \|\operatorname{div} u\|^2 = \|\operatorname{grad} u\|^2, \quad u \in H_n^1.$$

It follows that the restriction of the Y norm (the norm of $H(\text{div}) \cap H(\text{curl})$) to H_n^1 is equivalent to the full H^1 norm there. Since H_n^1 is closed in H^1 , it is a *closed* subspace of Y. Moreover, $H_n^1 \neq Y$. This follows from the fact that a generic solution of the Neumann problem for Poisson's equation on a nonconvex polygon exhibits singular behavior near the reentrant corners, which precludes it belonging to H^2 . Thus there exists $\phi \in H^1(\Omega)$ such that $\Delta \phi \in L^2(\Omega)$ and $\partial \phi / \partial n = 0$ on $\partial\Omega$, but $\phi \notin H^2(\Omega)$. Setting $u = \operatorname{grad} \phi$ we see that $u \in Y$ but $u \notin H^1$. In fact, this lack of H^1 regularity is typical of solutions of (5.2). For any such solution, we have $\inf_{v \in H_n^1} \|u - v\|_Y = \delta_u > 0$, where δ_u is the distance of u from the proper closed subspace H_n^1 . This tells us that if the subspaces V_h are contained in H_n^1 , then the numerical solution $u_h \in V_h$ cannot converge to u. This is the case when V_h is any subspace of Y which consists of piecewise polynomials with respect to some triangulation. Indeed, the condition $V_h \subset H(\text{curl})$ implies that the tangential components of any $v \in V_h$ must be continuous across triangle boundaries, and the fact that $V_h \subset H(div)$ implies that the normal components must be continuous. Thus V_h consists of continuous piecewise polynomial functions,

³This identity follows from a more general one: $\int \operatorname{curl} u \operatorname{curl} v \, dx + \int \operatorname{div} u \operatorname{div} v \, dx = \int \operatorname{grad} u : \operatorname{grad} v \, dx + \sum_e \int_e \left(\frac{\partial u \cdot s}{\partial s} v \cdot n - \frac{\partial u \cdot n}{\partial s} v \cdot s \right) \, ds, \quad u, v \in H^1(\Omega; \mathbb{R}^2)$, where the sum is over all the edges *e* of the polygon Ω and *n* and *s* are the (constant) normal and tangent vectors on *e*. This is elementary for $u, v \in C^1(\overline{\Omega}; \mathbb{R}^2)$ (integration by parts). Each of the integrals appearing is a bilinear form of *u* and *v* bounded with respect to the H^1 norm, so we can extend by density. For example, $|\int_e \frac{\partial u \cdot s}{\partial s} v \cdot n| \leq ||\frac{\partial u \cdot s}{\partial s}||_{H^{-1/2}(e)} ||v \cdot n||_{H^{1/2}(e)} \leq ||u \cdot s||_{H^{1/2}(e)} ||v \cdot n||_{H^{1/2}(e)} \leq ||u||_{H^1(\Omega)} ||v||_{H^1(\Omega)}.$



Figure 5.1. Finite element solution to the Hodge Laplacian problem on an Lshaped domain (with f = (1, 0)). The left figure is calculated with a mixed method which is known to converge to the solution in L^2 . The right figure is based on the primal formulation using 24,576 piecewise linear elements. The primal-based numerical solution entirely misses the dominant behavior at the reentrant corner and produces a wholly inaccurate solution.

and so $V_h \subset H_n^1$. In short, a standard finite element method based on the primal formulation does not generally converge to the solution u. This behavior is illustrated in Figure 5.1, which depicts a numerical analyst's nightmare. The figure on the right shows the numerical solution computed with what would seem to be a natural discretization: the finite element method based on the primal weak formulation, using continuous piecewise linear vector fields as trial functions. The method converges, as the mesh is refined, to what might be taken to be the solution but is, in fact, an altogether different vector field. In three dimensions, the situation is similar but worse, because then H_n^1 is a closed subspace of infinite codimension. See [34] for an illuminating discussion.

We encountered this problem already in Chapter 1 when we computed the eigenvalues of the vector Laplacian on a polygonal domain with a hole. Because of the reentrant corners, the true eigenfunctions are singular, and the standard finite element method cannot compute them accurately. This results in nonconvergence of the eigenvalues as well, as we saw in the example.

An independent difficulty with the primal formulation arises when there are harmonic forms. The difficulty is that we usually cannot arrange that the harmonic forms belong to the subspace V_h . The harmonic forms will generally be smooth functions, but not polynomials. Therefore, if the space V_h is a finite element space, consisting of piecewise polynomials, it will not contain the any nonzero harmonic forms. The discrete problem, to find $u_h \in V_h$ such that

$$\langle du_h, dv \rangle + \langle d^*u_h, d^*v \rangle = \langle f, v \rangle, \quad v \in V_h,$$

therefore has a unique solution (the bilinear form on the left-hand side defines an operator $V_h \rightarrow V'_h$ with zero kernel). Stated in other terms, the space of discrete harmonic forms vanishes. Therefore, the side condition $u \perp \mathfrak{H}$ is lost on passage to the discrete level. The result, illustrated in Figure 5.2, is that the discrete solution does not resemble the exact solution, even in order of magnitude. Another nightmare!



Figure 5.2. Approximation of the Hodge Laplacian problem on an annulus (f = (0, x)). The true solution shown here on the left is an (accurate) approximation by a mixed method. The standard Galerkin solution using continuous piecewise linear vector fields, shown on the right, is totally different.

5.2 • Mixed methods of discretization

5.2.1 • The mixed Galerkin method

In view of the deficiencies of Galerkin methods based on the primal weak formulation, we now turn to mixed methods, i.e., Galerkin methods based on the *mixed* weak formulation. Here the situation is much better. Indeed, if the Galerkin subspaces are carefully chosen, then a mixed method returns an accurate, stable, convergent approximation, avoiding both of the problems just discussed for the primal formulation. To achieve this, we will impose three conditions on the choice of Galerkin subspaces, and from these we will obtain a variety of results, including stability and convergence.

Let V_h^j be a finite dimensional subspace of V^j . This gives rise to spaces of discrete cocycles and discrete coboundaries:

$$\mathfrak{Z}_h^j = \{ \, v \in V_h^j \, | \, dv = 0 \, \} \subset \mathfrak{Z}^j, \quad \mathfrak{B}_h^{j+1} = \{ \, dv \, | \, v \in V_h^j \, \} \subset \mathfrak{B}^{j+1}$$

Given $V_h^{k-1} \subset V^{k-1}$ and $V_h^k \subset V^k$, the associated space of *discrete harmonic* forms is defined in analogy with (4.2) as

$$\mathfrak{H}_{h}^{k} = \{ v \in \mathfrak{Z}_{h}^{k} \, | \, v \perp \mathfrak{B}_{h}^{k} \, \}.$$

Using the spaces V_h^{k-1} , V_h^k , and \mathfrak{H}_h^k , we define the mixed discretization of the Hodge Laplace problem, based on the mixed weak formulation (4.16) in the obvious way. Namely, given $f \in W^k$, we seek $\sigma_h \in V_h^{k-1}$, $u_h \in V_h^k$, and $p_h \in \mathfrak{H}_h^k$ such that

$$\langle \sigma_h, \tau \rangle - \langle u_h, d\tau \rangle = 0, \qquad \tau \in V_h^{k-1}, \langle d\sigma_h, v \rangle + \langle du_h, dv \rangle + \langle p_h, v \rangle = \langle f, v \rangle, \quad v \in V_h^k, \langle u_h, q \rangle = 0, \qquad q \in \mathfrak{H}_h^k.$$
 (5.3)

Note that we generally do *not* expect that \mathfrak{H}_h^j is contained in \mathfrak{H}^j . (In finite element applications to the Hodge Laplacian, the former will consist of piecewise polynomials and the latter of smooth nonpolynomial functions.) When $\mathfrak{H}_h^j \nsubseteq \mathfrak{H}^j$, the

mixed discretization (5.3) is not a standard Galerkin discretization of (4.16) but rather a generalized Galerkin method. Specifically, the exact solution u of (4.16) need not satisfy the final equation of (5.3). This is a consistency error, which we will need to account for in our error analysis. In the language of finite elements, (5.3) is, in this case, a *nonconforming* method. Often, however, we have $\mathfrak{H}_{h}^{k} = \mathfrak{H}^{k} = 0$. In such cases, the mixed discretization is a conforming Galerkin method.

The mixed method is well-defined in the sense that (5.3) has a unique solution, no matter how the subspaces V_h^{k-1} and V_h^k are chosen. To prove this, it suffices to show that the only solution for f = 0 is $\sigma_h = 0$, $u_h = 0$, $p_h = 0$. Choosing $\tau = \sigma_h$, $v = u_h$, and $q = p_h$, and combining the equations, we get that $\sigma_h = 0$ and $du_h = 0$. Thus $u_h \in \mathfrak{Z}_h$. Moreover, since $\sigma_h = 0$, the first equation implies that $u_h \perp \mathfrak{B}_h$. Thus $u_h \in \mathfrak{H}_h$, and so the last equation implies that u_h vanishes. Finally, we take $v = p_h$ in the second equation and conclude that p_h vanishes as well.

Next we turn to the analysis of this mixed method. Although it is not used in the method itself, for the analysis we require the specification of a third space, $V_h^{k+1} \subset V^{k+1}$.

5.2.2 • Three fundamental properties of the subspaces

An important realization arising from FEEC is that there are three properties which the subspaces V_h^j should satisfy in order to guarantee that the mixed Galerkin method provides accurate approximation of the true solution. We now specify these properties.

Approximation property

The first property we need for these spaces is the obviously necessary condition that the subspaces used to approximate u and σ provide good approximation. This can be quantified in various ways. For example, we might assume that for j = k - 1 and k, we have a family of spaces V_h^j , indexed by h decreasing to 0, and require that

$$\lim_{h \to 0} \inf_{v \in V_h^j} \| w - v \|_V = 0, \quad w \in V^j.$$
(5.4)

Or we might require that the best approximation error be $O(h^r)$ for some power r for w belonging to some dense subspace of V^j . Rather than pin down a specific requirement, we will make an appropriate approximation assumption when needed.

Two other properties relate to the complex (5.1), and not just to the individual spaces.

Subcomplex property

We require that $dV_h^{k-1} \subset V_h^k$ and $dV_h^k \subset V_h^{k+1}$ or, in other terminology, that

$$V_h^{k-1} \xrightarrow{d} V_h^k \xrightarrow{d} V_h^{k+1}$$
(5.5)

is a subcomplex of (5.1). The discrete complex (5.5) then is itself a Hilbert complex. As the base spaces we choose $W_h^j = V_h^j$, but with the W inner product, rather than the V inner product. (Of course all inner products are equivalent on V_h^j since it is finite dimensionsal.) Being finite dimensional, (5.5) is a closed Hilbert complex.

As a Hilbert complex in its own right, the discrete complex exhibits all the structures discussed in the previous chapter. The discrete differential

$$d_h^j: V_h^j \to V_h^{j+1}$$

is just the restriction of d^j (for j = k - 1 or k). The kernel of d_h^j is denoted by \mathfrak{Z}_h^j and the range of d_h^{j-1} by \mathfrak{B}_h^j . As for any complex,

$$\mathfrak{B}_h^j \subset \mathfrak{Z}_h^j.$$

The orthogonal complement of the former space inside the latter is the space of discrete harmonic forms \mathfrak{H}_h^k . Since the discrete differential d_h^j is just a restriction of d^j , we have

$$\mathfrak{Z}_h^j\subset\mathfrak{Z}^j,\quad\mathfrak{B}_h^j\subset\mathfrak{B}^j.$$

However, it is generally *not true* that $\mathfrak{H}_h^j \subset \mathfrak{H}_h^j$ since elements of the latter space must be orthogonal to all of \mathfrak{B}^j , not just to \mathfrak{B}_h^j .

The discrete differential d_h^{j-1} has an adjoint operator d_{jh}^* . Since d_h^{j-1} is bounded, its adjoint is everywhere defined, so the spaces V_{jh}^* coincide with $W_h^j = V_h^j$ and

$$d_{jh}^*: V_h^j \to V_h^{j-1}.$$

Unlike for the discrete exterior derivative itself, the discrete adjoint operator is generally *not* the restriction of the continuous adjoint operator d_j^* , and the null spaces \mathfrak{Z}_{jh}^* and ranges \mathfrak{B}_{jh}^* cannot be expected to be contained in their continuous counterparts.

The Hodge decomposition for the discrete complex, which is guaranteed to exist by Theorem 4.5, is called the *discrete Hodge decomposition*:

$$V_h^k = \mathfrak{B}_h^k \oplus \mathfrak{H}_h^k \oplus \mathfrak{B}_{kh}^*.$$
(5.6)

Bounded cochain projections

Finally, we come to the third and most important assumption, which relates the original complex (5.1) to the subcomplex (5.5). We assume that there exists a *bounded cochain projection* π_h from the former to the latter. That it is a *projection* means that for each j, π_h^j is a linear map from V^j onto V_h^j which restricts to the identity on V_h^j . That it is a *cochain* projection means that the following diagram commutes:

$$\begin{array}{cccc} V^{k-1} & \stackrel{d}{\longrightarrow} V^k & \stackrel{d}{\longrightarrow} V^{k+1} \\ & & & \downarrow \pi_h^{k-1} & & \downarrow \pi_h^k & & \downarrow \pi_h^{k+1} \\ V_h^{k-1} & \stackrel{d}{\longrightarrow} V_h^k & \stackrel{d}{\longrightarrow} V_h^{k+1} \end{array}$$

Finally, that π_h is *bounded* could mean one of two things. It might refer to boundedness in V norm: $\|\pi_h v\|_V \leq c \|v\|_V$ for all $v \in V^j$. Alternatively, we might require W-boundedness: $\|\pi_h v\| \leq c \|v\|$. Since π_h is a cochain map, W-boundedness implies V-boundedness:

$$\|\pi_h v\|_V^2 = \|\pi_h v\|^2 + \|d\pi_h v\|^2 = \|\pi_h v\|^2 + \|\pi_h dv\|^2 \le c\|v\|^2 + \|dv\|^2 = c\|v\|_V^2,$$

where we used W-boundedness to obtain the inequality. For now we require only the weaker assumption of V-boundedness, although this will change when we consider improved estimates. In either case, if we are considering a family of subcomplexes indexed by h, we will want the bounds to be uniform in h in the sense that a single constant c works for all h.

A simple, but important, property of a bounded projection is that it provides quasi-optimal approximation, in the sense that, up to a constant multiple, the approximation of an element by its projection is bounded by the best approximation in the space. That is,

$$\|v - \pi_h v\|_V = \inf_{w \in V_h^j} \|(v - w) - \pi_h (v - w)\|_V \le c \inf_{w \in V_h^j} \|v - w\|_V, \quad v \in V^j,$$
(5.7)

with $c = ||I - \pi_h||_{\mathcal{L}(V,V)} = ||\pi_h||_{\mathcal{L}(V,V)}$. (This last identity goes back to [38]. See also [69].) In the rest of this chapter, we derive a series of important further conclusions from these three hypotheses.

5.2.3 • Consistency, stability, and convergence of the mixed method

The gap between \mathfrak{H} and \mathfrak{H}_h

First we show that the cohomology space \mathcal{H}_{h}^{k} for the subcomplex (5.5) is isomorphic to that of the original complex (5.1).

Theorem 5.1 (isomorphism of cohomology). Assume that the complex (5.5) is a finite dimensional subcomplex of the (5.1) which admits a V-bounded cochain projection π_h . In addition, assume the approximation property

$$\|q - \pi_h q\| < \|q\|, \quad 0 \neq q \in \mathfrak{H}^k.$$
(5.8)

Then π_h induces an isomorphism from \mathcal{H}^k onto \mathcal{H}^k_h .

Proof. As discussed in Section 2.4, the cochain map π_h induces a map $\mathcal{H}^k \to \mathcal{H}_h^k$. Since π_h is a projection, the induced map on cohomology is a surjection by Proposition 2.1. Hence, we need only show that it is an injection. Since \mathcal{H}^k consists of equivalence classes of elements of \mathfrak{Z}^k modulo \mathfrak{B}^k , and \mathcal{H}_h^k consists of equivalence classes of elements of \mathfrak{Z}_h^k modulo \mathfrak{B}_h^k , we need to prove that if $z \in \mathfrak{Z}^k$ and $\pi_h z \in \mathfrak{B}_h^k$, then $z \in \mathfrak{B}^k$. For $z \in \mathfrak{Z}^k$ the Hodge decomposition takes the form z = q + b with $q \in \mathfrak{H}_h^k$ and $b \in \mathfrak{B}^k$, so our goal is to show that q vanishes. Since π_h is a cochain map, $\pi_h b \in \mathfrak{B}_h^k$, and we assumed that $\pi_h z \in \mathfrak{B}_h^k$ as well, so $\pi_h q \in \mathfrak{B}_h^k \subset \mathfrak{B}^k$. But q, being harmonic, is orthogonal to \mathfrak{B}^k , so $q \perp \pi_h q$, and, by the Pythagorean theorem,

$$\|q - \pi_h q\|^2 = \|q\|^2 + \|\pi_h q\|^2 \ge \|q\|^2.$$
(5.9)

If $q \neq 0$, then (5.8) and (5.9) lead to ||q|| > ||q||, a contradiction, and we conclude that q = 0 as desired.

Since the homology spaces are isomorphic to the spaces of harmonic forms on both the continuous and the discrete levels, this result tells us that the continuous and discrete harmonic form spaces, \mathfrak{H}^k and \mathfrak{H}^k_h , are isomorphic. The next result

shows not only that they are isomorphic but also that they are close in the sense that any element of one of the spaces can be well-approximated by an element of the other. This will enable us to control the consistency error coming from the nonconformity of the mixed method when the two spaces are not equal.

Theorem 5.2 (gap between harmonic forms). Assume that the complex (5.5) is a finite dimensional subcomplex of (5.1) which admits a V-bounded cochain projection π_h . Then

$$\|(I - P_{\mathfrak{H}_h})q\|_V \le \|(I - \pi_h^k)q\|_V, \quad q \in \mathfrak{H}^k,$$
(5.10)

$$\|(I - P_{\mathfrak{H}})q\|_{V} \le \|(I - \pi_{h}^{k})P_{\mathfrak{H}}q\|_{V}, \quad q \in \mathfrak{H}_{h}^{k}.$$
(5.11)

Proof. For $q \in \mathfrak{H}$, $P_{\mathfrak{H}_h}q = P_{\mathfrak{H}_h}q$, since $\mathfrak{Z}_h = \mathfrak{H}_h \oplus \mathfrak{B}_h$ and $q \perp \mathfrak{B} \supset \mathfrak{B}_h$. This means that $P_{\mathfrak{H}_h}q$ is the element of \mathfrak{Z}_h closest to q as measured in the V norm. But $\pi_h q$ also belongs to \mathfrak{Z}_h , since π_h is a cochain map. Therefore, (5.10) holds.

If $q \in \mathfrak{H}_h \subset \mathfrak{Z}_h \subset \mathfrak{Z}$, the Hodge decomposition gives us $q - P_{\mathfrak{H}}q \in \mathfrak{B}$, so $\pi_h(q - P_{\mathfrak{H}}q) \in \mathfrak{B}_h$, and so is orthogonal to both the discrete harmonic form q and the harmonic form $P_{\mathfrak{H}}q$. Therefore

$$\|q - P_{\mathfrak{H}}q\|_{V} \le \|q - P_{\mathfrak{H}}q - \pi_{h}(q - P_{\mathfrak{H}}q)\|_{V} = \|(I - \pi_{h})P_{\mathfrak{H}}q\|_{V}.$$

The *gap* between two closed subspaces E and F of a Hilbert space V is defined [56, Chapter IV, Section 2.1]

$$gap(E,F) = \max\left(\sup_{\substack{u \in E \\ \|u\|=1}} \inf_{v \in F} \|u - v\|_V, \sup_{\substack{v \in F \\ \|v\|=1}} \inf_{u \in E} \|u - v\|_V\right).$$
(5.12)

An equivalent expression is $gap(E, F) = ||P_E - P_F||$, where $P_E, P_F : V \rightarrow V$ are the orthogonal projections onto E and F [56, footnote 1, p. 198]. From Theorem 5.2, we immediately get an estimate of the gap between the spaces of harmonic and discrete harmonic forms:

$$\operatorname{gap}(\mathfrak{H}, \mathfrak{H}_h) \leq \sup_{\substack{q \in \mathfrak{H} \\ \|q\|=1}} \|(I - \pi_h)q\|_V.$$

The discrete Poincaré inequality and stability

The *stability* of the mixed method (5.3) is measured by the norm of the discrete solution operator $f \mapsto (\sigma_h, u_h, p_h)$. Since (5.3) is just the mixed weak formulation of the abstract Hodge Laplacian problem associated to the discrete complex (5.5), Theorems 4.7 and 4.8, which apply to any closed Hilbert complex, give the estimate

$$\|\sigma_h\|_V + \|u_h\|_V + \|p_h\| \le c_h \|f\|, \tag{5.13}$$

with the stability constant c_h depending only on the constant in the Poincaré inequality for the discrete complex. Thus the key to stability is a bound on the discrete Poincaré constant. This is the content of the next result, which gives such a bound in terms of the Poincaré constant for the original complex and the bound on the cochain projection. **Theorem 5.3 (discrete Poincaré inequality).** Assume that the complex (5.5) is a finite dimensional subcomplex of (5.1) which admits a V-bounded cochain projection π_h . Then

$$||v||_V \le c_P ||\pi_h^k|| ||dv||_V, \quad v \in \mathfrak{Z}_h^{k\perp} \cap V_h^k,$$

where c_P is the constant appearing in the Poincaré inequality (4.7) and $\|\pi_h^k\|$ denotes the V^k operator norm of π_h^k .

Proof. Given $v \in \mathfrak{Z}_h^{k\perp} \cap V_h^k$, define $z \in \mathfrak{Z}^{k\perp} \cap V^k$ by dz = dv. By Theorem 4.6, $||z||_V \leq c_P ||dv||$. Now, $v - \pi_h z \in V_h^k$ and $d\pi_h z = \pi_h dz = \pi_h dv = dv$ with the three equalities coming from, first, the fact that π_h is a cochain map, second, the definition of z, and, third, the subcomplex property $dV_h^k \subset V_h^{k+1}$ and the fact that π_h^{k+1} is a projection. Thus $v - \pi_h z \in \mathfrak{Z}_h^k$ and so is orthogonal to v. Therefore

$$\|v\|_V^2 = \langle v, \pi_h^k z \rangle_V + \langle v, v - \pi_h^k z \rangle_V = \langle v, \pi_h^k z \rangle_V \le \|v\|_V \|\pi_h^k\| \|z\|_V,$$

so

$$\|v\|_{V} \le \|\pi_{h}^{k}\| \, \|z\|_{V} \le \|\pi_{h}^{k}\|c_{P}\|dv\|_{V}. \qquad \Box$$

As an immediate consequence of this result and Theorem 4.9, we obtain the discrete inf-sup condition and the stability estimate.

Theorem 5.4. Assume that the complex (5.5) is a finite dimensional subcomplex of the (5.1) which admits a V-bounded cochain projection π_h . Then the inf-sup condition for the mixed bilinear form B, given in (4.18), over the subspace $V_h^{k-1} \times V_h^k \times \mathfrak{H}_h^k$ holds, and the inf-sup constant γ_h can be bounded below by a positive constant depending only on c_P and $\|\pi_h^k\|$. Consequently the stability estimate (5.13) holds with a constant c_h depending only on the same quantities.

Error estimates for the mixed method

In the numerical analysis of differential equations, error estimates are obtained from the stability and consistency of the discretization scheme. In the present case, stability comes from Theorem 5.4 and consistency comes from the assumed approximation properties of the space V_h^j and from Theorem 5.2 on the gap between the spaces of continuous and discrete harmonic forms (the latter required in the case \mathfrak{H}_h^k is not contained in \mathfrak{H}^k , so the mixed method is nonconforming). The following theorem gives the resulting error estimates. It was proven in the concrete case of the de Rham complex in [11, Theorem 7.4] and then in the case of general Hilbert complexes (essentially with the same proof) in [13, Theorem 3.9]. We repeat this proof here.

Theorem 5.5 (basic error estimate for the mixed method). Assume that the complex (5.5) is a finite dimensional subcomplex of (5.1) which admits a V-bounded cochain projection π_h . Let $(\sigma, u, p) \in V^{k-1} \times V^k \times \mathfrak{H}^k$ solve the Hodge Laplace problem (4.16) and let $(\sigma_h, u_h, p_h) \in V_h^{k-1} \times V_h^k \times \mathfrak{H}_h^k$ solve the

discrete problem (5.3). Then

$$\|\sigma - \sigma_h\|_V + \|u - u_h\|_V + \|p - p_h\| \le C \left(\inf_{\tau \in V_h^{k-1}} \|\sigma - \tau\|_V + \inf_{v \in V_h^k} \|u - v\|_V + \inf_{q \in V_h^k} \|p - q\|_V + \mu \inf_{v \in V_h^k} \|P_{\mathfrak{B}}u - v\|_V \right),$$

where

$$\mu = \mu_h^k = \sup_{\substack{r \in \mathfrak{H}^k \\ \|r\| = 1}} \|(I - \pi_h^k)r\|$$

and the constant C depends only on c_P and $||\pi_h^k||$.

Proof. With B defined as in (4.18), we first observe that (σ, u, p) satisfies

$$B(\sigma, u, p; \tau_h, v_h, q_h) = \langle f, v_h \rangle - \langle u, q_h \rangle, \quad (\tau_h, v_h, q_h) \in V_h^{k-1} \times V_h^k \times \mathfrak{H}_h^k.$$

Let τ , v, and q be the V-orthogonal projections of σ , u, and p into V_h^{k-1} , V_h^k , and \mathfrak{H}_h^k , respectively. Then, for any $(\tau_h, v_h, q_h) \in V_h^{k-1} \times V_h^k \times \mathfrak{H}_h^k$, we have

$$B(\sigma_{h} - \tau, u_{h} - v, p_{h} - q; \tau_{h}, v_{h}, q_{h})$$

= $B(\sigma - \tau, u - v, p - q; \tau_{h}, v_{h}, q_{h}) + \langle u, q_{h} \rangle$
= $B(\sigma - \tau, u - v, p - q; \tau_{h}, v_{h}, q_{h}) + \langle P_{\mathfrak{H}_{h}}u, q_{h} \rangle$
 $\leq C(\|\sigma - \tau\|_{V} + \|u - v\|_{V} + \|p - q\| + \|P_{\mathfrak{H}_{h}}u\|)(\|\tau_{h}\|_{V} + \|v_{h}\|_{V} + \|q_{h}\|).$

Theorem 5.4 then gives

$$\begin{aligned} \|\sigma_h - \tau\|_V + \|u_h - v\|_V + \|p_h - q\| \\ &\leq C(\|\sigma - \tau\|_V + \|u - v\|_V + \|p - q\| + \|P_{\mathfrak{H}_h}u\|). \end{aligned} (5.14)$$

Using (5.10) and the boundedness of the projection π_h we have

$$\|p - q\| \le \|(I - \pi_h)p\| \le C \inf_{q \in V_h^k} \|p - q\|_V.$$
(5.15)

Next we show that

$$\|P_{\mathfrak{H}_h}u\| \le \mu \|(I-\pi_h)u_{\mathfrak{B}}\|_V$$

Now $u \perp \mathfrak{H}^k$, so $u = u_{\mathfrak{B}} + u_{\perp}$, with $u_{\mathfrak{B}} \in \mathfrak{B}^k$ and $u_{\perp} \in \mathfrak{Z}^{k\perp}$. Since $\mathfrak{H}^k_h \subset \mathfrak{Z}^k$, $P_{\mathfrak{H}_h} u_{\perp} = 0$, and since $\pi_h u_{\mathfrak{B}} \in \mathfrak{B}^k_h$, $P_{\mathfrak{H}_h} \pi_h u_{\mathfrak{B}} = 0$. Let $q = P_{\mathfrak{H}_h} u / || P_{\mathfrak{H}_h} u|| \in \mathfrak{H}^k_h$. By Theorem 5.2, there exists $r \in \mathfrak{H}^k$ (and so $r \perp \mathfrak{B}^k$) with $||r|| \leq 1$ and

$$||q - r|| \le ||(I - \pi_h)r|| \le \sup_{\substack{r \in \mathfrak{H}^k \\ ||r|| = 1}} ||(I - \pi_h)r||.$$

Therefore

$$\begin{aligned} \|P_{\mathfrak{H}_{h}}u\| &= \langle u_{\mathfrak{B}} - \pi_{h}u_{\mathfrak{B}}, q - r \rangle \\ &\leq \|(I - \pi_{h})u_{\mathfrak{B}}\| \sup_{\substack{r \in \mathfrak{H}^{k} \\ \|r\| = 1}} \|(I - \pi_{h})r\| \leq c \, \mu \inf_{v \in V_{h}^{k}} \|P_{\mathfrak{B}}u - v\|_{V}), \quad (5.16) \end{aligned}$$

since π_h is a bounded projection. The theorem follows from (5.14)–(5.16) and the triangle inequality.

Theorem 5.5 is typically applied when we have a family of Galerkin subspaces V_h^j indexed by h (usually a mesh size parameter), with uniformly bounded cochain projections. Then the rate of the convergence depends only on approximation properties of the spaces V_h^{k-1} and V_h^k . It is important to note that the third term on the right-hand side of the error estimate is the approximation for p by the full space V_h^k , not just by \mathfrak{H}_h^k . The final term on the right-hand side is typically much smaller than the other terms. If there are no harmonic forms, it is not present at all. Otherwise, it is the product of two terms, both of which are determined by the approximation properties of V_h^k .

The left-hand side of the error estimate couples together the W norm of various solution quantities: σ , $d\sigma$, u, du, and p. It is possible to get sharper estimates by considering them individually. An analogous situation (actually a special case) is the error estimate for the standard piecewise linear finite element method for Poisson's equation, which takes the form

$$||u - u_h||_{H^1} \le C \inf_{v \in V_h} ||u - v||_{H^1} \le Ch ||u||_{H^2}.$$

The left-hand side combines the L^2 error of u with that of grad u and so gives a suboptimal rate for the L^2 error in u. If we apply the Aubin–Nitsche duality argument, however, we can prove the sharper estimate $||u - u_h|| \le Ch^2 ||u||_{H^2}$.

As a simple example of such an improved error estimates for the mixed Hodge Laplacian, consider the error in $d\sigma$ measured in the norm of W^k . From the relation between the mixed formulation and the Hodge decomposition, we have

$$d\sigma_h = P_{\mathfrak{B}_h} f = P_{\mathfrak{B}_h} P_{\mathfrak{B}} f = P_{\mathfrak{B}_h} d\sigma.$$

Since $P_{\mathfrak{B}_h} d\sigma$ is the best approximation of $d\sigma$ from \mathfrak{B}_h and $\pi_h d\sigma$ is another element of \mathfrak{B}_h , we have

$$\|d\sigma - d\sigma_h\| = \|d\sigma - P_{\mathfrak{B}_h} d\sigma\| \le \|d\sigma - \pi_h d\sigma\|.$$

Assuming that the π_h^k is W-bounded, we obtain

$$\|d\sigma - d\sigma_h\| \le \|\pi_h^k\|_{\mathcal{L}(W,W)} \inf_{v \in V_h^k} \|d\sigma - v\|,$$

just as in (5.7). Thus, approximation of $d\sigma$ by $d\sigma_h$ is, up to a constant factor, as good as the best approximation afforded by the subspace V_h^k where $d\sigma_h$ lives.

Theorem 5.6 below, which is taken from [13], gives the full set of such improved error estimates for the mixed method for the Hodge Laplacian. The theorem requires not just that the Hilbert complex on the continuous level be closed but rather that it satisfy the compactness property (see Definition 4.3), which is a stronger assumption. Moreover, it requires that the discrete subcomplex admit not just V-bounded cochain projections but W-bounded. The theorem is stated in terms of the best approximation error in the W^j norm by elements of V_h^j for various quantities $w \in W^j$, for which we introduce the notation

$$E(w) = E_h^j(w) = \inf_{v \in V_h^j} ||w - v||, \quad w \in W^j.$$

Finally, the statement of the theorem uses the notation

$$\mu = \mu_h^k = \|(I - \pi_h) P_{\mathfrak{H}}\|_{\mathcal{L}(W,W)}, \quad \delta = \delta_h^k = \|(I - \pi_h) K\|_{\mathcal{L}(W,W)},$$

$$\eta = \eta_h^k = \max_{j=k-1,k} [\|(I - \pi_h^j) d^* K\|_{\mathcal{L}(W,W)}, \|(I - \pi_h^{j+1}) dK\|_{\mathcal{L}(W,W)}],$$
Box 5.1. Summary of Chapter 5 on approximation of Hilbert complexes.

If $V^{k-1} \xrightarrow{d} V^k \xrightarrow{d} V^{k+1}$ is a segment of the domain complex of a closed Hilbert complex, we approximate the associated Hodge Laplacian problem by discretizing the mixed weak formulation via Galerkin's method with finite dimensional subspaces $V_h^j \subset V^j$.

The key requirements on the Galerkin subspaces are (1) they afford good approximation, (2) they form a subcomplex $(dV_h^j \subset V_h^{j+1})$, and (3) they admit a *bounded cochain projection*, i.e., bounded linear projections $V^j \to V_h^j$ which commute with the d^j .

If these requirements are satisfied, then the discrete complex is itself a Hilbert complex with the same cohomology, the space of discrete harmonic forms is close to that of the true harmonic forms, and the Poincaré constant for the discrete complex is bounded by the true Poincaré constant. It follows that the Galerkin method is consistent and stable and converges with the rate of the best approximation. Assuming also a compactness condition and boundedness of the cochain projections in W norm, we further obtain optimal order error estimates for each variable separately with sharp regularity requirements.

of which the first already appeared in Theorem 5.5. The operator K here is the solution operator for the Hodge Laplacian defined just after Remark 4.10. Due to the compactness property, K, dK, and d^*K are compact operators on W. Together with an approximation property, which ensures that π_h converges pointwise to the identity on W as $h \to 0$, this implies that $\eta, \delta, \mu \to 0$ with h.

Theorem 5.6 (improved error estimates for the mixed method). Assume that the Hilbert complex (5.1) satisfies the compactness property and that (5.5) is a finite dimensional subcomplex which admits a W-bounded cochain projection π_h . Let $(\sigma, u, p) \in V^{k-1} \times V^k \times \mathfrak{H}^k$ solve the Hodge Laplace problem (4.16) and let $(\sigma_h, u_h, p_h) \in V_h^{k-1} \times V_h^k \times \mathfrak{H}_h^k$ solve the discrete problem (5.3). Then

$$\|d(\sigma - \sigma_h)\| \le CE(d\sigma),\tag{5.17}$$

$$\|\sigma - \sigma_h\| \le C[E(\sigma) + \eta E(d\sigma)],\tag{5.18}$$

$$||p - p_h|| \le C[E(p) + \mu E(d\sigma)],$$
 (5.19)

$$||d(u - u_h)|| \le C(E(du) + \eta[E(d\sigma) + E(p)]),$$
(5.20)

$$||u - u_h|| \le C(E(u) + \eta[E(du) + E(\sigma)]$$
(5.21)

$$+ (\eta^2 + \delta)[E(d\sigma) + E(p)] + \mu E(P_{\mathfrak{B}}u)]),$$

where the constant C depends on the W-bounds on π_h^{k-1} , π_h^k , and π_h^{k+1} but otherwise is independent of the choice of subcomplex.

The proof of this theorem is given in [13]. It requires several pages and will not be repeated here.

Chapter 6 Basic notions of exterior calculus

We have thus far studied Hilbert complexes in the abstract and mostly illustrated the concepts via the de Rham complex on a domain in three dimensions. In that example the complex has four spaces and three differential operators, which are the vector calculus operators gradient, curl, and divergence, whose properties we studied individually, employing very analogous arguments for each. In this second part of the book, we will focus on the de Rham complex in its most natural setting, on an *n*-dimensional manifold. The complex then has n + 1 spaces connected by *n* distinct operators. Using the language of exterior calculus we can introduce all these operators in a uniform way, obtaining a theory which is at once more elegant, more general, and more insightful. In this chapter, we present the relevant exterior calculus, culminating in the L^2 de Rham complex on a Riemannian manifold. Then, in the following chapter, we turn to the discretization of the corresponding Hodge Laplace boundary value problems using finite elements.

6.1 • Exterior algebra

The primary objects of study in exterior calculus are differential k-forms, which we take up in the next section. These are functions whose values are alternating multilinear forms, also called algebraic k-forms. The exterior calculus is naturally concerned with the operations of differentiation and integration of these functions. But first we establish the basic terminology and properties of the algebraic k-forms which arise at each point, just as, before studying vector calculus, we must first understand finite dimensional vector spaces. This is the subject of exterior algebra, initiated by Grassman in the middle of the 19th century. Here, we briefly review the definitions and main properties, all of which are elementary.

6.1.1 • Multilinear forms

For a finite dimensional vector space V, and a nonnegative integer k, we consider the vector space $\operatorname{Lin}^k V$ consisting of all k-linear maps

$$\omega: \underbrace{V \times \cdots \times V}^{k \text{ times}} \to \mathbb{R}.$$

By convention, we understand $\operatorname{Lin}^0 V$ to be the space of scalars \mathbb{R} . If $\omega \in \operatorname{Lin}^j V$ and $\mu \in \operatorname{Lin}^k V$, the tensor product $\omega \otimes \mu \in \operatorname{Lin}^{j+k} V$ is given by

$$(\omega \otimes \mu)(v_1, \dots, v_{j+k}) = \omega(v_1, \dots, v_j)\mu(v_{j+1}, \dots, v_{j+k}), \quad v_1, \dots, v_{j+k} \in V.$$

From the definition, the space $\operatorname{Lin}^1 V$ is just the dual space V'. We refer to its elements as either *covectors* or 1-forms. If we choose a basis u_1, \ldots, u_n of V $(n = \dim V)$, this induces a dual basis u^1, \ldots, u^n of V', defined by $u^i(u_j) = \delta_{ij}$. A general element of $\operatorname{Lin}^k V$ is determined by the value it assigns to each of the k-tuples $(u_{\sigma_1}, \ldots, u_{\sigma_k})$ with the σ_i integers between 1 and n. Thus dim $\operatorname{Lin}^k V = n^k$. The element of $\operatorname{Lin}^k V$ which takes the k-tuple $(u_{\sigma_1}, \ldots, u_{\sigma_k})$ to 1 and all the other k-tuples $(u_{\tau_1}, \ldots, u_{\tau_k})$ to 0 is simply $u^{\sigma_1} \otimes \cdots \otimes u^{\sigma_k}$, and so the set of all of these gives a basis for $\operatorname{Lin}^k V$.

The case $V = \mathbb{R}^n$ differs from the general case only in that there is a canonical basis e_1, \ldots, e_n . As notation for the corresponding dual basis we use dx^1, \ldots, dx^n . Thus, dx^i is simply the linear functional which assigns to each vector \mathbb{R}^n its *i*th component. The reason for the odd notation $(dx^i \text{ rather than, say, } e^i)$ will become clear as we develop exterior calculus.

6.1.2 • Alternating multilinear forms

Multilinear forms $\omega \in \operatorname{Lin}^k V$ that change sign under the interchange of two variables are called *skew-symmetric* or *alternating*:

$$\omega(v_1, \dots, v_i, \dots, v_j, \dots, v_k) = -\omega(v_1, \dots, v_j, \dots, v_i, \dots, v_k),$$

$$1 \le i < j \le n, \quad v_1, \dots, v_k \in V.$$

In particular, an alternating form vanishes if any of its arguments is repeated. To any multilinear form ω there is associated the alternating form which is its skew-symmetric part:

$$(\operatorname{skw}\omega)(v_1,\ldots,v_k) = \frac{1}{k!} \sum_{\sigma \in \Sigma_k} \operatorname{sign}(\sigma)\omega(v_{\sigma_1},\ldots,v_{\sigma_k})$$

where the sum is over all the permutations of the integers 1 to k. For example, if $\omega \in \operatorname{Lin}^3 V$, then

$$(\operatorname{skw}\omega)(v_1, v_2, v_3) = \frac{1}{6} [\omega(v_1, v_2, v_3) + \omega(v_2, v_3, v_1) + \omega(v_3, v_1, v_2) - \omega(v_1, v_3, v_2) - \omega(v_3, v_2, v_1) - \omega(v_2, v_1, v_3)].$$

We define the space $\operatorname{Alt}^k V$ as the space of all skew-symmetric k-linear forms. The spaces $\operatorname{Alt}^0 V$ and $\operatorname{Alt}^1 V$ coincide with $\operatorname{Lin}^0 V$ and $\operatorname{Lin}^1 V$, i.e., are just the space of scalars and the space of covectors, respectively. An element of $\operatorname{Alt}^k V$ is determined by the value it assigns to each of the k-tuples $(e_{\sigma_1}, \ldots, e_{\sigma_k})$ with $1 \le \sigma_1 < \cdots < \sigma_k \le n$, so dim $\operatorname{Alt}^k V = \binom{n}{k}$. For k > n, $\operatorname{Alt}^k V = 0$.

The *exterior product* or wedge product of $\omega \in \operatorname{Alt}^j V$ and $\mu \in \operatorname{Alt}^k V$ is given by

$$(\omega \wedge \mu)(v_1, \dots, v_{j+k}) = \sum_{\sigma} \operatorname{sign}(\sigma) \omega(v_{\sigma_1}, \dots, v_{\sigma_j}) \mu(v_{\sigma_{j+1}}, \dots, v_{\sigma_{j+k}}),$$

where the sum is over $\sigma \in \Sigma_{j+k}$ for which $\sigma_1 < \cdots < \sigma_j$ and $\sigma_{j+1} < \cdots < \sigma_{j+k}$. Equivalently,

$$\omega \wedge \mu = \binom{j+k}{j} \operatorname{skw}(\omega \otimes \mu).$$

The exterior product satisfies the anticommutativity law

$$\omega \wedge \mu = (-)^{jk} \mu \wedge \omega, \quad \omega \in \operatorname{Alt}^j V, \ \mu \in \operatorname{Alt}^k V.$$

Let $\{u_i\}$ be a basis for V and $\{u^i\}$ the corresponding dual basis. For given $1 \leq \sigma_1 < \cdots < \sigma_k \leq n$, the element of $\operatorname{Alt}^k V$ which takes the k-tuple $(u_{\sigma_1}, \ldots, u_{\sigma_k})$ to 1 and takes all the other k-tuples with increasing indices to 0 is

$$u^{\sigma_1} \wedge \dots \wedge u^{\sigma_k}. \tag{6.1}$$

The set of all of such form a basis for $\operatorname{Alt}^k V$, and so the general alternating k-form on V is a linear combination of these. In particular, the general alternating k-form on \mathbb{R}^n can be written

$$\sum_{1 \le \sigma_1 < \dots < \sigma_k \le n} a_\sigma \, dx^{\sigma_1} \wedge \dots \wedge dx^{\sigma_k}$$

for a unique choice of coefficients $a_{\sigma} \in \mathbb{R}$.

6.1.3 • Operations on alternating forms

Contraction

Besides the vector space operations on $\operatorname{Alt}^k V$ and the exterior product, there are several other useful algebraic operations. Given $\omega \in \operatorname{Alt}^k V$ with $k \ge 1$ and $v \in V$, the *contraction* of ω and v, also known as the interior product, is the (k-1)-form

$$\omega \lrcorner v(v_1, \ldots, v_{k-1}) = \omega(v, v_1, \ldots, v_{k-1}).$$

(If k = 0, $\omega \lrcorner v := 0$.) Since ω is alternating, $(\omega \lrcorner v) \lrcorner v$ always vanishes. Moreover, it is easy to check that contraction with a vector is an antiderivation in the sense that

$$(\omega \wedge \eta) \lrcorner v = (\omega \lrcorner v) \wedge \eta + (-1)^k \omega \wedge (\eta \lrcorner v), \quad \omega \in \operatorname{Alt}^k V, \ \eta \in \operatorname{Alt}^l V.$$

Inner product

If the vector space V is endowed with an *inner product* $\langle \cdot, \cdot \rangle$, then this induces the Riesz isomorphism V to V' and so an inner product on V'. An *inner product* on Alt^k V is then defined by

$$\langle u^1 \wedge \dots \wedge u^k, v^1 \wedge \dots \wedge v^k \rangle = \det \left[\left(\langle u^i, v^j \rangle \right)_{i,j=1}^n \right],$$

which extends to general elements of $\operatorname{Alt}^k V$ by bilinearity. If the u_i form an *orthonormal* basis for V, then the corresponding dual basis elements u^i are orthonormal in V', as are the resulting basis functions (6.1) of $\operatorname{Alt}^k V$.

Volume form

When k is equal to the dimension n of V, the space $\operatorname{Alt}^n V$ is one-dimensional, and so a basis consists of a single element. If the space V is endowed with both an inner product and an *orientation*, then there is a canonical choice of basis element. The ordered bases of a finite dimensional vector space V divide into two classes according to whether the change-of-basis matrix connecting two of them has a positive or negative determinant. We orient the vector space by designating the ordered bases of one class as the positively oriented ones. (In \mathbb{R}^3 , it is conventional to take as the positively oriented bases those which satisfy the "right-hand rule.") If $\omega \in \operatorname{Alt}^n V$, then the value of $\omega(u_1, \ldots, u_n)$ is the same for any orthonormal positively oriented basis u_1, \ldots, u_n , and we determine a unique element of $\operatorname{Alt}^n V$ by requiring this value to be 1. The resulting alternating *n*form is called the *volume form* on V, which we typically denote by vol. It may be expressed in terms of the dual basis to any positively oriented orthonormal basis as vol = $u^1 \wedge \cdots \wedge u^n$. On \mathbb{R}^n , the canonical basis e_1, \ldots, e_n is orthonormal and positively oriented, so

$$\mathsf{vol} = dx^1 \wedge \dots \wedge dx^n$$

Hodge star

We continue with an oriented inner product space V. Let $\omega \in \operatorname{Alt}^k V$ for some $0 \le k \le n$, and let $\mu \in \operatorname{Alt}^{n-k} V$. Then $\omega \wedge \mu = c$ vol for some real number c, and the map $\mu \mapsto c$ is a linear functional on Alt^{n-k} . By the Riesz representation theorem it is given by the inner product of μ with some other element of $\operatorname{Alt}^{n-k} V$. That element depends on ω and is called the *Hodge star* or Hodge dual of ω , denoted $\star \omega$. The Hodge star is thus characterized by the relation

$$\omega \wedge \mu = \langle \star \omega, \mu \rangle \text{vol}, \quad \omega \in \operatorname{Alt}^k V, \ \mu \in \operatorname{Alt}^{n-k} V.$$
(6.2)

For each k, we thus have a linear map \star : Alt^k $V \to$ Alt^{n-k} V. If we choose a positively oriented orthonormal basis of 1-forms, u^1, \ldots, u^n , so vol $= u^1 \wedge \cdots \wedge u^n$, we find that

$$\star u^{\sigma_1} \wedge \dots \wedge u^{\sigma_k} = \pm u^{\tau_1} \wedge \dots \wedge u^{\tau_{n-k}}$$

where σ and τ are increasing sequences which partition $\{1, \ldots, n\}$ and the sign is that of the permutation (σ, τ) . For example, in \mathbb{R}^3 ,

$$\star dx^1 = dx^2 \wedge dx^3, \quad \star dx^2 = -dx^1 \wedge dx^3, \quad \star dx^3 = dx^1 \wedge dx^2.$$

Note that \star is an *isometry* and $\star \star \omega = (-)^{k(n-k)}\omega$, or, equivalently, if $\omega \in \operatorname{Alt}^k V$, then

 $\star \star \omega = \omega$ in odd dimensions, $\star \star \omega = (-)^k \omega$ in even dimensions.

Combining this with (6.2) gives an alternate characterization of the Hodge star:

$$\omega \wedge \star \nu = \langle \omega, \nu \rangle \text{vol}, \quad \omega, \nu \in \text{Alt}^k V.$$
(6.3)

Scalar and vector proxies

If V is an oriented inner product space of dimension n, then there is a natural isomorphism of each of the spaces $\operatorname{Alt}^0 V$ and $\operatorname{Alt}^n V$ onto the scalar space \mathbb{R} ,

and of each of the spaces $\operatorname{Alt}^1 V$ and $\operatorname{Alt}^{n-1} V$ onto the vector space V. In this way we can associate either a *scalar proxy* or a *vector proxy* to each member of $\operatorname{Alt}^k V$ for k = 0, 1, n-1, or n (and so for all cases if dim $V \leq 3$). We now describe these isomorphisms.

For any vector space V, $\operatorname{Alt}^0 V$ is just the space \mathbb{R} of scalars and the isomorphism is the identity. The inner product on V defines the isomorphism of $\operatorname{Alt}^1 V = V'$ onto V via the Riesz representation. The Hodge star then maps $\operatorname{Alt}^{n-1} V$ isomorphically to $\operatorname{Alt}^1 V$ (which is isomorphic to V) and similarly maps $\operatorname{Alt}^n V$ to $\operatorname{Alt}^0 V = \mathbb{R}$.

The case n = 2 is particular, since in that case we have two ways of associating a vector proxy with an alternating 1-form: either as a 1-form by applying the Riesz map or as an (n-1)-form, that is, by applying first the Hodge star and then the Riesz map. These do not coincide, since the Hodge star $\operatorname{Alt}^1 V \to \operatorname{Alt}^1 V$ is not the identity map: it takes u^1 to u^2 and u^2 to $-u^1$. Therefore the two vector proxies differ by a rotation through a right angle.

Pullback under linear maps

If $L: V \to W$ is a linear map between finite dimensional vector spaces, then there is an obvious *pullback* operation $L^*: \operatorname{Alt}^k W \to \operatorname{Alt}^k V$ given by

$$L^*\omega(v_1,\ldots,v_k) = \omega(Lv_1,\ldots,Lv_k).$$
(6.4)

The pullback acts contravariantly: if $K : W \to X$ is another map, then $(K \circ L)^* = L^* \circ K^*$. It also respects the algebraic structure, including the exterior product:

$$L^*(\omega \wedge \mu) = L^*\omega \wedge L^*\mu. \tag{6.5}$$

In the case when V is a subspace of W and L is the inclusion, the pullback of a k-form on W is its restriction to $V \times \cdots \times V$. Assuming that W has an inner product, we may pass to vector proxies to view the pullback $\operatorname{Alt}^1 W \to \operatorname{Alt}^1 V$ as a map from $W \to V$, which is nothing other than the orthogonal projection. Another example, is when W is an oriented inner product space of dimension n and V an oriented subspace of dimension n-1. In that case, we may use proxies to view the pullback $\operatorname{Alt}^{n-1} W \to \operatorname{Alt}^{n-1} V$ as a map from $W \to \mathbb{R}$. Let us calculate this map. We start with a positive orthonormal basis u_2, \ldots, u_n of V and prepend to it the unit normal u_1 of V in W with the sign chosen so that u_1, \ldots, u_n is positively oriented. Let u^1, \ldots, u^n denote the dual basis of W'. The pullback of u^1 to V' vanishes, while the pullbacks v^2, \ldots, v^n of the remaining u^i give the basis of V' dual to u_2, \ldots, u_n . Since pullback respects the exterior product, it follows that

$$L^*(u^2 \wedge \dots \wedge u^n) = v^2 \wedge \dots \wedge v^n = \operatorname{vol}_V, \quad L^*(u^1 \wedge u^2 \wedge \dots \wedge \widehat{u^i} \wedge \dots \wedge u^n) = 0.$$

Transferred to proxies, the pullback map L^* thus becomes

$$u_1 \mapsto 1, \qquad u_i \mapsto 0, \quad i = 2, \dots, n.$$

This map is simply the inner product of a vector in W with the oriented unit normal u_1 . To summarize, in terms of proxies,

- the pullback of a 1-form to a subspace is the orthogonal projection, and
- the pullback of an n-1 form to a subspace of codimension 1 is the inner product with the normal.

Proxies					
$\operatorname{Alt}^0 \mathbb{R}^3 = \mathbb{R}$	$ c \leftrightarrow c$				
$\operatorname{Alt}^1 \mathbb{R}^3 \xrightarrow{\cong} \mathbb{R}^3$	$u_1 dx^1 + u_2 dx^2 + u_3 dx^3 \leftrightarrow u$				
$\operatorname{Alt}^2 \mathbb{R}^3 \xrightarrow{\cong} \mathbb{R}^3$	$ u_1 dx^2 \wedge dx^3 - u_2 dx^1 \wedge dx^3 + u_3 dx^1 \wedge dx^2 \leftrightarrow u $				
$\operatorname{Alt}^3 \mathbb{R}^3 \xrightarrow{\cong} \mathbb{R}$	$c \leftrightarrow c dx^1 \wedge dx^2 \wedge dx^3$				
Exterior product					
$ \begin{array}{c} \wedge: \operatorname{Alt}^1 \mathbb{R}^3 \times \operatorname{Alt}^1 \mathbb{R}^3 \to \operatorname{Alt}^2 \mathbb{R}^3 \\ \wedge: \operatorname{Alt}^1 \mathbb{R}^3 \times \operatorname{Alt}^2 \mathbb{R}^3 \to \operatorname{Alt}^3 \mathbb{R}^3 \end{array} $	$\begin{vmatrix} \times : \mathbb{R}^3 \times \mathbb{R}^3 \to \mathbb{R}^3 \\ \cdot : \mathbb{R}^3 \times \mathbb{R}^3 \to \mathbb{R} \end{vmatrix}$				
Contraction with a vector $v \in \mathbb{R}^3$					
	$\begin{vmatrix} v \cdot : \mathbb{R}^3 \to \mathbb{R} \\ v \times : \mathbb{R}^3 \to \mathbb{R}^3 \\ v : \mathbb{R} \to \mathbb{R}^3 (c \mapsto cv) \end{vmatrix}$				
Inner product					
Inner product on $\operatorname{Alt}^k \mathbb{R}^3$	Euclidean inner product on \mathbb{R} or \mathbb{R}^3				
Volume form					
$vol = dx^1 \wedge dx^2 \wedge dx^3$	$\left \left(v_1, v_2, v_3 \right) \mapsto \det[v_1 v_2 v_3] \right.$				
Hodge star					
$ \begin{array}{c} \star : \operatorname{Alt}^{0} \mathbb{R}^{3} \to \operatorname{Alt}^{3} \mathbb{R}^{3} \\ \star : \operatorname{Alt}^{1} \mathbb{R}^{3} \to \operatorname{Alt}^{2} \mathbb{R}^{3} \\ \star : \operatorname{Alt}^{2} \mathbb{R}^{3} \to \operatorname{Alt}^{1} \mathbb{R}^{3} \\ \star : \operatorname{Alt}^{2} \mathbb{R}^{3} \to \operatorname{Alt}^{1} \mathbb{R}^{3} \end{array} $	$ \begin{aligned} & \operatorname{id} : \mathbb{R} \to \mathbb{R} \\ & \operatorname{id} : \mathbb{R}^3 \to \mathbb{R}^3 \\ & \operatorname{id} : \mathbb{R}^3 \to \mathbb{R}^3 \\ & \operatorname{id} : \mathbb{R} \to \mathbb{R} \end{aligned} $				
Pullback by a linear map $L: \mathbb{R}^3 \to \mathbb{R}^3$					
$L^* : \operatorname{Alt}^0 \mathbb{R}^3 \to \operatorname{Alt}^0 \mathbb{R}^3$ $L^* : \operatorname{Alt}^1 \mathbb{R}^3 \to \operatorname{Alt}^1 \mathbb{R}^3$ $L^* : \operatorname{Alt}^2 \mathbb{R}^3 \to \operatorname{Alt}^2 \mathbb{R}^3$ $L^* : \operatorname{Alt}^3 \mathbb{R}^3 \to \operatorname{Alt}^3 \mathbb{R}^3$	$\begin{vmatrix} \operatorname{id} : \mathbb{R} \to \mathbb{R} \\ L^T : \mathbb{R}^3 \to \mathbb{R}^3 \\ (\det L)L^{-1} : \mathbb{R}^3 \to \mathbb{R}^3 \\ (\det L) : \mathbb{R} \to \mathbb{R} (c \mapsto c \det L) \end{vmatrix}$				

Table 6.1. *Exterior algebra of* \mathbb{R}^3 *in terms of scalar and vector proxy fields.*

Exterior algebra of \mathbb{R}^3

Since the spaces $\operatorname{Alt}^0 \mathbb{R}^n$ and $\operatorname{Alt}^n \mathbb{R}^n$ may be identified with \mathbb{R} and the spaces $\operatorname{Alt}^1 \mathbb{R}^n$ and $\operatorname{Alt}^{n-1} \mathbb{R}^n$ with \mathbb{R}^n , in the case of n = 3 (or fewer) dimensions, all the operations of exterior algebra may be viewed as operations on scalars or vectors. These operations are tabulated in Table 6.1, which is adapted from [11]. We remark that the matrix $(\det L)L^{-1}$ which shows up as the pullback of a 2-form in the penultimate line of the table is the adjugate matrix of L, which is defined (as the transpose of the cofactor matrix) even when L is singular.

6.2 • Exterior calculus

The primary objects of study in vector calculus are functions $\Omega \to \mathbb{R}$ and vector fields $\Omega \to \mathbb{R}^n$, defined on a domain $\Omega \subset \mathbb{R}^n$. In exterior calculus on \mathbb{R}^n , the primary objects are differential k-forms which are functions $\Omega \to \operatorname{Alt}^k \mathbb{R}^n$. In particular, the cases of functions and vector fields is included. Indeed, as we have seen, there are two different but natural ways to think of the functions of vector calculus as differential forms (either as 0-forms or *n*-forms), and similarly vector fields may be identified with either 1-forms or (n-1)-forms. This distinction, which is lost in vector calculus, turns out to be a strength of exterior calculus. Maxwell referred to it when he wrote that "physical vector quantities may be divided into two classes, in one of which the quantity is defined with reference to a line, while in the other the quantity is defined with reference to an area" [58]. In higher dimensions other types of fields arise. For example, a differential 2-form in four dimensions is a field with six independent components (the number of independent components in a 4×4 skew-symmetric matrix).

Actually exterior calculus develops most naturally if we are more general, permitting as the domain Ω any manifold (finite dimensional, with or without boundary, and smooth enough that tangent spaces are defined), not just open subsets in \mathbb{R}^n . Much of the theory works in this degree of generality, depending only on the differentiable structure of the domain. Other concepts require additional structure: either an inner product (i.e., a Riemannian manifold), or an orientation (an oriented manifold), or both, and some formulas simplify if we use a preferred basis. Of course, a domain on \mathbb{R}^n is endowed with all these things. But it is enlightening to hold clearly in mind what structures are required for what operations, as we shall do. From a practical point of view, even if we are only interested in solving boundary value problems in \mathbb{R}^n , we will need differential forms not only on the domain where the differential equation is defined but also on its boundary—which is a manifold, but not a domain in Euclidean space.

6.2.1 • Differential forms

Let Ω be a manifold of dimension n. At each point $x \in \Omega$, the tangent space $T_x\Omega$ is a vector space of dimension n, and the collection of all pairs (x, v) with $x \in \Omega$ and $v \in T_x\Omega$ defines the tangent bundle, which is a manifold of dimension 2n. A *section* of this bundle is a function from Ω to its tangent bundle which assigns to each $x \in \Omega$ a tangent vector $v_x \in T_x\Omega$. That is, it is a tangent vector field.

Applying the exterior algebra constructions of the previous section, we may define a bundle (the *k*th exterior power of the cotangent bundle) consisting of pairs (x, μ) with $\mu \in \operatorname{Alt}^k T_x \Omega$. A differential *k*-form is a section of this bundle, i.e., a function which assigns to each $x \in \Omega$ an element of $\operatorname{Alt}^k T_x \Omega$. Thus ω being a differential *k*-form means that

$$\omega_x(v_1,\ldots,v_k) \in \mathbb{R} \ \forall \ x \in \Omega, v_1,\ldots,v_k \in T_x\Omega,$$

with the dependence on the vectors v_i being multilinear and alternating. We will usually require some degree of smoothness for the dependence on x, varying with the context. In the case where Ω is indeed a subdomain of \mathbb{R}^n , we may identify all the tangent spaces with \mathbb{R}^n , and so a differential k-form is just a function $\Omega \to \operatorname{Alt}^k \mathbb{R}^n$. It may be written as

$$\omega = \sum_{1 \le \sigma_1 < \dots < \sigma_k \le n} a_\sigma \, dx^{\sigma_1} \wedge \dots \wedge dx^{\sigma_k}, \tag{6.6}$$

with the coefficients a_{σ} real-valued functions on Ω .

We write $\Lambda^k(\Omega)$, or just Λ^k , for the vector space of differential k-forms on Ω when we do not need to emphasize the exact degree of smoothness expected. When we do, we will precede the notation with a function space. For example, $C\Lambda^k$ denotes the space of continuous differential forms, $C^1\Lambda^k$ of continuously differentiable forms, etc. We will introduce notation for spaces with different specified levels of smoothness below.

Exterior product, contraction, inner product, Hodge star, vector proxies

The algebraic operations on alternating multilinear forms discussed in the previous section can be applied to differential forms pointwise. For example, the exterior product of $\omega \in \Lambda^j$ and $\mu \in \Lambda^k$ is a differential (j + k)-form given by

$$(\omega \wedge \mu)_x = \omega_x \wedge \mu_x.$$

Similarly, if v is a vector field on Ω and $\omega \in \Lambda^k(\Omega)$, then $\omega \lrcorner v$ is a differential (k-1)-form. In order to extend the operations that need an inner product, we require the manifold Ω to be Riemannian, so that at each point x there is given an inner product on $T_x\Omega$. The inner product should depend smoothly on x in the sense that if v_x and w_x are smooth vector fields, then $x \mapsto \langle v_x, w_x \rangle_{T_x\Omega}$ should be a smooth function on Ω (the exact degree of smoothness again depending on the context). Several of the algebraic operations also require an orientation of the tangent spaces $T_x\Omega$ and we suppose that these have been assigned at each point in a consistent fashion. This requires that the manifold be orientable. It is not possible, for example, on a Möbius strip. On an oriented Riemannian manifold, there is a unique volume n-form, characterized by the fact that at each point x it assigns the value 1 to each positively oriented orthonormal basis of $T_x\Omega$. The Hodge star operator from Λ^k to Λ^{n-k} is then defined pointwise by (6.2), and we can compute the inner product in $L^2\Lambda^k(\Omega)$ by

$$\langle \omega, \nu \rangle_{L^2 \Lambda^k} = \int_{\Omega} \langle \omega, \nu \rangle \mathrm{vol} = \int_{\Omega} \omega \wedge \star \nu, \quad \omega, \nu \in L^2 \Lambda^k(\Omega).$$

Applying the vector proxies pointwise, we can identify differential 0-forms and differential *n*-forms with real-valued functions on the manifold and identify differential 1-forms and differential (n - 1)-forms with tangential vector fields.

Pullback under smooth maps and traces

If $\phi: M \to N$ is a differentiable map between manifolds, its derivative $\phi'(x)$ is a linear map of $T_x M$ to $T_{\phi(x)} N$. If $\omega \in \Lambda^k(N)$, we may therefore define a differential k-form on Ω by

$$(\phi^*\omega)_x = \phi'(x)^*\omega_{\phi(x)},\tag{6.7}$$

where, on the right-hand side, we are using the algebraic pullback defined by (6.4). This defines the pullback $\phi^* \omega \in \Lambda^k(M)$ of a differential form under a smooth map. As in the algebraic case, pullback acts contravariantly and respects the exterior product.

The pullback of the inclusion of a submanifold M into another manifold N is called the trace operator, denoted $\operatorname{tr} : \Lambda^k(N) \to \Lambda^k(M)$. In the case of 0-forms it is just restriction of functions. The derivative of the inclusion at a point $x \in M$ is the inclusion of $T_x M \subset T_x N$. This fact and our calculations at the end of Section 6.1.3 provide formulas for the trace in terms of proxies in two important cases. If w is a vector field on N corresponding to some 1-form $\omega \in \Lambda^1(N)$, then the vector field \overline{w} on M corresponding to $\operatorname{tr} \omega \in \Lambda^1(M)$ is given by

$$\bar{w}(x) = \pi w(x), \quad x \in M,$$

where $\pi = \pi_{T_xM} : T_xN \to T_xM$ is the orthogonal projection. If instead $\dim M = \dim N - 1$ and the vector field w corresponds to $\omega \in \Lambda^{\dim N - 1}(N)$,

then tr ω corresponds to the scalar function $\langle w, n \rangle$ on M, where n is the appropriately oriented unit normal to M in N.

In the case $M = \mathbb{R}^{n-1}$, $N = \mathbb{R}^n$, with the inclusion given by $x \mapsto (0, x)$, we can write the corresponding trace operator in terms of coordinates. If ω is the differential form given by (6.6), then

$$\operatorname{tr} \omega = \sum_{\{\sigma \mid \sigma_1 > 1\}} (\operatorname{tr} a_{\sigma}) \, dx^{\sigma_1} \wedge \dots \wedge dx^{\sigma_k}.$$
(6.8)

The trace operator appearing in the term tr a_{σ} on the right-hand side is the 0-form trace, which is just the restriction to \mathbb{R}^{n-1} .

6.2.2 • Exterior derivative

The exterior derivative, which maps $\Lambda^k(\Omega)$ to $\Lambda^{k+1}(\Omega)$, is the fundamental operator of exterior calculus. It is simplest to define when Ω is a domain in \mathbb{R}^n . If ω is a smooth differential k-form, and v_1, \ldots, v_k are k vectors in \mathbb{R}^n , then $\omega_x(v_1, \ldots, v_k)$ is a differentiable function of x, which is alternating multilinear in the k vectors. If v_0 is an additional vector, then we may apply $\partial_{v_0} := v_0 \cdot \text{grad}$, the directional derivative in the direction v_0 , to this function, obtaining a quantity which depends on x and all (k + 1)-vectors. The dependence on the vectors is linear in each but is not alternating. To get a (k + 1)-form we take the alternating part. Multiplying this by k + 1 gives the exterior derivative $d\omega$:

$$(d\omega)_x(v_0,\ldots,v_k) = \sum_{j=0}^k (-)^j \partial_{v_j} \omega_x(v_0,\ldots,\widehat{v_j},\ldots,v_k).$$

Written in terms of the standard coordinates on \mathbb{R}^n , a general element of $\Lambda^k(\Omega)$ can be written as

$$\omega = \sum_{1 \le \sigma_1 < \dots < \sigma_k \le n} f_\sigma \, dx^{\sigma_1} \wedge \dots \wedge dx^{\sigma_k},$$

where the coefficients f_{σ} are functions on Ω . Then

$$d\omega = \sum_{\sigma} \sum_{j=1}^{n} \frac{\partial f_{\sigma}}{\partial x^{j}} dx^{j} \wedge dx^{\sigma_{1}} \wedge \dots \wedge dx^{\sigma_{k}}.$$
 (6.9)

In the case k = 0, for example, we have the formula

$$df = \sum_{j=1}^{n} \frac{\partial f}{\partial x^j} \, dx^j$$

Identifying the 1-form on the right-hand side with a vector field, we see that the exterior derivative on a 1-form on \mathbb{R}^n is just the gradient operator. Similar considerations show that the exterior derivative on an (n - 1)-form can be identified with the divergence operator. In three dimensions, the exterior derivative on 1-forms can be viewed as a mapping from vector fields to vector fields and is found to be the curl operator. Thus, in this case the entire de Rham complex (2.2) can be expressed in the terminology of vector calculus as in (2.3).

An important property of the exterior derivative is that it is invariant under pullback. If $\phi : \Omega \to M$ is a map from a subdomain of \mathbb{R}^n to a subdomain of \mathbb{R}^m , and $\omega \in \Lambda^k(M)$, then $\phi^* \omega \in \Lambda^k(\Omega)$ and

$$d\phi^*\omega = \phi^* d\omega. \tag{6.10}$$

In the case where ϕ is a diffeomorphism between two domains of equal dimension, we conclude that

$$d\omega = (\phi^{-1})^* d\phi^* \omega. \tag{6.11}$$

We may use this formula to define $d\omega$ when M is a manifold rather than a domain in Euclidean space. We choose any chart ϕ mapping some $\Omega \subset \mathbb{R}^n$ onto an open set of M and then use the formula (6.11) to define $d\omega$ on that open subset. Since M is covered by charts, in this way we define $d\omega \in \Lambda^k(M)$. That the definition does not depend on the choice of charts follows from (6.10). Moreover this invariance under pullback holds for any smooth maps between manifolds.

Another essential property of the exterior derivative is that it is a differential in the sense that $d \circ d = 0$. It is enough to verify this on a domain in \mathbb{R}^n , where it follows easily from (6.9) thanks to the commutativity of partial derivatives. Equally essential is the Leibniz rule

$$d(\omega \wedge \mu) = d\omega \wedge \mu + (-)^k \omega \wedge d\mu, \quad \omega \in \Lambda^k(\Omega), \ \mu \in \Lambda^j(\Omega).$$
(6.12)

It is worth emphasizing that the exterior derivative is defined without any reference to an inner product (metric) or an orientation on the manifold. A differential k-form on any manifold can be differentiated to give a differential (k + 1)-form.

6.2.3 Integration of differential forms

A differential k-form can be integrated over an oriented k-dimensional manifold to give a real number. Notice that the form degree and the manifold dimension must be the same. No metric or measure is required on the manifold to define the integral, and if they exist, then the integral is independent of them. However, the orientation does matter. If we switch it, the sign of the integral changes.

To get a basic idea of how the integral is defined, consider a parallelepiped P determined by a point $x \in \mathbb{R}^k$ and k linearly independent vectors v_1, \ldots, v_k , as in Figure 6.1(a). If ω is a constant k-form on P, then the integral of ω over P is simply $\omega_x(v_1, \ldots, v_k)$. If ω is smooth, but not constant, then this will be a good approximation of the integral if the parallelepiped is sufficiently small. Now any manifold can be approximated arbitrarily well by a collection of small tangent parallelopideds using coordinate charts, as illustrated in Figure 6.1(b). We then get an approximate value for the integral by evaluating ω on each parallelopiate and adding

$$\int_{\Omega} \omega \approx \sum_{i} \omega_{x^{(i)}}(v_1^{(i)}, \dots, v_n^{(i)}),$$

where the $x^{(i)}$ and $v_j^{(i)}$ are the points and tangent vectors determining the parallelograms. Passing to the limit we obtain $\int_{\Omega} \omega$.



Figure 6.1. (a) *A parallelogram determined by a point and two tangent vectors at the point.* (b) *A two-dimensional manifold can be approximated arbitrarily closely by a collection of tangent parallelograms.*

One easy consequence of the definition is that the integral is invariant under orientation-preserving diffeomorphisms. That is, if $\phi : M \to N$ is such a diffeomorphism of *n*-dimensional manifolds, then

$$\int_{M} \phi^* \omega = \int_{N} \omega, \quad \omega \in \Lambda^n(N).$$
(6.13)

In particular, we may use charts to reduce the computation of the integral over a manifold to integrals over subdomains of \mathbb{R}^n , and most of its properties can be verified in this way. If Ω is such a domain, then an *n*-form on Ω can be written $f(x^1, \ldots, x^n) dx^1 \wedge \cdots \wedge dx^n$, and its integral is exactly what the notation

$$\int_{\Omega} f(x^1, \dots, x^n) \, dx^1 \wedge \dots \wedge dx^n$$

suggests (i.e., it equals the Lebesgue integral of the function f over Ω). More generally, if Ω is an oriented Riemannian manifold, then $\int_{\Omega} f$ vol coincides with the integral of f with respect to its associated Riemannian measure.

If ω is a k-form on an n-dimensional manifold Ω , and M is an oriented kdimensional submanifold of Ω , then the trace of ω on M is a k-form on a kdimensional manifold and so can be integrated over M. Technically, the result should be written as $\int_M \operatorname{tr}_M \omega$, but frequently one just writes $\int_M \omega$. For example, suppose M is an oriented curve in some domain in \mathbb{R}^n and ω is a 1-form on the domain. We may view ω as a vector field, and then $\operatorname{tr}_M \omega$ corresponds to the tangential projection of the vector field onto the curve. Thus $\int_M \omega$ is the usual line integral from vector calculus.

6.2.4 • Stokes theorem and integration by parts

The capstone of exterior calculus is the Stokes theorem, a far-ranging generalization of the fundamental theorem of calculus. Under relatively mild smoothness requirements on the compact oriented *n*-dimensional manifold with boundary Ω and the differential form ω it states that

$$\int_{\Omega} d\omega = \int_{\partial \Omega} \operatorname{tr} \omega, \quad \omega \in \Lambda^{n-1}(\Omega).$$

Again, the result is usually stated without writing the trace operator explicitly. The boundary orientation is inherited from that of the manifold (namely, a positively oriented basis for the tangent space of the boundary at some point is one which, when prepended by the outward normal at the point, is positively oriented for Ω).

If we replace ω by $\omega \wedge \mu$ in this equation and use the Leibniz rule (6.12) we obtain an integration-by-parts formula:

$$\int_{\Omega} d\omega \wedge \mu = (-)^{k+1} \int_{\Omega} \omega \wedge d\mu + \int_{\partial \Omega} \operatorname{tr} \omega \wedge \operatorname{tr} \mu, \quad \omega \in \Lambda^k(\Omega), \ \mu \in \Lambda^{n-k-1}(\Omega),$$
(6.14)

valid for sufficiently smooth forms (we shall investigate the required smoothness below). Using the Hodge star operator on the boundary, the boundary integral can be written as the inner product $\langle \star \operatorname{tr} \omega, \operatorname{tr} \mu \rangle_{L^2 \Lambda^{n-k-1}(\partial \Omega)}$.

6.2.5 • An example application

Before continuing, we use the exterior calculus framework to derive a result of vector calculus that plays a large role in the development of mixed finite elements. (See, e.g., [19, Lemma 2.1.6] for a statement of the result in two dimensions.) Let $p: \Omega \to \mathbb{R}$ and $v: \Omega \to \mathbb{R}^3$ be scalar and vector fields on a domain Ω in \mathbb{R}^3 , and consider the inner product $(p, \operatorname{div} v)_{L^2(\Omega)}$. Recognizing the divergence as the exterior derivative of a 2-form, we let $\nu = v^1 dx^2 \wedge dx^3 - v^2 dx^1 \wedge dx^3 + v^3 dx^1 \wedge dx^2$ be the differential 2-form with vector proxy v. Then $d\nu = (\operatorname{div} v)vol$ and $(p, \operatorname{div} v)_{L^2(\Omega)} = \int_{\Omega} p \wedge d\nu$, where the integral on the right-hand side is in the sense of differential forms. Next, we suppose that $\phi: \hat{\Omega} \to \Omega$ is an orientation-preserving diffeomorphism and apply invariance under pullback of the integral (6.13), the exterior product (6.5), and the exterior derivative (6.10) to transform the integral over Ω to one over $\hat{\Omega}$:

$$\int_{\Omega} p \wedge d\nu = \int_{\hat{\Omega}} \phi^*(p \wedge d\nu) = \int_{\hat{\Omega}} \phi^* p \wedge \phi^* d\nu = \int_{\hat{\Omega}} \phi^* p \wedge d\phi^* \nu.$$

From the definition (6.7) for the pullback of a field and the formulas given at the bottom of Table 6.1 for the pullback of a linear map, we see that $\phi^* p = \hat{p} := p \circ \phi$ and the proxy field for $\phi^* \nu$ is the vector field

$$\hat{v} = (\det \phi')(\phi')^{-1}(v \circ \phi) : \hat{\Omega} \to \mathbb{R}^3$$

The vector field \hat{v} is called the *Piola transform* of v. We see that the Piola transform arises naturally as the pullback on 2-forms and that, as a consequence, it satisfies the crucial property

$$(p,\operatorname{div} v)_{L^2(\Omega)} = (\hat{p},\operatorname{div} \hat{v})_{L^2(\hat{\Omega})}.$$

6.2.6 • The L^2 theory of differential forms

In this section we suppose that Ω is a smooth compact *n*-dimensional oriented Riemannian manifold with or without boundary and we view the exterior derivatives $d = d^k$ as closed densely defined unbounded operators from $L^2\Lambda^k(\Omega)$ to $L^2\Lambda^{k+1}(\Omega)$ which fit together to create the L^2 de Rham complex. We show that the exterior derivative operators all have closed range and compute their adjoints. The techniques are completely parallel to those used in Section 3.4 for the special case of a domain in \mathbb{R}^3 .

Since we assume that an orientation and a Riemannian structure are given, the Hodge star is well-defined. From (6.3), the L^2 inner product of the two k-forms is

$$\langle \omega, \mu \rangle_{L^2 \Lambda^k} = \int_{\Omega} \langle \omega, \mu \rangle \mathrm{vol} = \int_{\Omega} \omega \wedge \star \mu, \quad \omega, \mu \in \Lambda^k(\Omega).$$

We may drop the Λ^k from the subscript on the inner product or drop the subscript altogether when there is little risk of confusion.

The *codifferential* operator is defined as $\delta = \pm \star d \star$. The sign is determined by requiring that $\delta = \delta^k : \Lambda^k \to \Lambda^{k-1}$ satisfy

$$\star \delta \omega = (-)^k d \star \omega, \quad \omega \in \Lambda^k.$$

If we set $\mu = \star \nu$ in (6.14) and use this definition of δ , we get a second form of the integration-by-parts formula:

$$\langle d\omega, \nu \rangle = \langle \omega, \delta \nu \rangle + \int_{\partial \Omega} \operatorname{tr} \omega \wedge \operatorname{tr} \star \nu, \quad \omega \in \Lambda^k, \ \nu \in \Lambda^{k+1}.$$

This formula shows that δ is the formal adjoint of d, which explains the choice of sign in its definition.

If $\omega \in L^2 \Lambda^k(\Omega)$ and $\mu \in L^2 \Lambda^{k+1}(\Omega)$, then the equation $d\omega = \mu$ has meaning in the sense of distributions, namely, it means that

$$\langle \omega, \delta \nu \rangle = \langle \mu, \nu \rangle, \quad \nu \in C_0^{\infty} \Lambda^{k+1}(\Omega).$$

We then define $H\Lambda^k(\Omega)$ as the set of $\omega \in L^2\Lambda^k(\Omega)$ for which $d\omega \in L^2\Lambda^{k+1}(\Omega)$, i.e., such a $\mu \in L^2\Lambda^{k+1}(\Omega)$ exists. In other words, we define d as the adjoint of the operator δ from $L^2\Lambda^{k+1}$ to $L^2\Lambda^k$ with the domain of the latter taken as $C_0^{\infty}\Lambda^{k+1}(\Omega)$. Being the adjoint of a densely defined operator, d so defined is a closed operator. It is densely defined in $L^2\Lambda^k$ as well, since clearly $C_0^{\infty}\Lambda^k(\Omega)$ is contained in $H\Lambda^k(\Omega)$ and dense in $L^2\Lambda^k(\Omega)$. Since d is a first order linear differential operator, we have $H^1\Lambda^k(\Omega) \subset H\Lambda^k(\Omega)$ (when k = 0, they are equal, both coinciding with the standard Sobolev space $H^1(\Omega)$). We will also use the fact that $C^{\infty}\Lambda^k(\overline{\Omega})$ is dense in $H\Lambda^k(\Omega)$. This is proved via smoothing just as is done for ordinary Sobolev spaces [41, Section 5.3.3].

The Hilbert spaces $L^2\Lambda^k(\Omega)$ and the closed unbounded operators d^k from $L^2\Lambda^k(\Omega)$ to $L^2\Lambda^{k+1}(\Omega)$ with domains $H\Lambda^k(\Omega)$ form a Hilbert complex, the L^2 de Rham complex on the manifold:

$$0 \to L^2 \Lambda^0(\Omega) \xrightarrow{(d^0, H\Lambda^0)} L^2 \Lambda^1(\Omega) \xrightarrow{(d^1, H\Lambda^1)} \cdots \xrightarrow{(d^{n-1}, H\Lambda^{n-1})} L^2 \Lambda^n(\Omega) \to 0.$$

To obtain something less unwieldy, we may abbreviate this to

$$0 \to L^2 \Lambda^0 \xrightarrow{(d,H\Lambda)} L^2 \Lambda^1 \xrightarrow{(d,H\Lambda)} \cdots \xrightarrow{(d,H\Lambda)} L^2 \Lambda^n \to 0$$
(6.15)

or just give the domain complex

$$0 \to H\Lambda^0 \xrightarrow{d} H\Lambda^1 \xrightarrow{d} \cdots \xrightarrow{d} H\Lambda^n \to 0.$$
(6.16)

We now look to the integration-by-parts formula (6.14) and establish conditions on ω and μ under which it is valid. It certainly holds for smooth ω and μ and therefore, by density, for $\omega \in H^1\Lambda^k$, $\mu \in H^1\Lambda^{n-k-1}$. In order to go further, we now show that the trace operator extends to a bounded operator on $H\Lambda^k(\Omega)$ with values in $H^{-1/2}\Lambda^k(\partial\Omega)$. Once this is accomplished, we can extend (6.14) to the case when one of the forms belongs to $H\Lambda$ and the other to $H^1\Lambda$. (We cannot allow both forms in $H\Lambda$ unless at least one has vanishing trace, since otherwise we cannot make sense of the boundary integral.)

By definition,

$$\langle d\omega, \nu \rangle - \langle \omega, \delta \nu \rangle = 0, \quad \omega \in H\Lambda^k, \ \nu \in C_0^{\infty}\Lambda^{k+1}.$$

Equivalently, after substituting $\mu = \star \nu$,

$$\int d\omega \wedge \mu - (-)^{k+1} \int \omega \wedge d\mu = 0, \quad \omega \in H\Lambda^k, \ \mu \in C_0^{\infty} \Lambda^{n-k-1}$$

For fixed $\omega \in H\Lambda^k$, the left-hand side defines a bounded linear functional on $\mu \in H\Lambda^{n-k-1}$. Defining $\mathring{H}\Lambda^{n-k-1}$ as the closure of $C_0^{\infty}\Lambda^{n-k-1}$ in $H\Lambda^{n-k-1}$, this equation then holds for all $\omega \in H\Lambda^k$, $\mu \in \mathring{H}\Lambda^{n-k-1}$.

To continue with the study of the trace map, we assume that the boundary is Lipschitz in the sense that we can cover it with open sets U of Ω that can be mapped by Lipschitz charts to a half ball $B = \{x \in \mathbb{R}^n \mid |x| = 1, x^n \ge 0\}$ with $U \cap \partial \Omega$ mapping to $B_0 := B \cap \{x^n = 0\}$. The spaces $H^1\Lambda^k(U)$ and $H^{1/2}\Lambda^k(U \cap \partial \Omega)$ may then be defined as the pullbacks of the corresponding spaces on B and B_0 , respectively, and the spaces $H^1\Lambda^k(\Omega)$ and $H^{1/2}\Lambda^k(\partial\Omega)$ by piecing together these spaces with a partition of unity.

Theorem 6.1. The trace map on k-forms maps $H^1\Lambda^k(\Omega)$ onto $H^{1/2}\Lambda^k(\partial\Omega)$.

Proof. In light of the preceding discussion, it suffices to prove this when Ω is the half ball B in \mathbb{R}^n and we are computing the trace on B_0 . This follows from the formula (6.8) for the trace and the surjectivity of the scalar trace $H^1(B) \to H^{1/2}(B_0)$. \Box

Applying the duality Lemma 3.10 as in the proof of Lemma 3.11, we find that the dual of the trace map is a bounded linear injection whose range is the annihilator of the null space of the trace map, which gives the following corollary.

Corollary 6.2. If $L : H^1\Lambda^k(\Omega) \to \mathbb{R}$ is a bounded linear functional for which $L\omega = 0$ whenever tr $\omega = 0$, then there exists a unique $g \in H^{-1/2}\Lambda^k(\partial\Omega)$ such that

$$L\omega = \langle g, \operatorname{tr} \omega \rangle, \quad \omega \in H^1 \Lambda^k(\Omega).$$

Moreover, $\|g\|_{H^{-1/2}\Lambda^k(\partial\Omega)} \leq c \|L\|_{H^1\Lambda^k(\Omega)'}$.

$$L\omega = \int_{\partial\Omega} \operatorname{tr} \omega \wedge h, \quad \omega \in H^1 \Lambda^k(\Omega).$$

For the last statement we have just taken $h = \star g$, where the Hodge star is on the boundary (and the integral is meaningful by dense extension).

Now, if $\mu \in H\Lambda^{n-k-1}$ and $\omega \in H^1\Lambda^k$ satisfies tr $\omega = 0$, then we have

$$\int_{\Omega} d\omega \wedge \mu - (-)^k \int_{\Omega} \omega \wedge d\mu = 0.$$

(It is certainly so if $\mu \in H^1 \Lambda^{n-k-1}$ and extends to $\mu \in H \Lambda$ by density.) Therefore, for any such μ , we may apply the corollary to the functional

$$L_{\mu}\omega := \int_{\Omega} d\omega \wedge \mu - (-)^k \int_{\Omega} \omega \wedge d\mu$$

and conclude that there exists a unique $\gamma \mu \in H^{-1/2} \Lambda^{n-k-1}(\partial \Omega)$ such that

$$\int_{\Omega} d\omega \wedge \mu - (-)^k \int_{\Omega} \omega \wedge d\mu = \int_{\partial \Omega} \operatorname{tr} \omega \wedge \gamma \mu, \quad \omega \in H^1 \Lambda^k.$$

Clearly $\gamma \mu = \operatorname{tr} \mu$ if μ is smooth, say, in $H^1 \Lambda^{n-k-1}$. In this way we have extended the trace operator on $H^1 \Lambda^{n-k-1}$ to a bounded operator mapping

$$H\Lambda^{n-k-1}(\Omega) \to H^{-1/2}\Lambda^{n-k-1}(\partial\Omega)$$

Henceforth we shall denote the extended operator by tr rather than γ . We summarize in the following theorem.

Theorem 6.3. The trace map on $H^1\Lambda^k(\Omega)$ extends to a bounded linear operator tr : $H\Lambda^k(\Omega) \to H^{-1/2}\Lambda^k(\partial\Omega)$. Moreover, we have an extended version of the integration-by-parts formula:

$$\int_{\Omega} d\omega \wedge \mu = (-)^k \int_{\Omega} \omega \wedge d\mu + \int_{\partial \Omega} \operatorname{tr} \mu \wedge \operatorname{tr} \omega, \quad \omega \in H^1 \Lambda^k, \ \mu \in H \Lambda^{n-k-1}.$$
(6.17)

We close this chapter with a few more useful results.

Theorem 6.4. The space $\mathring{H}\Lambda^k(\Omega)$, defined as the closure of $C_0^{\infty}\Lambda^k(\Omega)$ in $H\Lambda^k$, consists of all $\omega \in H\Lambda^k(\Omega)$ such that tr $\omega = 0$.

Proof. The proof proceeds via a partition of unity and flattening the boundary, together with a direct construction in the case of traces from a half space, just as for ordinary Sobolev spaces. See [41, Theorem 5.5.2]. \Box

The operators $(d^k, \mathring{H}\Lambda^k)$, i.e., the operator d^k taken with this smaller space as domain, leads to the de Rham complex with boundary conditions, another Hilbert complex:

$$0 \to L^2 \Lambda^0 \xrightarrow{(d, \dot{H}\Lambda)} L^2 \Lambda^1 \xrightarrow{(d, \dot{H}\Lambda)} \cdots \xrightarrow{(d, \dot{H}\Lambda)} L^2 \Lambda^n \to 0.$$
 (6.18)

The L^2 theory of the operator δ can be obtained analogously to the treatment just given for d. Alternatively, the properties of δ may be derived from those of dvia the Hodge star. We may view δ as a closed densely defined operator on $L^2 \Lambda^k$ with domain

$$H^*\Lambda^k(\Omega) = \{ \omega \in L^2\Lambda^k \, | \, \delta\omega \in L^2\Lambda^{k-1} \, \} = \star H\Lambda^{n-k}(\Omega).$$

By Theorem 6.3, the operator $tr \star acts$ boundedly from

$$H^*\Lambda^k(\Omega) \to H^{-1/2}\Lambda^{n-k}(\partial\Omega).$$

The closure of $C_0^{\infty} \Lambda^k$ in $H^* \Lambda^k$ is

$$\mathring{H}^*\Lambda^k = \{ \, \omega \in H^*\Lambda^k \mid \operatorname{tr} \star \omega = 0 \, \}.$$

A different operator is obtained by taking this space as the domain of δ . The following theorem summarizes the adjoints of these operators.

Theorem 6.5. The adjoint of

- 1. $(d, H\Lambda^k)$ is $(\delta, \mathring{H}^*\Lambda^{k+1})$;
- 2. $(d, \mathring{H}\Lambda^k)$ is $(\delta, H^*\Lambda^{k+1})$;
- 3. $(\delta, H^*\Lambda^{k+1})$ is $(d, \mathring{H}\Lambda^k)$;
- 4. $(\delta, \mathring{H}^* \Lambda^{k+1})$ is $(d, H\Lambda^k)$.

For each of the first two operators, we have corresponding Hilbert cochain complexes, namely, the de Rham cochain complexes (6.15) and (6.18) without and with boundary conditions, respectively. (The last two operators similarly form Hilbert chain complexes that are essentially identical to the corresponding de Rham complexes except for a change of indexing.) As a consequence of the result of Picard, discussed in Section 4.3, each of these complexes satisfies the compactness property, and so all the operators have closed range, and all the complexes are Fredholm. Knowing the adjoints, we can write down the strong form of the Hodge Laplace boundary value problem associated to the two complexes. For each the differential equation is the Hodge Laplace equation

$$(d^{k-1}\delta_k + \delta_{k+1}d^k)u = f$$

or, dropping sub- and superscripts, $(d\delta + \delta d)u = f$. For the complex (6.15) the boundary conditions, implied by the requirements that u and du must belong to the domain d^* , are

$$\operatorname{tr} \star u = 0, \quad \operatorname{tr} \star du = 0 \quad \text{on } \partial \Omega.$$

The interpretation of these boundary conditions in three dimensions was discussed in Section 4.5 and summarized in Table 4.1. For k = 0, the case of the scalar Laplacian, the first boundary condition is vacuous, while the second states that (grad u) $\cdot n = 0$, given *Neumann boundary conditions*. For k = 1, these are magnetic boundary conditions and for k = 2 electric boundary conditions, and for k = 3 they impose Dirichlet boundary conditions.

We can also consider the complex with boundary conditions (6.18). In this case, the boundary conditions are essential, coming from the requirements that u belong to $\mathring{H}\Lambda^k$ and σ to $\mathring{H}\Lambda^{k-1}$:

$$\operatorname{tr} u = 0, \quad \operatorname{tr} \delta u = 0 \quad \text{on } \partial \Omega.$$

Their interpretation in three dimensions is summarized in Table 4.2.

Box 6.1. Summary of Chapter 6 on exterior calculus.

For an *n*-dimensional vector space V, the space $\operatorname{Alt}^k V$ of alternating klinear forms on V has dimension $\binom{n}{k}$. It admits various algebraic operations:

- Wedge product: \wedge : Alt^k $V \times$ Alt^j $V \rightarrow$ Alt^{k+j} V.
- Contraction: \exists : Alt^k $V \times V \to \operatorname{Alt}^{k-1} V.$
- Pullback L^* : Alt^k $W \to$ Alt^k V induced by a linear map $L: V \to W$.
- Inner product on $\operatorname{Alt}^k V$ induced by one on V.
- Volume form, vol, a canonical unit norm element vol spanning Altⁿ V (requires inner product and orientation on V).
- *Hodge star*, a canonical isometry \star : Alt^k $V \to$ Alt^{n-k} V (requires inner product and orientation on V).
- *Proxies*: Alt⁰ V and Altⁿ V are canonically isomorphic to \mathbb{R} , Alt¹ V and Altⁿ⁻¹ V are canonically isomorphic to V.

A differential k-form on a manifold assigns an element of Alt^k of the tangent space to each point. All the operations above carry over to differential forms. In addition, a differential k-form ω admits an exterior derivative $d\omega$, which is a differential (k + 1)-form. The exterior derivative satisfies $d \circ d = 0$ and the Leibniz rule:

$$d(\omega \wedge \mu) = d\omega \wedge \mu + (-)^k \omega \wedge d\mu, \quad \omega \in \Lambda^k(\Omega), \ \mu \in \Lambda^j(\Omega).$$

A differential *n*-form may be integrated over an oriented *n*-dimensional manifold and satisfies the Stokes theorem

$$\int_{\Omega} d\omega = \int_{\partial \Omega} \operatorname{tr} \omega,$$

where the trace operator tr is the pullback of the inclusion $\partial \Omega \subset \Omega$.

On a smooth compact oriented Riemannian manifold with boundary, the integration-by-parts formula (6.17) holds. There we may view the exterior derivative as an unbounded operator $L^2\Lambda^k \rightarrow L^2\Lambda^{k+1}$ with domain $H\Lambda^k$. The L^2 de Rham complex is a closed Hilbert complex with domain complex

$$0 \to H\Lambda^0 \xrightarrow{d} H\Lambda^1 \xrightarrow{d} \cdots \xrightarrow{d} H\Lambda^n \to 0.$$

The adjoint complex has domain complex

$$0 \leftarrow \mathring{H}^* \Lambda^0 \xleftarrow{\delta} \mathring{H}^* \Lambda^1 \xleftarrow{\delta} \cdots \xleftarrow{\delta} \mathring{H}^* \Lambda^n \leftarrow 0,$$

where $\delta = \pm \star d \star$ is the codifferential. Alternatively, we may include boundary conditions in the de Rham complex, whence they do not arise in the adjoint.

Chapter 7 Finite element differential forms

Consider now a Lipschitz polyhedral domain $\Omega \subset \mathbb{R}^n$ and its L^2 de Rham complex (6.15) with domain complex (6.16). For any k, the segment

$$H\Lambda^{k-1}(\Omega) \xrightarrow{d} H\Lambda^{k}(\Omega) \xrightarrow{d} H\Lambda^{k+1}(\Omega)$$
 (7.1)

determines a Hodge Laplace boundary value problem involving the two exterior derivatives and their adjoints, as discussed in Section 4.4. In Theorem 6.5 we have identified the adjoint of $(d, H\Lambda^k)$ as $(\delta, \mathring{H}^*\Lambda^k)$, so the strong form of the *k*-form Hodge Laplace boundary value problem is

$$(\delta d + d\delta)u = f - P_{\mathfrak{H}}^{k} f \text{ in } \Omega, \quad \operatorname{tr} \star u = 0, \ \operatorname{tr} \star du = 0 \text{ on } \partial\Omega, \quad u \perp \mathfrak{H}^{k}, \ (7.2)$$

where \mathfrak{H}^k are the harmonic k-forms on Ω , determined by the differential equations dp = 0, $\delta p = 0$ and the boundary condition $\operatorname{tr} \star p = 0$. In the case of three dimensions, we have interpreted this problem, and related ones, in Section 4.5. In Chapter 5 we derived the requirements needed for a successful discretization of this problem using the mixed formulation and Galerkin's method. For this we need to construct Galerkin subspaces V_h^{k-1} of $H\Lambda^{k-1}(\Omega)$ and V_h^k of $H\Lambda^k(\Omega)$. The key requirements of these subspaces, besides from being efficiently implementable, are that they afford good approximation, form a subcomplex, and admit a bounded cochain projection. In that case, as shown in Chapter 5, the mixed Galerkin method is consistent and stable and we obtain convergence with an optimal rate.

In this chapter, we shall obtain efficiently implementable Galerkin subspaces of $H\Lambda^k$ by constructing them as spaces of finite elements. The desired degree of approximation will then be obtained by using a sufficiently fine mesh and/or sufficiently high degree piecewise polynomial shape functions. The challenge is to ensure that the spaces form a subcomplex admitting a bounded cochain projection.

Specifically, given any simplicial triangulation \mathcal{T}_h of Ω (in any number of dimensions $n \geq 1$), any form degree $0 \leq k \leq n$, and any polynomial degree $r \geq 1$, we shall construct two finite element subspaces of $H\Lambda^k(\Omega)$, denoted

$$\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$$
 and $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$

The former space uses as shape functions the *complete* space of polynomial k-forms with degree at most r, while the latter uses a somewhat smaller space of polynomials which, following [29], we refer to as *trimmed*. Then we show several ways to select two such spaces, one for (k - 1)-forms, the other for k-forms, to solve the Hodge Laplacian. These complete and trimmed families of finite element spaces of differential forms can be viewed as the most canonical for discretization of the de Rham complex. This is one of the important conclusions of FEEC.

The chapter is organized as follows. Before defining the complete and trimmed spaces $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$ and $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$ in Section 7.4, we establish some properties of differential forms with polynomial coefficients in Sections 7.1–7.3. Then, in Section 7.5, we show how these spaces may be combined into complexes which fulfill the hypotheses required for the error analysis in Chapter 5, thus establishing a full set of error bounds for our methods. In Section 7.7 we briefly describe analogous results for finite elements based on meshes of boxes rather than simplices. Finally, in Section 7.8, we relate the elements described in this chapter to various finite elements in the literature.

Before proceeding, we make a brief remark on notation. The finite element spaces $\mathcal{P}_r\Lambda^k(\mathcal{T}_h)$ and $\mathcal{P}_r^-\Lambda^k(\mathcal{T}_h)$ which we will define in the next section consist of piecewise polynomial differential k-forms defined over the domain Ω . We will give a unified construction of the spaces, but each space depends on the polynomial degree $r \geq 0$, the form degree k between 0 and $n = \dim \Omega$, and the simplicial mesh \mathcal{T}_h . We shall also use the notation $\mathcal{P}_r\Lambda^k$ and $\mathcal{P}_r^-\Lambda^k$ for the spaces of complete and trimmed polynomial (not piecewise polynomial) differential kforms of degree r that are used on each simplex to construct the finite element spaces. These spaces of polynomials are defined for each n and for each $r \geq 0$ and $0 \leq k \leq n$. Finally, we write $\mathcal{P}_r\Lambda^k(T)$ and $\mathcal{P}_r^-\Lambda^k(T)$ for the spaces of restrictions of these polynomials to a single simplex T. Thus our notation is analogous to a common notation for standard Lagrange finite elements of degree r; $\mathcal{P}_r(\mathcal{T}_h)$ for the finite element space, \mathcal{P}_r for the space of polynomials of degree r, and $\mathcal{P}_r(T)$ for their restriction to an element.

7.1 • The complete polynomial spaces of differential forms

The spaces $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$ and $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$ will be constructed through the usual process of *finite element assembly*. In general, a finite element space is determined by the specification of the triangulation \mathcal{T}_h , and for each element $T \in \mathcal{T}_h$:

- a finite dimensional space of shape functions V(T) on T, and
- a set of degrees of freedom (DOFs) for V(T) (i.e., a basis for its dual space), with each DOF associated to a particular face of T.

The assembled finite element space is then defined to consist of all functions on Ω whose restriction to each T belongs to the shape function space V(T) and for which the DOFs are single-valued (in the sense that when two elements share a common face, the corresponding DOFs on the face take the same value).

In the present case, the shape function spaces will be spaces of differential kforms with polynomial coefficients. For $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$, the choice of shape function space V(T) is an obvious one, namely, the complete polynomial space

$$\mathcal{P}_r \Lambda^k(T) = \left\{ \sum_{1 \le \sigma_1 < \dots < \sigma_k \le n} p_\sigma \, dx^\sigma \, | \, p_\sigma \in \mathcal{P}_r(T) \right\}$$

of a differential k-form with polynomial coefficients of degree at most r. Its dimension is easily calculated:

$$\dim \mathcal{P}_r \Lambda^k(T) = \binom{n}{k} \times \dim \mathcal{P}_r(T) = \binom{n}{k} \binom{n+r}{n} = \binom{r+n}{r+k} \binom{r+k}{r}.$$
(7.3)

Note that $d\mathcal{P}_r\Lambda^k \subset \mathcal{P}_{r-1}\Lambda^{k+1}$, i.e., the exterior derivative lowers the polynomial degree at the same time as it raises the form degree. Therefore, for each r we have a polynomial subcomplex of the de Rham complex:

$$0 \to \mathcal{P}_r \Lambda^0 \xrightarrow{d} \mathcal{P}_{r-1} \Lambda^1 \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{P}_{r-n} \Lambda^n \to 0.$$

(We understand $\mathcal{P}_s\Lambda^k$ to be 0 if s < 0.) Moreover, as proved below in Corollary 7.3, this complex is exact, except for the constant functions in $\mathcal{P}_r\Lambda^0$, which form the kernel of the first d. That is, if $\omega \in \mathcal{P}_s\Lambda^k$ for some k > 0 and $d\omega = 0$, then $\omega = d\mu$ for some $\mu \in \mathcal{P}_{s+1}\Lambda^{k-1}$. The proof of this fact will involve a very valuable tool called the *Koszul complex*. The same tool will be used to define DOFs and to define an alternative space of shape functions.

7.2 - The Koszul complex

For $x \in \Omega \subset \mathbb{R}^n$, the tangent space $T_x\Omega$ may be identified with \mathbb{R}^n , so we obtain a vector field by assigning to each point $x \in \Omega$ that same x, now viewed as a vector in its tangent space. This identity vector field is display in Figure 7.1.

The Koszul differential $\kappa \omega = \omega \lrcorner id$ of a k-form ω is the (k-1)-form obtained by contracting ω with the identity vector field. That is,

$$(\kappa \,\omega)_x(v_1,\ldots,v_{k-1}) = \omega_x(x,v_1,\ldots,v_{k-1}), \quad x \in \Omega, \ v_1,\ldots,v_{k-1} \in \mathbb{R}^n,$$

where inside the parentheses on the right-hand side the point x is viewed as a tangent vector. On 0-forms, κ is understood to be zero. Since $\omega(x, x, ...) \equiv 0$, $\kappa \circ \kappa = 0$, so κ is indeed a differential. Note that

$$\kappa(dx^i) = x^i$$

i.e., the Koszul differential of the basic 1-form dx^i is the *i*th coordinate function. Since κ is defined by contraction with a vector field, it satisfies a Leibniz rule:

$$\kappa(\omega \wedge \mu) = (\kappa \, \omega) \wedge \mu + (-1)^k \omega \wedge (\kappa \, \mu).$$

In particular, $\kappa(f\omega) = f \kappa \omega$ for any differential form ω and any function f, so, for example, $\kappa(f dx^i) = f x^i$. These rules determine κ completely. For example,

$$\begin{aligned} \kappa(dx^i \wedge dx^j) &= x^i \, dx^j - x^j \, dx^i, \\ \kappa(dx^i \wedge dx^j \wedge dx^k) &= x^i \, dx^j \wedge dx^k - x^j \, dx^i \wedge dx^k + x^k \, dx^i \wedge dx^j. \end{aligned}$$



Figure 7.1. The identity vector field.

If we identify 1-forms with vector fields, then κ corresponds to the dot product of the vector field with x (or, more properly, with the identity vector field). Applied to 2-forms in three dimensions, κ is the cross product with x, and on 3-forms it is the product of a scalar field with x to get a vector field.

The Koszul differential κ maps $\mathcal{P}_r \Lambda^k$ to $P_{r+1} \Lambda^{k-1}$, increasing the polynomial degree and decreasing the form degree, exactly the reverse of d. Thus both κd and $d\kappa$ map $\mathcal{P}_r \Lambda^k$ to itself. The following theorem points to an intimate relation between κ and d, called the *homotopy formula*. In it, we write $\mathcal{H}_r \Lambda^k$ for the k-forms with *homogeneous* polynomial coefficients of degree r.

Theorem 7.1.

$$(d \kappa + \kappa d)\omega = (k+r)\omega, \quad \omega \in \mathcal{H}_r \Lambda^k.$$

Proof. We prove this by induction on k, the case k = 0 being Euler's identity $x \cdot \operatorname{grad} p = r p$ for p a homogeneous polynomial of degree r. Thus it is enough to prove the result for the (k+1)-form $dx^i \wedge \omega$ assuming that it is true for $\omega \in \mathcal{H}_r \Lambda^k$. Using the Leibniz rules for κ and d and the identity $d^2 = 0$, we have

$$d\kappa(dx^i \wedge \omega) = d(x^i\omega - dx^i \wedge \kappa\,\omega) = dx^i \wedge \omega + x^i d\omega + dx^i \wedge d\,\kappa\,\omega$$

and

$$\kappa \, d(dx^i \wedge \omega) = - \, \kappa (dx^i \wedge d\omega) = - x^i d\omega + dx^i \wedge \kappa \, d\omega$$

Therefore

$$\begin{aligned} (d\kappa + \kappa d)(dx^i \wedge \omega) &= dx^i \wedge \omega + dx^i \wedge (d\kappa + \kappa d)\omega \\ &= dx^i \wedge \omega + (k+r)dx^i \wedge \omega = (k+r+1)dx^i \wedge \omega, \end{aligned}$$

where we have invoked the inductive hypothesis to obtain the second equality. This completes the proof. \Box

Remark 7.2. An alternative proof is based on Cartan's *homotopy formula* of differential geometry. See Theorem 3.1 of [11].

Corollary 7.3. The polynomial de Rham complex

$$\mathcal{P}_r \Lambda^0 \xrightarrow{d} \mathcal{P}_{r-1} \Lambda^1 \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{P}_{r-n} \Lambda^n \to 0$$

and the Koszul complex

$$0 \to \mathcal{P}_{r-n}\Lambda^n \xrightarrow{\kappa} \mathcal{P}_{r-n+1}\Lambda^{n-1} \xrightarrow{\kappa} \cdots \xrightarrow{\kappa} \mathcal{P}_r\Lambda^0$$

are both exact.

Note that we have left the 0 off from one end of each complex, since we would not have exactness there.

Proof. We establish this for the homogeneous polynomial de Rham complex

$$\mathcal{H}_r \Lambda^0 \xrightarrow{d} \mathcal{H}_{r-1} \Lambda^1 \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{H}_{r-n} \Lambda^n \to 0.$$

We can then just sum to get the result. We must show that if $\omega \in \mathcal{H}_s \Lambda^k$ and $d\omega = 0$, then ω is in the range of d. Indeed,

$$\omega = (s+k)^{-1} (d\kappa + \kappa d) \omega = (s+k)^{-1} d\kappa \omega.$$

A similar proof holds for the Koszul complex.

Another important consequence is a direct sum decomposition.

Corollary 7.4.

$$\mathcal{H}_r \Lambda^k = \kappa \,\mathcal{H}_{r-1} \Lambda^{k+1} \oplus d\mathcal{H}_{r+1} \Lambda^{k-1}. \tag{7.4}$$

Π

Proof. By the homotopy formula, any element of $\mathcal{H}_r \Lambda^k$ belongs to $\kappa \mathcal{H}_{r-1} \Lambda^{k+1} + d\mathcal{H}_{r+1} \Lambda^{k-1}$. Moreover the intersection of these two spaces is zero, since if ω belongs to the intersection, then $d\omega = 0$, $\kappa \omega = 0$, so $\omega = 0$ by the homotopy formula.

7.3 • The trimmed polynomial spaces of differential forms

In view of the results just established, the complete space of polynomial k-forms of degree r can be decomposed as

$$\mathcal{P}_{r}\Lambda^{k} = \mathcal{P}_{r-1}\Lambda^{k} \oplus \mathcal{H}_{r}\Lambda^{k} = \mathcal{P}_{r-1}\Lambda^{k} \oplus \kappa \mathcal{H}_{r-1}\Lambda^{k+1} \oplus d\mathcal{H}_{r+1}\Lambda^{k-1}.$$

If we drop the last summand, we get a space intermediate between $\mathcal{P}_{r-1}\Lambda^k$ and $\mathcal{P}_r\Lambda^k$, the trimmed spaces of polynomial forms of degree r:

$$\mathcal{P}_r^-\Lambda^k := \mathcal{P}_{r-1}\Lambda^k + \kappa \,\mathcal{H}_{r-1}\Lambda^{k+1}$$

Note that $\mathcal{P}_r^-\Lambda^0 = \mathcal{P}_r\Lambda^0$ and $\mathcal{P}_r^-\Lambda^n = \mathcal{P}_{r-1}\Lambda^n$, but for 0 < k < n, $\mathcal{P}_r^-\Lambda^k$ is contained strictly between $\mathcal{P}_{r-1}\Lambda^k$ and $\mathcal{P}_r\Lambda^k$. One can compute the dimension

of the dropped term $d\mathcal{H}_{r+1}\Lambda^{k-1}$ using the exactness of the Koszul complex and induction (see [11, Theorem 3.3]). This gives a formula for the dimension of the trimmed polynomial spaces:

$$\dim \mathcal{P}_r^- \Lambda^k = \binom{r+n}{r+k} \binom{r+k-1}{k}.$$

Comparing with (7.3), we see that

$$\dim \mathcal{P}_r^- \Lambda^k = \frac{r}{r+k} \dim \mathcal{P}_r \Lambda^k.$$

We have $d\mathcal{P}_r^-\Lambda^k \subset d\mathcal{P}_r\Lambda^k \subset \mathcal{P}_{r-1}\Lambda^{k+1} \subset \mathcal{P}_r^-\Lambda^k$. So we obtain another polynomial de Rham complex in which the spaces are trimmed:

$$\mathcal{P}_r^- \Lambda^0 \xrightarrow{d} \mathcal{P}_r^- \Lambda^1 \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{P}_r^- \Lambda^n \to 0.$$
(7.5)

Note that in this sequence, unlike in the polynomial de Rham complex of Corollary 7.3, the degree r stays constant. This complex is also exact. To see this, first note that if $d \kappa \omega = 0$ for some differential form ω , then, by the homotopy formula applied to $\kappa \omega$, we have $\kappa \omega = 0$. In short, d is injective on the range of κ . It follows that

$$\mathcal{N}(d|_{\mathcal{P}_{r-1}\Lambda^k}) = \mathcal{N}(d|_{\mathcal{P}_r^-\Lambda^k}).$$
(7.6)

Next, since $\mathcal{P}_r \Lambda^{k-1} = \mathcal{P}_r^- \Lambda^{k-1} + d\mathcal{H}_{r+1} \Lambda^{k-2}$, we obtain

$$d\mathcal{P}_r \Lambda^{k-1} = d\mathcal{P}_r^- \Lambda^{k-1}. \tag{7.7}$$

By the exactness of the full polynomial de Rham complex, the left-hand sides of (7.6) and (7.7) are equal. Therefore, the right-hand sides are equal as well:

$$\mathcal{N}(d|_{\mathcal{P}^-_r\Lambda^k}) = d\mathcal{P}^-_r\Lambda^{k-1}.$$

This establishes the exactness of (7.5).

Remark 7.5. The Koszul differential was defined as the contraction of a differential form at each point x with the vector in $T_x \mathbb{R}^n$ obtained by translating the vector in $\mathbb{R}^n \equiv T_0 \mathbb{R}^n$ which points from the origin to x. Thus the Koszul differential depends on the choice of origin in \mathbb{R}^n . However, one can check [13, p. 331] that the polynomial space $\mathcal{P}_r^- \Lambda^k$ is unaffected by this choice. This observation is significant, because it allows us to define the trimmed space $\mathcal{P}_r^- \Lambda^k$ not only on \mathbb{R}^n but on any affine subspace. It is then easy to see that if $\omega \in \mathcal{P}_r^- \Lambda^k(\mathbb{R}^n)$, then the trace of ω on an affine subspace V belongs to $\mathcal{P}_r^- \Lambda^k(V)$.

7.4 • Finite element differential forms

We are now ready to define our finite element spaces. As shape functions we use one of the complete polynomial spaces $\mathcal{P}_r \Lambda^k(T)$ or one of the trimmed polynomial spaces $\mathcal{P}_r^- \Lambda^k(T)$. An important observation is that the two spaces are related and are best understood if treated together. Indeed, the DOFs for the complete space involves the trimmed spaces on the faces and vice versa. Specifically, the DOFs for $\omega \in \mathcal{P}_r \Lambda^k(T)$ are

$$\omega \mapsto \int_{f} (\operatorname{tr}_{f} \omega) \wedge \mu, \quad \mu \in \mathcal{P}_{r+k-d}^{-} \Lambda^{d-k}(f), \quad f \in \Delta_{d}(T), \ d \ge k, \quad (7.8)$$

while those for $\omega \in \mathcal{P}_r^- \Lambda^k(T)$ are

$$\omega \mapsto \int_{f} (\operatorname{tr}_{f} \omega) \wedge \mu, \quad \mu \in \mathcal{P}_{r+k-d-1} \Lambda^{d-k}(f), \quad f \in \Delta_{d}(T), \ d \ge k.$$
(7.9)

Note that each DOF is associated to a specific face f, of some dimension between 0 and n, of the simplex $T(\Delta_d(T)$ denotes the set of all faces of T of dimension d). The assembled finite element space is then determined by the requirement that the DOFs are single-valued.

In order that the DOF specifications (7.8) and (7.9) well define finite element spaces, we must verify *unisolvence* for each of them. This requires showing, first, that the total number of DOFs we have associated to the shape function space is no more than its dimension and, second, that if all the DOFs vanish for some ω in the shape function space, then ω itself vanishes. This can be accomplished for both choices of shape function spaces. Here we verify unisolvence for the trimmed space $\mathcal{P}_r^- \Lambda^k(T)$, for any polynomial degree $r \ge 1$, any dimension $n \ge 1$, and any form degree $0 \le k \le n$. The proof, which is based on the properties of the Koszul differential, is taken directly from [6]. A very similar proof applies in the case of the complete polynomial spaces.

The number of DOFs defined in (7.9) is

$$\sum_{d \ge k} \#\Delta_d(T) \dim[\mathcal{P}_{r+k-d-1}\Lambda^{d-k}(\mathbb{R}^d)] = \sum_{d \ge k} \binom{n+1}{d+1} \binom{r+k-1}{d} \binom{d}{d-k} = \sum_{j \ge 0} \binom{n+1}{j+k+1} \binom{r+k-1}{j+k} \binom{j+k}{j}.$$

Simplifying with the binomial identity

$$\binom{a}{b}\binom{b}{c} = \binom{a}{c}\binom{a-c}{a-b}$$

(which is immediate from the definition of the binomial coefficients), and Vandermonde's identity

$$\sum_{j\geq 0} \binom{a}{b+j} \binom{c}{j} = \binom{a+c}{a-b},$$

the right-hand side becomes

$$\binom{r+n}{r+k}\binom{r+k-1}{k} = \dim \mathcal{P}_r^- \Lambda^k(T),$$

as required.

It remains to show that if all the DOFs vanish for some $\omega \in \mathcal{P}_r^- \Lambda^k(T)$, then ω vanishes. The proof proceeds by induction on the dimension n (the case n = 1 being easy). For any facet (face of codimension 1) f of T, $\operatorname{tr}_f \omega$ belongs to $\mathcal{P}_r^- \Lambda^k(f)$, and the DOFs for $\operatorname{tr}_f \omega$ vanish, since they are among the DOFs for ω . Therefore we can invoke the inductive hypothesis to conclude that $\operatorname{tr}_f \omega$ vanishes for all facets f, i.e., that $\omega \in \mathring{\mathcal{P}}_r^- \Lambda^k(T)$, the subspace of elements of $\mathcal{P}_r^- \Lambda^k(T)$ whose traces vanish on the entire boundary. Since the trace operator commutes with the exterior derivative (as does any pullback), we have that $\nu := d\omega \in \mathring{\mathcal{P}}_{r-1}^- \Lambda^{k+1}(T)$. Moreover,

$$\int_{T} \nu \wedge p = \pm \int_{T} \omega \wedge dp = 0, \quad p \in \mathcal{P}_{r+k-n} \Lambda^{n-k-1}(T),$$

where the first equality comes from integration by parts and the second from the vanishing DOFs of ω . From this one can show, by a concrete construction based on barycentric coordinates [6, Lemma 3.4], that $\nu = 0$. Since $\omega \in \mathcal{P}_r^- \Lambda^k$ and $d\omega = 0$, it follows that $\omega \in \mathcal{P}_{r-1}\Lambda^k$, which, together with the vanishing of its traces, means $\omega \in \mathring{\mathcal{P}}_{r-1}\Lambda^k$. A second application of the construction of [6, Lemma 3.4] shows that ω vanishes and completes the unisolvence proof.

Thus we have defined a family of spaces of polynomial differential k-forms on each element T for use as shape functions, namely, the spaces $\mathcal{P}_r \Lambda^k(T)$ and $\mathcal{P}_r^- \Lambda^k(T)$, defined for each $r \geq 1$. Moreover, for each of these spaces we have specified a unisolvent set of DOFs. Together with the triangulation \mathcal{T}_h , each such choice then determines a corresponding finite element space $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$ or $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$ of piecewise polynomial differential forms.

The DOFs defined in (7.8) and (7.9) have many nice properties, which are reflected in properties of the assembled finite element spaces. In particular,

 the DOFs enforce exactly the continuity needed for the assembled finite element function to belong to HΛ^k,

$$\mathcal{P}_r\Lambda^k(\mathcal{T}_h) = \{ \omega \in H\Lambda^k(\Omega) \, | \, \omega|_T \in \mathcal{P}_r\Lambda^k(T) \, \};$$

• the canonical projection operators $\Pi_{r,h}^k : C\Lambda^k(\bar{\Omega}) \to \mathcal{P}_r\Lambda^k(\mathcal{T}_h)$ determined by the DOFs form a cochain map and similarly for the \mathcal{P}^- spaces with their canonical projections $\Pi_{r,h}^{-,k}$.

The commutativity of the canonical projections with the exterior derivative can be verified using the Stokes theorem. It is a crucial property: the main step toward the construction of bounded cochain projections.

7.5 • Properties of the finite element spaces

The approximation properties of the finite element spaces $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$ and $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$ (asymptotically, as the mesh size *h* tends to 0) can be established as is usual for finite elements, using the fact that the canonical projection preserves polynomials locally, and applying the Bramble–Hilbert lemma and scaling. This is carried out in detail in [11, Theorem 5.3], where the following result is proved.

Theorem 7.6. Denote by Π_h the canonical projection of $\Lambda^k(\Omega)$ onto either $\mathcal{P}_r\Lambda^k(\mathcal{T}_h)$ or $\mathcal{P}_{r+1}^-\Lambda^k(\mathcal{T}_h)$. Let $1 \le p \le \infty$ and $(n-k)/p < s \le r+1$. Then Π_h extends boundedly to $W_p^s\Lambda^k(\Omega)$, and there exists a constant C independent of h, such that

$$\|\omega - \Pi_h \omega\|_{L^p \Lambda^k(\Omega)} \le Ch^s |\omega|_{W^s_p \Lambda^k(\Omega)}, \quad \omega \in W^s_p \Lambda^k(\Omega).$$

Note that the restriction s > (n - k)/p comes from the Sobolev embedding theorem: it is what is required for functions in W_p^s to having traces on k-dimensional faces.

For a smooth k-form ω and the complete polynomial space $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$, we may take p = 2, s = r + 1 to obtain an L^2 estimate of order r + 1:

$$\|\omega - \Pi_{r,h}^k \omega\|_{L^2 \Lambda^k} \le C h^{r+1} |\omega|_{H^{r+1} \Lambda^k}.$$

For the trimmed polynomial space $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$, which contains the full polynomial space $\mathcal{P}_{r-1}\Lambda^k(\mathcal{T}_h)$ but not $\mathcal{P}_r\Lambda^k(\mathcal{T}_h)$, the order is reduced to r:

$$\|\omega - \Pi_{r,h}^{-,k}\omega\|_{L^2\Lambda^k} \le Ch^r |\omega|_{H^r\Lambda^k}.$$

For the complete polynomial space, we have the cochain relation

$$d\Pi_{r,h}^k \omega = \Pi_{r-1,h}^{k+1} d\omega,$$

and so

$$\|d(\omega - \Pi_{r,h}^{k}\omega)\|_{L^{2}\Lambda^{k+1}} = \|d\omega - \Pi_{r-1,h}^{k+1}d\omega\|_{L^{2}\Lambda^{k+1}} \le Ch^{r}|d\omega|_{H^{r}\Lambda^{k+1}},$$

a loss of one order of convergence for $d\omega$ as compared to ω . For the trimmed space the cochain relation is

$$d\Pi_{r,h}^{-,k}\omega = \Pi_{r,h}^{-,k+1}d\omega$$

so there is no loss of order:

$$\|d(\omega - \Pi_{r,h}^{-,k}\omega)\|_{L^2\Lambda^{k+1}} \le Ch^r |d\omega|_{H^r\Lambda^{k+1}}$$

We are now ready to choose the pair of Galerkin subspaces V_h^{k-1} , V_h^k . Two obvious choices are

$$\mathcal{P}_r\Lambda^{k-1}(\mathcal{T}_h), \ \mathcal{P}_{r-1}\Lambda^k(\mathcal{T}_h) \quad ext{and} \quad \mathcal{P}_r^-\Lambda^{k-1}(\mathcal{T}_h), \ \mathcal{P}_r^-\Lambda^k(\mathcal{T}_h).$$

using either complete polynomial spaces for both subspaces or trimmed spaces for both. There are two other choices as well, which use one complete space and one trimmed. We may reduce V_h^{k-1} in the complete polynomial pair to obtain

$$\mathcal{P}_r^- \Lambda^{k-1}(\mathcal{T}_h), \ \mathcal{P}_{r-1}\Lambda^k(\mathcal{T}_h).$$

This does not change the range dV_h^{k-1} and again the canonical projections commute: $d\Pi_{r,h}^{-,k} = \Pi_{r-1,h}^k d$. Similarly, we may expand V_h^{k-1} in the trimmed polynomial pair to get

$$\mathcal{P}_r \Lambda^{k-1}(\mathcal{T}_h), \ \mathcal{P}_r^- \Lambda^k(\mathcal{T}_h),$$

which, again, does not affect the range of d. Altogether, we obtain four possible choices for each degree r:

$$V_{h}^{k-1} = \begin{cases} \mathcal{P}_{r}\Lambda^{k-1}(\mathcal{T}_{h}) \\ \text{or} \\ \mathcal{P}_{r}^{-}\Lambda^{k-1}(\mathcal{T}_{h}) \end{cases}, \quad V_{h}^{k} = \begin{cases} \mathcal{P}_{r}^{-}\Lambda^{k}(\mathcal{T}_{h}) \\ \text{or} \\ \mathcal{P}_{r-1}\Lambda^{k}(\mathcal{T}_{h}) \text{ (if } r > 1) \end{cases}.$$
(7.10)

The L^2 rates of approximation they achieve for the variables $u \in \Lambda^k$ and $\sigma = \delta u \in \Lambda^{k-1}$ are shown in Table 7.1. That these estimates hold not only for the projection of σ and u but also for the Galerkin solution (when the true solution is smooth) can be deduced from the improved error estimates of Theorem 5.6.

Of course, to obtain these results we need to fulfill the hypotheses of Theorem 5.6. For this we have to specify a third space $V_h^{k+1} \subset H\Lambda^{k+1}$. (This space is not used in the Galerkin method but is required for theory.) For this we take

$$V_h^{k+1} = \begin{cases} \mathcal{P}_r^- \Lambda^{k+1}(\mathcal{T}_h) & \text{if } V_h^k = \mathcal{P}_r^- \Lambda^k(\mathcal{T}_h), \\ \mathcal{P}_{r-1}^- \Lambda^{k+1}(\mathcal{T}_h) & \text{if } V_h^k = \mathcal{P}_{r-1} \Lambda^k(\mathcal{T}_h). \end{cases}$$

V_h^{k-1}	V_h^k	σ	$d\sigma$	u	du
$\mathcal{P}_r \Lambda^{k-1}$	$\mathcal{P}_r^- \Lambda^k$	r+1	r	r	r
$\mathcal{P}_r^- \Lambda^{k-1}$	$\mathcal{P}_r^- \Lambda^k$	r	r	r	r
$\mathcal{P}_r \Lambda^{k-1}$	$\mathcal{P}_{r-1}\Lambda^k$	r+1	r	r	r-1
$\mathcal{P}_r^- \Lambda^{k-1}$	$\mathcal{P}_{r-1}\Lambda^k$	r	r	r	r-1

 Table 7.1. Rates of approximation for stable choices of Galerkin subspaces.

This defines the subcomplex (5.5) of the corresponding segment of the de Rham complex, required by Theorem 5.6. The remaining hypothesis is that there exist W-bounded cochain projections, i.e., bounded projection operators $\pi_h^j: L^2\Lambda^j \to$ V_{k}^{j} for j = k - 1, k, k + 1 which commute with d on its domain. We have already remarked that the canonical projections, defined through the DOFs, form a cochain projection, that is, they commute with d on smooth enough forms. However, the canonical projection is not in general bounded on $L^2\Lambda^k$, nor even on $H\Lambda^k$. For example, $\Pi^0_{1,h}$ is the usual interpolant into the space of piecewise linear functions. It is not bounded on $H\Lambda^0 = H^1$ in more than one dimension, because it depends on the point values at the vertices, and it is not L^2 bounded even in one dimension. However, there are several ways a bounded cochain projection can be constructed. One approach, which is presented in detail in [11, Section 5.4], is to apply a mollifier before applying the canonical projection (following [28]), and then to correct the result to obtain a projection with the desired properties (following [66]). This construction provides the desired W-bounded cochain projections, and hence the error estimates of Theorems 5.5 and 5.6 hold.

We remark that there are various other properties one might desire from the cochain projections which are not provided by the construction just recalled. For example, one might want the construction to be local or one might ask that the bounds be uniform in the degree r. For this reason, the construction of bounded cochain projections remains an active research area [30, 43].

7.6 • The Whitney forms $\mathcal{P}_1^- \Lambda^k$ and de Rham's theorem

The lowest degree trimmed polynomial finite element space $\mathcal{P}_1^-\Lambda^k(\mathcal{T}_h)$ is a subspace of $H\Lambda^k(\Omega)$. Its elements are called the *Whitney k-forms*. Its shape functions $\mathcal{P}_1^-\Lambda^k = \mathcal{P}_0\Lambda^k + \kappa \mathcal{P}_0\Lambda^{k+1}$ form a subspace of the complete space $\mathcal{P}_1\Lambda^k$ of linear polynomial k-forms (a proper subspace if k > 0). Figure 7.2 shows element diagrams representing the four spaces $\mathcal{P}_1^-\Lambda^k$, $0 \le k \le 3$, in three dimensions. Each element diagram shows a single simplex with a symbol placed on a face to mark each DOF occuring there. Notice that the DOFs of the Whitney forms are highly geometric: for the Whitney k-forms, there is precisely one DOF on each k-dimensional subsimplex—namely, the integral of the trace over the faces—and none on the faces of other dimensions. In fact, these spaces appeared in the book [70] of Hassler Whitney in 1957, independently of finite elements.

As we have seen, the Whitney forms assemble to form a cochain complex:

$$0 \to \mathcal{P}_1^- \Lambda^0(\mathcal{T}_h) \xrightarrow{d} \mathcal{P}_1^- \Lambda^1(\mathcal{T}_h) \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{P}_1^- \Lambda^1(\mathcal{T}_h) \to 0,$$
(7.11)

which is a discrete subcomplex of the L^2 de Rham complex associated to the triangulation \mathcal{T}_h . We also recall that the Whitney forms in $\mathcal{P}_1^-\Lambda^k(\mathcal{T}_h)$ have



Figure 7.2. DOFs for the Whitney forms in three dimensions.

well-defined traces on the k-faces. Indeed, that the traces on k-faces are singlevalued is part of the definition of the space.⁴ Therefore, the de Rham map F^k , which takes differential k-forms to k-cochains, is well-defined on the Whitney forms:

$$F^k \phi(g) = \int_g \operatorname{tr}_g \phi, \quad \phi \in \mathcal{P}_1^- \Lambda^k(\mathcal{T}_h), \quad g \in \Delta_k(\mathcal{T}_h).$$

Notice that the values $F^k \phi(g)$ are just the DOFs of ϕ : they can be assigned arbitrarily and uniquely determine an element of ϕ . This is just another way to say that the de Rham map F^k defines an isomorphism of space $\mathcal{P}_1^- \Lambda^k(\mathcal{T}_h)$ of Whitney *k*-forms onto the space $C^k(\mathcal{T}_h)$ of simplicial *k*-cochains. The collection of these isomorphisms, as we vary *k*, form a cochain map, as we already saw in Chapter 2. In short, *the de Rham map is a cochain isomorphism between the Whitney form complex and the simplicial cochain complex.*

We can use the Whitney forms to prove a version of de Rham's theorem, which we stated at the end of Chapter 2. The cohomology associated to the Whitney form complex (7.11) is certainly isomorphic to the simplical cohomology. Indeed, the complexes themselves, not just their cohomology, are isomorphic. But now consider a bounded cochain projection, which we know exists, from the L^2 de Rham complex to the Whitney complex. By Theorem 5.1, this induces an isomorphism of the de Rham cohomology onto the cohomology of the Whitney complex, as long as the approximation property (5.8) holds. This is certainly true if the mesh size is sufficiently small, which can always be obtained by refining the mesh. But it is in fact true for any mesh, since the simplicial cohomology is easily shown to be unchanged by refinement.

Remark 7.7. We have proven that the L^2 de Rham cohomology is isomorphic to the simplicial cohomology. De Rham's theorem is more typically stated in terms of the smooth de Rham cohomology, but this is not a significant difference, since the smooth and L^2 de Rham cohomology are themselves isomorphic via the inclusion map, essentially due to elliptic regularity. See [24, Theorem 2.12]. More significant is that the isomorphism we have given is induced by the smoothed projection operators (bounded cochain projections), rather than by the de Rham map, which is not defined on the full $H\Lambda^k$ spaces.

7.7 • Cubical elements

Although we have only discussed simplicial triangulations to this point, finite element spaces based on meshes of *n*-dimensional boxes (products of intervals)

⁴In fact, the trace of a Whitney k-form on a k-face f is a constant k-form, since it belongs to $\mathcal{P}_1^-\Lambda^k(f)$ and on a k-dimensional domain, $\mathcal{P}_1^-\Lambda^k$ reduces to $\mathcal{P}_0\Lambda^k$.

are also popular. It turns out that there are two families of finite element spaces on boxes which are in many ways parallel to the $\mathcal{P}_r\Lambda^k$ and $\mathcal{P}_r^-\Lambda^k$ spaces on simplices and which we summarize in this section.

The simpler of the two families is denoted by $Q_r^- \Lambda^k$. It is defined on an arbitrary mesh of boxes \mathcal{T}_h in \mathbb{R}^n and for every r > 0 and $0 \le k \le n$. To define the space, we need to supply shape functions and DOFs on an arbitrary box $T \subset \mathbb{R}^n$. This can be done using the appropriate notion of a tensor product of chain complexes. Here we quickly recall the results, which are extracted from [9], to which we refer for details.

Without loss of generality, we take $T = I^n$, where I = [0, 1]. Then the shape function space is

$$\mathcal{Q}_r^-\Lambda^k(I^n) = \bigoplus_{1 \le \sigma_1 < \dots < \sigma_k \le n} \left[\bigotimes_{i=1}^n \mathcal{P}_{r-\delta_{i,\sigma}}(I) \right] dx^{\sigma_1} \wedge \dots \wedge dx^{\sigma_k},$$

where

$$\delta_{i,\sigma} = \begin{cases} 1, & i \in \{\sigma_1, \dots, \sigma_k\}, \\ 0 & \text{otherwise.} \end{cases}$$

In the case k = 0, this space is understood to be

$$\mathcal{Q}_r^- \Lambda^0(I^n) = \bigotimes_{i=1}^n \mathcal{P}_r(I),$$

i.e., the space conventionally referred to as $Q_r(I^n)$. This definition makes sense also if r = 0, so $Q_0^- \Lambda^0(I^n) = Q_0(I^n) = \mathbb{R}$ is the space of constant functions. For k > 0, $Q_0^- \Lambda^k(I^n) = 0$. We also allow n = 0 in this definition, i.e., when I^n reduces to a single point. We then understand $Q_r^- \Lambda^0(I^0)$ to be the space \mathbb{R} of constants. It is easy to count the dimension of these spaces:

$$\dim \mathcal{Q}_r^- \Lambda^k(I^n) = \binom{n}{k} (r+1)^{n-k} r^k, \quad 0 \le k \le n, \ r \ge 0$$

Finally we specify a set of DOFs for $\mathcal{Q}_r^- \Lambda^k(I^n)$ $(r \ge 1, 0 \le k \le n)$,

$$v \mapsto \int_{f} \operatorname{tr}_{f} v(x) \wedge q(x), \quad q \in \mathcal{Q}_{r-1}^{-} \Lambda^{d-k}(f),$$
 (7.12)

for each face f of I^n of degree $d \ge k$. As shown in [9] their unisolvence follows from tensor product considerations. For each $r \ge 1$, the resulting finite element spaces combine to give a finite element de Rham subcomplex:

$$0 \to \mathcal{Q}_r^- \Lambda^0(\mathcal{T}_h) \xrightarrow{d} \mathcal{Q}_r^- \Lambda^1(\mathcal{T}_h) \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{Q}_r^- \Lambda^n(\mathcal{T}_h) \to 0.$$

Note that, like the $\mathcal{P}_r \Lambda^k$ sequence, the polynomial degree remains constant across the complex. Moreover, the canonical projections associated to the DOFs form a cochain map.

The second family of cubical finite elements was derived in [7], using the tools of FEEC. We refer to it as the *serendipity family* of finite element differential forms and denote the finite element space by $S_r \Lambda^k$. It has a somewhat complicated definition of the shape functions. First we define the space of polynomial differential forms $\mathcal{H}_{r,s}\Lambda^k$ as the span of those monomials $(x^1)^{\alpha_1}\cdots(x^n)^{\alpha_n} dx^{\sigma_1} \wedge \cdots \wedge dx^{\sigma^k}$ which are

- (i) of polynomial degree r (i.e., $\alpha_1 + \cdots + \alpha_n = r$),
- (ii) are linear in at least s variables, not counting the variables x^{σ_i} , i.e.,

$$#(\{\alpha_i=1\}\setminus\{\sigma_1,\ldots,\sigma_k\})\geq s.$$

Then we take as shape function space the space of polynomial differential forms

$$\mathcal{S}_r\Lambda^k := \mathcal{P}_r\Lambda^k \oplus \bigoplus_{s \ge 1} [\kappa \,\mathcal{H}_{r+s-1,s}\Lambda^{k+1} \oplus d\,\kappa \,\mathcal{H}_{r+s,s}\Lambda^k].$$

In contrast to this intricate definition of the shape functions, the DOFs for the $S_r \Lambda^k$ family are mercifully simple. To a face f of dimension d we associate the DOFs

$$u \mapsto \int_{f} (\operatorname{tr}_{f} \omega) \wedge \mu, \quad q \in \mathcal{P}_{r-2d} \Lambda^{d-k}(f),$$

where the weighting functions are taken from ordinary full polynomial spaces, but their degrees drop by two each time the dimension of the face increases by one. In [7], we show that these DOFs are unisolvent, and the main properties of the family are determined. In particular, it is shown that the assembled finite element spaces combine into a de Rham subcomplex with commuting canonical projections:

$$0 \to \mathcal{S}_r \Lambda^0(\mathcal{T}_h) \xrightarrow{d} \mathcal{S}_{r-1} \Lambda^1(\mathcal{T}_h) \xrightarrow{d} \cdots \xrightarrow{d} \mathcal{S}_{r-n} \Lambda^n(\mathcal{T}_h) \to 0.$$

Note that the polynomial degrees decrease as for the $\mathcal{P}_r \Lambda^k$ family.

We remark that some additional families of cubical finite element differential forms have been proposed at least in two or three dimensions, sharing many of the same properties. In particular, we note the TNT element of Cockburn and Qiu [33] and the trimmed serendipity family of Gillette and Kloefkorn [44]. Another remark is that for all these cubical elements, including these discussed above, there is a major loss of accuracy if the cube is mapped not to a box or a parallelepided but to a more general shape (such as a convex quadrilateral in two dimensions, or some sort of hexahedron in three dimensions). This effect, which becomes more pronounced the larger the form degree of the elements, was studied in [9]. Some elements which avoid this have recently been constructed by Arbogast and collaborators [3, 4].

7.8 - Historical antecedents

We have defined the simplicial finite element spaces $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$ and $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$ in a systematic and unified fashion in any dimension n for every $0 \le k \le n$ and $r \ge 1$. In one, two, and three dimensions these spaces of finite element differential forms may be identified via proxies with spaces of scalar- and vector-valued finite element spaces. In this way, they unify a broad range of finite element methods, developed relatively independently by many researchers over a period of many decades.

In any number of dimensions, the spaces $\mathcal{P}_r \Lambda^0$ and $\mathcal{P}_r^- \Lambda^0$ coincide with each other and identify with the most classical finite element space, the Lagrange space of continuous piecewise polynomials of degree at most r, going back in the case r = 1 to Courant [36] in 1943. These are the most natural finite elements to use to discretize the space H^1 . The spaces $\mathcal{P}_r \Lambda^n$ and $\mathcal{P}_{r+1}^- \Lambda^n$ (note the change

of index) also coincide with each other, and they identify with the larger space of all piecewise polynomials of degree at most r, whether continuous or not, the natural finite element discretization of L^2 . Such spaces are used in finite element calculations in the discontinuous Galerkin method, initiated by Reed and Hill [65] in 1973. The spaces $\mathcal{P}_r^- \Lambda^1$ in two dimensions identify with the famous space of Raviart and Thomas [64], introduced in 1975, to discretize H(div). That space was generalized in two different ways to three dimensions in 1980 by Nédélec [59], once as a discretization of H(div), the other as a discretization of H(curl). These correspond to the spaces $\mathcal{P}_r^-\Lambda^1$ and $\mathcal{P}_r^-\Lambda^2$ in three dimensions, respectively. Finally, the two-dimensional space $\mathcal{P}_r\Lambda^1$ identifies with the Brezzi–Douglas–Marini space from 1985 and the three-dimensional spaces $\mathcal{P}_r \Lambda^1$ and $\mathcal{P}_r \Lambda^2$ with their generalization by Nédélec in 1986, referred to as the Nédélec edge and face elements of the second kind. Thus, in $n \leq 3$ dimensions, all these spaces have appeared in the finite element literature. The FEEC treatment revealed the unity among these spaces, which had arisen in diverse contexts. In some cases, the most natural DOFs for the spaces were not known before the development of FEEC.

It is also worth noting that some of these spaces of piecewise polynomial differential forms appeared independently in the context of geometry. The spaces $\mathcal{P}_1^- \Lambda^k$ are highly geometric, having one DOF per k-simplex in the triangulation, and so being isomorphic to the space of simplicial k-cochains. Indeed, these spaces were constructed by Whitney in his 1957 book [70], decades before their appearance as mixed finite elements. The connection to mixed finite elements was pointed out much later yet, by Bossavit [20] in 1988. Further, the complete polynomial spaces $\mathcal{P}_r \Lambda^k$ associated to a triangulation make an appearance in the topological work of Sullivan [67, 68] in the 1970s, a decade before their reinvention as Brezzi-Douglas-Marini elements. In a largely overlooked conference proceedings from the early 1980s [17], Baker called these Sullivan–Whitney forms and he analyzed the convergence of a discretization of an eigenvalue problem for the Hodge Laplacian using them. A first unified treatment of the $\mathcal{P}_r^- \Lambda^k$ spaces was made in a seminal paper of Hiptmair [50] in 1999. In the 2006 paper of Arnold, Falk, and Winther [11] in which the term *finite element exterior calculus* first appeared, the Koszul complex was first applied to finite elements, allowing a major simplification and a unified treatment of both the $\mathcal{P}_r^- \Lambda^k$ and $\mathcal{P}_r \Lambda^k$ spaces.

For the cubic finite element spaces, we briefly described the two families $Q_r^- \Lambda^k(\mathcal{T}_h)$ and $\mathcal{S}_r \Lambda^k(\mathcal{T}_h)$. The former can be considered a tensor product analogue of the $\mathcal{P}_r^- \Lambda^k(\mathcal{T}_h)$ spaces and had been derived alongside the corresponding simplicial spaces. For example, in two dimensions the $Q_r^- \Lambda^1$ spaces are the quadrilateral Raviart–Thomas spaces from [64]. The $\mathcal{S}_r \Lambda^k$ family is different in that it was largely absent from the literature and its discovery was inspired by FEEC. The elements for 0-forms, ordinary H^1 finite elements, are well-known in two dimensions. They are the popular serendipity finite elements to higher dimensions. The $\mathcal{S}_r \Lambda^0$ gives the systematic generalization of these elements to higher dimensions. The \mathcal{S}_r spaces for 1-forms and 2-forms in three and more dimensions were likewise not previously known.

Box 7.1. Summary of Chapter 7 on finite element differential forms.

Give an *n*-dimension simplicial triangulation \mathcal{T}_h for any $n \ge 1$, any form degree $0 \le k \le n$, and any polynomial degree $r \ge 1$, there are two canonical spaces of finite element differential *k*-forms consisting of piecewise polynomials of degree *r* or less, the complete polynomial space $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$ and the trimmed polynomial space $\mathcal{P}_r \Lambda^k(\mathcal{T}_h)$. Like any finite element spaces, these can be specified by giving their shape functions and DOFs and showing that these are unisolvent.

As the name suggests, the shape functions for the complete polynomial space consist of all k-forms with coefficients which are polynomials of degree at most r. The trimmed polynomial space restricts the coefficient polynomials to a subspace of all polynomials of degree at most r but which contains at least all polynomials of degree at most r - 1. It is precisely defined in terms of the Koszul complex. In the case of 0-forms, the trimmed space coincides with the complete polynomial space of the same degree, while for the case of n-forms, it coincides with the complete polynomial space is not a complete polynomial space. In the lowest order case r = 1, the trimmed spaces are the spaces of Whitney forms.

The complete polynomial spaces with r decreasing as k increases combine to form a subcomplex of the L^2 de Rham complex with bounded cochain projections. The trimmed polynomial spaces similarly form a de Rham subcomplex, but with degree r constant.

Two such spaces, one for k-forms and one for (k-1)-forms, either both complete, both trimmed, or one of each, can be used to solve the k-form Hodge Laplacian, as indicated in (7.10).

Similar families exist for cubical meshes rather than simplicial triangulations.

Chapter 8 Further directions and applications

So far we have considered as our Hilbert complex the L^2 de Rham complex on a domain in \mathbb{R}^n , so that the resulting operator is the Hodge Laplacian. For $0 \le k \le n$, the segment (7.1) gives rise to the weak formulation: given $f \in L^2\Lambda^k$, find $\sigma \in H\Lambda^{k-1}$, $u \in H\Lambda^k$, and $p \in \mathfrak{H}^k$ such that

$$\langle \sigma, \tau \rangle - \langle u, d\tau \rangle = 0, \qquad \tau \in H\Lambda^{k-1}, \\ \langle d\sigma, v \rangle + \langle du, dv \rangle + \langle p, v \rangle = \langle f, v \rangle, \quad v \in H\Lambda^k, \\ \langle u, q \rangle = 0, \qquad q \in \mathfrak{H}^k,$$

$$(8.1)$$

for which the corresponding strong formulation is given in (7.2). We have shown stability, consistency, and convergence for the Galerkin solution to this problem using any of the four pairs of spaces in (7.10) and any polynomial degree $r \ge 1$. The corresponding rates of convergence, for smooth solutions, are shown in Table 7.1.

In this chapter, we quickly survey some other problems that can be addressed in the FEEC framework.

8.1 - Eigenvalue problems

Besides the Hodge Laplacian source problem, we are often interested in solving the associated eigenvalue problem

$$(\delta d + d\delta)u = \lambda u$$

together with boundary conditions, for $\lambda \in \mathbb{R}$, $u \neq 0$. (Note that $\lambda = 0$ is an eigenvalue if and only if there exist nonzero harmonic forms, and the harmonic forms form the eigenspace associated to the zero eigenvalue.)

Setting $\sigma = d^*u$, we obtain the mixed weak formulation of the eigenvalue problem: find $\lambda \in \mathbb{R}$, $\sigma \in H\Lambda^{k-1}$, $0 \neq u \in H\Lambda^k$ such that

$$\begin{split} \langle \sigma, \tau \rangle - \langle u, d\tau \rangle &= 0, \qquad \tau \in H\Lambda^{k-1}, \\ \langle d\sigma, v \rangle + \langle du, dv \rangle &= \lambda \langle u, v \rangle, \quad v \in H\Lambda^k. \end{split}$$

The finite element discretization is then clear. We replace the spaces $H\Lambda^{k-1}$ and $H\Lambda^k$ with a stable pair of spaces of finite element differential forms, V_h^{k-1} and
V_h^k , which we have derived for the source problem. This leads to a symmetric sparse generalized matrix eigenvalue problem, which can be solved by a variety of iterative eigensolvers, such as the Krylov–Schur method.

The convergence theory for eigenvalue problems is quite subtle. A good convergence result must of course show that any of the (infinitely many) eigenvalues of the continuous problem is approximated arbitrarily closely by an eigenvalue of the discrete problem, if the mesh is sufficiently fine. It must also take into account multiplicity, so that an eigenvalue of the continuous problem of multiplicity m is approximated by m eigenvalues of the discrete problem (counting their multiplicity, which will usually be different from m). It must also rule out spurious eigenvalues, that is, show that the discrete eigenvalues do not accumulate, as the mesh is refined, at any number which is not an eigenvalue of the continuous problem. The approximation of eigenvectors has to be stated particularly carefully in the case of multiple eigenvalues.

Fortunately, the theory of convergence of Galerkin approximations to eigenvalue problems has been carefully developed and it applies to the case of the Hodge Laplacian. The basic theory is due Babuška and Osborn [16]. An excellent exposition is in [18]. A brief discussion of its application to the Hodge Laplacian eigenvalue problem is given in [11, Section 8] and [13, Section 3.6] and so will not be repeated here. The main message is that the hypotheses of the Babuška–Osborn theory can be deduced from the improved error estimates of Theorem 5.6, and consequently whenever the hypotheses of that theorem—the compactness property of the Hilbert complex and W-bounded cochain projectors for the subcomplex of Galerkin spaces—are met, we obtain convergence of the eigenvalues and eigenvectors. The convergence rates for the eigenvectors are optimal (determined by the approximation properties of the spaces), and the convergence rates for the eigenvalues are double those of the eigenvectors.

We have already seen an example of eigenvalue convergence for the Hodge Laplacian (1.5) in Table 1.3 of the introduction. By contrast, Table 1.2 showed that a naive method, not based on the mixed formulation, does not converge. In the introduction, we also considered the Maxwell eigenvalue problem curl curl $u = \lambda u$. A naive method based on Lagrange elements led to erroneous results, with the true spectrum either not recognizable from the computation or polluted by persistent spurious eigenvalues, as shown in Table 1.4. Computations made with the weak formulation, but using the lowest order Whitney 1-forms, shown in the rows labeled FEEC in the table, converge nicely. This can be proven by relating the Maxwell eigenvalue problem to the Hodge Laplacian eigenvalue problem. The former arises as the \mathfrak{B}_1^* problem after a Hodge decomposition, and so the convergence analysis for the Hodge Laplacian eigenvalue problem carries over to the Maxwell problem. See [13, Section 3.6.1] for details.

8.2 • Variable coefficients

Next we begin to move beyond the Hodge Laplacian. A simple variation is to change the inner product on one or more of the base Hilbert spaces $W^k = L^2 \Lambda^k$ from the standard L^2 inner product to a weighted inner product:

$$\langle \omega, \mu \rangle = \int_{\Omega} \langle A \omega, \mu \rangle \mathrm{vol}, \quad \omega, \mu \in L^2 \Lambda^k(\Omega).$$

Here A is a function on Ω for which A(x) is a symmetric positive definite operator on Alt^k $T_x\Omega$ for each $x \in \Omega$ (e.g., if Ω is a domain in \mathbb{R}^n , A can be viewed as a function from Ω to the space of symmetric positive definite matrices of size $\binom{n}{k}$). Assuming an upper bound on A and a positive lower bound on its least eigenvalue, the weighted inner product on W^k is equivalent to the usual one, and so the new Hilbert complex inherits all essential properties as the old, such as having closed ranges. The abstract Hodge Laplacian of the complex, however, is different. Leaving the general case to the reader, consider the simple example of the final segment of the de Rham complex on a domain in \mathbb{R}^n . We take as the base spaces the weighted space $L^2\Lambda^{n-1}(\Omega, A)$ described above, where A is an $n \times n$ symmetric positive definite matrix-valued function, and leave the rest of the complex untouched, so the relevant segment is

$$L^2 \Lambda^{n-1}(\Omega, A) \xrightarrow{(\operatorname{div}, H(\operatorname{div}))} L^2 \Lambda^n(\Omega) \to 0.$$
 (8.2)

The weak form of the abstract Hodge Laplace problem for this complex then seeks $\sigma \in H(\text{div}), u \in L^2$ such that

$$\begin{split} \langle A\sigma,\tau\rangle - \langle u,\operatorname{div}\tau\rangle &= 0, \qquad \tau \in H(\operatorname{div}), \\ \langle \operatorname{div}\sigma,v\rangle &= \langle f,v\rangle, \quad v \in L^2. \end{split}$$

This is the standard weak formulation corresponding to the problem

$$A\sigma + \operatorname{grad} u = 0, \quad \operatorname{div} \sigma = f_{\mathfrak{grad}}$$

or, equivalently, $-\operatorname{div} C \operatorname{grad} u = f$, $\sigma = -C \operatorname{grad} u$, where $C = A^{-1}$, together with the Dirichlet boundary condition u = 0 on $\partial\Omega$ which, in this weak formulation, is a natural boundary condition. In short, this variable coefficient partial differential operator is again the abstract Hodge Laplacian of a Hilbert complex, albeit a slightly different Hilbert complex than the L^2 de Rham complex. It may be discretized with the same finite element spaces as in the case where A is the identity, and the analysis carries over without additional considerations. Many other elliptic partial differential equations with variable coefficients can be treated in a similar fashion.

8.3 • Lower order terms

A basic model problem arising in magnetostatics seeks a (vector proxy) 2-form B for which

$$\operatorname{curl}\operatorname{curl} B = f, \quad \operatorname{div} B = 0 \quad \operatorname{in} \Omega$$

for f given in the range of the curl, together with boundary conditions imposed on B. The solution can be found by solving the Hodge Laplace equation

$$\operatorname{curl}\operatorname{curl} B - \operatorname{grad}\operatorname{div} B = f.$$

For time-harmonic problems, an additional zeroth order term occurs in this equation, giving a perturbed Hodge Laplace equation

$$\operatorname{curl}\operatorname{curl} B - \operatorname{grad}\operatorname{div} B - m^2 B = f,$$

where m is the wave number. In magnetohydrodynamics, there is an additional first order advection term induced by the velocity vector field v, so in the static case we have

$$\operatorname{curl}\operatorname{curl} B - \operatorname{grad}\operatorname{div} B - \operatorname{curl}(v \times B) = f.$$

When such lower order terms occur, the differential operator acting on B cannot be realized as the abstract Hodge Laplacian of any Hilbert complex. However, it is not difficult to extend the mixed finite element methods we have studied to obtain viable numerical methods for the Hodge Laplacian to such problems with lower order terms. What is less clear is the convergence analysis for these discretizations, since the analysis presented above does not directly apply to them. In [14] we studied the convergence of the mixed methods of FEEC for differential operators of the form

$$Lu = [(d+l_1)(\delta+l_2) + (\delta+l_3)(d+l_4) + l_5],$$

which includes the most natural lower order perturbations of the Hodge Laplace operator $d\delta + \delta d$. Here the l_i are operators which act as multipliers by smooth coefficient fields. For example, the zeroth order term occuring in the time-harmonic problem is accounted for in l_5 , while the first order perturbation arising in the magnetohydrodynamic problem is given by $l_2B = -v \times B$ (the contraction of a 2-form with a vector field).

The analysis of [14] is quite involved, and we will only briefly summarize the main results. It is possible, of course, that the operator L (including its boundary conditions) admits a kernel, just as the ordinary Helmholtz operator $-\Delta + l$ may have a kernel when l is negative. Suppose that this is not the case, i.e., 0 is not an eigenvalue of L, and so the continuous problem is well-posed. Then the finite element discretizations we have derived for the Hodge Laplacian, all four possibilities given in (7.10), are stable for the perturbed operator L as well, at least for sufficiently fine triangulations. The basic error estimates, as stated in Theorem 5.5, therefore hold. The surprise comes when trying to establish the improved error estimates of Theorem 5.6, which establish optimal L^2 rates of convergence for each of the quantities σ , $d\sigma$, u, and du. It turns out that for some choices of the finite element spaces and some lower order terms (l_1, \ldots, l_5) there can be a loss of one convergence order for either σ or $d\sigma$ or both. The results are summarized in Table 8.1, which should be compared to Table 7.1 (the case when all the l_i vanish). These rates were all established theoretically and verified as sharp computationally in [14], to which we refer for details.

V_h^{k-1}	V_h^k	σ	$d\sigma$	u	du
$\mathcal{P}_r \Lambda^{k-1}$	$\mathcal{P}_r^- \Lambda^k$	$\begin{cases} r+1 & \text{if } l_2 = l_5 = 0, \\ r & \text{otherwise.} \end{cases}$	r	r	r
$\mathcal{P}_r^- \Lambda^{k-1}$	$\mathcal{P}_r^- \Lambda^k$	r	r	r	r
$\mathcal{P}_r \Lambda^{k-1}$	$\mathcal{P}_{r-1}\Lambda^k$	$\begin{cases} r+1 & \text{if } l_2 = l_3 = l_5 = 0, \\ r & \text{otherwise.} \end{cases}$	$\begin{cases} r & \text{if } l_3 = 0, \\ r - 1 & \text{otherwise.} \end{cases}$	r	r-1
$\mathcal{P}_r^- \Lambda^{k-1}$	$\mathcal{P}_{r-1}\Lambda^k$	r	$\begin{cases} r & \text{if } l_3 = 0, \\ r - 1 & \text{otherwise.} \end{cases}$	r	r-1

Table 8.1. L^2 error rates for FEEC solution of the Hodge Laplacian.

8.4 • Differential equations on manifolds

We have mostly taken our domain to be a bounded open set in \mathbb{R}^n . However, the Hodge Laplace problem is defined on any Riemannian manifold, with or without boundary, with a unique solution up to appropriately defined harmonic forms. FEEC may then be used to solve the Hodge Laplace problem on a manifold, but this brings in substantial additional considerations. The solution of PDEs on manifolds is an important area of research itself, even in the case of the Laplace-Beltrami equation, which is another name for the 0-form Laplacian on a manifold. In one common approach, the manifold is itself approximated by a polyhedral (that is, piecewise flat) manifold or perhaps a higher order piecewise polynomial approximation, and the PDE on the manifold related to a PDE on the approximate manifold, which is solved by finite elements. (The implementation of the finite element method on a triangulated polyhedral manifold is straightforward, at least when the manifold is embedded in Euclidean space.) See [39] for an excellent survey on finite elements for such surface PDEs. For the treatment of the Hodge Laplacian for k-forms on a manifold, with k > 0, a similar approach can be used. However, the analysis of the errors requires a significant generalization of the theory of approximation of Hilbert complexes as we presented it in Chapter 5. Such a generalization, and its application to the Hodge Laplacian on an embedded hypersurface in Euclidean space, is developed by Holst and Stern in [51].

8.5 • Parabolic and hyperbolic problems

Besides the Hodge Laplace equation $(d\delta + \delta d)u = f$, FEEC can be used to formulate and solve the related parabolic equation, which seeks a *time-dependent* differential form $u : [0, T] \to \Lambda^k$ satisfying the *Hodge heat equation*

$$u_t + (\delta d + d\delta u) = f, \tag{8.3}$$

together with the *initial condition* $u(0) = u_0$ and, for instance, the natural boundary conditions tr $\star u = 0$ and tr $\star du = 0$. As in the elliptic case, the weak formulation introduces the auxiliary variable $\sigma = \delta u$ and determines $(\sigma, u) : [0, T] \to H\Lambda^{k-1} \times H\Lambda^k$ by

$$\langle \sigma, \tau \rangle - \langle d\tau, u \rangle = 0, \quad \tau \in H\Lambda^{k-1}, \ t \in [0, T],$$

$$\langle u_t, v \rangle + \langle d\sigma, v \rangle + \langle du, dv \rangle = \langle f, v \rangle, \quad v \in H\Lambda^k, \ t \in [0, T],$$

$$(8.4)$$

together with the initial condition. Note that the equations in (8.4) must hold at each time t. In comparison to (8.1), (8.4) includes an additional term involving the time-derivative u_t , but, unlike the elliptic case, even if there are harmonic forms, they do not need to be explicitly accounted for in the weak formulation of the parabolic problem.

The Hodge heat equation is discussed in the FEEC setting in [10, Theorem 4.4]. There we show how the results we have obtained for the Hodge Laplacian can be used together with the Hille–Yosida–Phillips theory to establish well-posedness of the parabolic problem. The paper goes on to analyze the Galerkin method, both semidiscrete and, using a simple time-stepping scheme, fully discrete. Optimal order error estimates are obtained based on an appropriate elliptic projection, whose properties again follow from the elliptic case.

We can also consider the *Hodge wave equation*, which is a second order hyperbolic equation determining a time-dependent k-form $u : [0, T] \to \Lambda^k$ from the differential equation

$$u_{tt} + (\delta d + d\delta)u = f. \tag{8.5}$$

Of course, we now require two initial conditions, $u(0) = u_0$ and $u_t(0) = v_0$. To solve this problem we introduce three auxiliary variables: $\sigma = \delta u$, $v = u_t$, and $\beta = du$, time-dependent differential forms of order k - 1, k, and k + 1, respectively. We may then derive the strong form equations

$$\sigma_t = \delta u_t = \delta v,$$

$$v_t = u_{tt} = -(\delta d + d\delta)u + f = -\delta\beta - d\sigma + f,$$

$$\beta_t = du_t = dv,$$

or, in matrix form,

$$\frac{d}{dt} \begin{pmatrix} \sigma \\ v \\ \beta \end{pmatrix} = \begin{pmatrix} 0 & \delta & 0 \\ -d & 0 & -\delta \\ 0 & d & 0 \end{pmatrix} \begin{pmatrix} \sigma \\ v \\ \beta \end{pmatrix} + \begin{pmatrix} 0 \\ f \\ 0 \end{pmatrix},$$

which reveals the structure of a symmetric hyperbolic system (the key point being that the matrix operator on the right-hand side is skew-symmetric). As initial data, we set $\sigma(0) = \delta u_0$, $v(0) = v_0$, and $\beta(0) = du_0$. Note that the original variable u is not part of the system but can be obtained by integrating $u_t = v$, i.e.,

$$u(t) = u_0 + \int_0^t v(s) \, ds.$$

Turning now to the mixed weak formulation, it seeks $(\sigma, v, \beta) : [0, T] \rightarrow H\Lambda^{k-1} \times H\Lambda^k \times H\Lambda^{k+1}$ satisfying

$$\langle \sigma_t, \tau \rangle - \langle v, d\tau \rangle = 0, \qquad \tau \in H\Lambda^{k-1}, \ t \in [0, T],$$

$$\langle v_t, w \rangle + \langle d\sigma, w \rangle + \langle \beta, dw \rangle = \langle f, w \rangle, \qquad w \in H\Lambda^k, \ t \in [0, T],$$

$$\langle \beta_t, \gamma \rangle - \langle dv, \gamma \rangle = 0, \qquad \gamma \in H\Lambda^{k+1}, \ t \in [0, T].$$

$$(8.6)$$

The well-posedness of this formulation and the convergence of its Galerkin solution are established in the Hilbert complex framework in [63, Chapter 4]. As for the parabolic problem, this is obtained using Hille–Yosida–Phillips theory, which is particularly simple when, as here, the time-independent operator is skew-symmetric. The same skew-symmetry leads immediately to energy conservation, as well. Namely, taking the test functions (τ, w, γ) equal to the solution (σ, v, β) at time t in the weak formulation, we see that

$$\|\sigma\|_{L^2}^2 + \|v\|_{L^2}^2 + \|\beta\|_{L^2}^2$$

remains constant in time.

8.6 • Maxwell's equations

In this section we show how to formulate Maxwell's equations as, essentially, a Hodge wave equation. The equations are, in traditional notation,

$$D_t - \operatorname{curl} H = -j, \tag{8.7}$$

$$B_t + \operatorname{curl} E = 0, \tag{8.8}$$

$$\operatorname{div} B = 0, \tag{8.9}$$

$$\operatorname{div} D = q. \tag{8.10}$$

We view the electric field E and the magnetizing field H as (vector proxies for) 1-forms and the magnetic field B and electric displacement D as 2-forms. Thus, both the curl operators in the equations are the exterior derivative d^1 and both the divergence operators are d^2 . The current density j is thus a 2-form, and the charge density q a 3-form. These equations are supplemented by the constitutive equations

$$D = \epsilon E, \quad B = \mu H,$$

where the permittivity ϵ and permeability μ are positive scalars or symmetric positive definite operators which may vary in space (but not time). Note that equation $D = \epsilon E$ has a clear meaning if we view both D and E as vector fields but less so if we think of it as relating differential forms of different orders. Written correctly in the language of differential forms, the equation is $D = \star_{\epsilon} E$, where the Hodge star operator \star_{ϵ} is formed with respect to an ϵ -weighted inner product rather than the usual unweighted $L^2 \Lambda^k$ inner product.

We note that (8.7) and (8.10) imply a necessary compatibility condition on the data, namely, that $q_t = -\operatorname{div} j$. In the notation of exterior calculus, this says that $q_t = -dj$ or, weakly,

$$\int q_t \wedge \tau - \int j \wedge d\tau = 0, \quad \tau \in \mathring{H}\Lambda^0.$$
(8.11)

The initial data we impose on E, D, B, and H must be compatible as well. Obviously we must have

$$D(0) = \epsilon E(0), \quad B(0) = \mu H(0),$$

and, from (8.9) and (8.10), we require

$$\operatorname{div} B(0) = 0, \quad \operatorname{div} D(0) = q(0),$$

as well. Finally the Maxwell problem is completed with boundary conditions. For simplicity, we take homogeneous electric (perfect conductor) conditions $E \times n = 0$, $B \cdot n = 0$, i.e., tr E = 0 and tr B = 0.

Our immediate goal is to write Maxwell's equations in the form of the 1-form Hodge wave equation. The unknown time-dependent 1-form, denoted by v in (8.6), will be the electric field E. The unknown 2-form, denoted by β above, will be the negative of the magnetic field B. The remaining dependent variable, the 0-form σ , will in fact turn out to vanish identically. It is a slack variable, whose vanishing reflects the compatibility of the data q and j. (Our weak formulation will be well-posed even when the data q and j are incompatible, but, in that case, σ will not vanish, and E and B will not satisfy Maxwell's equations.) As explained above, the variable coefficients ϵ and μ are handled by adjusting the inner product we use on the spaces of L^2 differential forms. For differential 1-forms we weight the inner product by ϵ ,

$$\langle E, F \rangle_{\epsilon} := \int \star_{\epsilon} E \wedge F,$$

which is simply $\int \epsilon E \cdot F \, dx$ when written in terms of the vector proxies. Similarly for differential 2-forms we weight the inner product by μ^{-1} . For the 0-forms we may weight with an arbitrary positive function α .

We are now ready to write the Maxwell system in weak form. With $\sigma \equiv 0$, we get from (8.7)–(8.10) that $(\sigma, E, B) : [0, T] \rightarrow \mathring{H}\Lambda^0 \times \mathring{H}\Lambda^1 \times \mathring{H}\Lambda^2$ satisfy

$$\langle \sigma_t, \tau \rangle_{\alpha} - \langle E, d\tau \rangle_{\epsilon} = \int q \wedge \tau, \qquad \tau \in \mathring{H}\Lambda^0, \ t \in [0, T],$$
(8.12)

$$\langle E_t, F \rangle_{\epsilon} + \langle d\sigma, F \rangle_{\epsilon} - \langle B, dF \rangle_{\mu^{-1}} = -\int j \wedge F, \quad F \in \mathring{H}\Lambda^1, \ t \in [0, T],$$
(8.13)

$$\langle B_t, C \rangle_{\mu^{-1}} + \langle dE, C \rangle_{\mu^{-1}} = 0, \qquad C \in \mathring{H}\Lambda^2, \ t \in [0, T],$$

(8.14)

with the initial conditions $\sigma(0) = 0$, $E(0) = E_0$, $B(0) = B_0$. After changing variables from B to -B, this problem is exactly of the form of the Hodge wave equation (8.6), except that the inner products are weighted, and the right-hand side of the first equation need not vanish. The well-posedness of the problem can then be established in the same way as for (8.6).

Let us show that, under the assumption of compatible data, the slack variable σ vanishes and that, with $D = \star_{\epsilon} E$ and $H = \star_{\mu=1} B$, the quadruple (E, B, H, D) satisfies Maxwell's equations. To see that σ vanishes, we first note that it vanishes at time t = 0 by assumption. Next, we evaluate (8.12) at time t = 0 and use the fact that $\star_{\epsilon} E(0) = D(0)$ and dD(0) = q(0) to conclude that $\sigma_t(0) = 0$. Next, we differentiate (8.12) in time and add it to (8.13), taking $F = d\tau$ in the latter. Using the compatibility condition (8.11), we obtain

$$\langle \sigma_{tt}, \tau \rangle_{\alpha} + \langle d\sigma, d\tau \rangle_{\epsilon} = 0, \quad \tau \in \mathring{H}\Lambda^0, \ t \in [0, T].$$

Taking $\tau = \sigma_t$, we obtain that

$$\frac{d}{dt}(\|\sigma_t\|_{\alpha}^2 + \|d\sigma\|_{\epsilon}^2) = 0.$$

Thus the quantity in parentheses is constant in time, and we have already seen that it vanishes when t = 0. Therefore, σ_t vanishes, i.e., σ is constant in time, and so it too must vanish.

Having established that σ vanishes, it is easy to see that (8.12)–(8.14) imply the Maxwell equations. Equation (8.12) is a weak restatement of (8.10), (8.13) of (8.7), and (8.14) of (8.8). Finally (8.9) follows from (8.8) since div B(0) = 0.

8.7 • The Stokes equations

Recall the mixed weak formulation of the Neumann problem for the Poisson equation on a domain Ω in \mathbb{R}^n . Given $f \in \hat{L}^2$, we seek $u \in \mathring{H}(\operatorname{div}, \Omega)$, $p \in \hat{L}^2(\Omega)$ such that

(The hat in \hat{L}^2 denotes the subspace orthogonal to constants, reflecting the fact that f must have integral zero and that p is determined only up to an additive constant. We could have also worked with L^2 and added terms for the space of harmonic forms, namely, the constants.) The problem (8.15) is the mixed weak formulation of the Hodge Laplacian associated to the segment

$$\overset{\text{div}}{\longrightarrow} \overset{\text{div}}{\longrightarrow} \overset{L^2}{\Omega}(\Omega) \to 0,$$
(8.16)

taken from the end of the L^2 de Rham complex with boundary conditions. This formulation can be easily discretized using the finite element differential forms developed in Chapter 7. We now compare this problem to the problem of Stokes flow.

The equations of Stokes flow, written in strong form, seek a vector field u, the velocity, and a scalar field p, the pressure, satisfying

$$-\mu\Delta u + \operatorname{grad} p = f$$
, div $u = 0$ in Ω .

Here μ is the dynamic viscosity, and the vector field f gives the imposed body force density. A simple choice of boundary conditions is no-slip, meaning that u = 0 on $\partial\Omega$. More realistically we might take the no-slip condition on a portion of the boundary and, on the remainder, the stress boundary condition $2\mu \epsilon(u)n - pn = 0$, stated in terms of the symmetric gradient $\epsilon(u)$ of u (the strain rate) and the outward unit normal n.

Sticking with pure no-slip boundary conditions for simplicity, the weak formulation seeks $u \in \mathring{H}^1(\Omega; \mathbb{R}^n)$, $p \in \hat{L}^2(\Omega)$ such that

$$b(u,v) - \langle p, \operatorname{div} v \rangle = \langle f, v \rangle, \quad v \in \mathring{H}^1(\Omega; \mathbb{R}^n),$$

$$\langle \operatorname{div} u, q \rangle = 0, \qquad q \in \hat{L}^2(\Omega).$$
(8.17)

The bilinear form $b: \mathring{H}^1(\Omega; \mathbb{R}^n) \times \mathring{H}^1(\Omega; \mathbb{R}^n) \to \mathbb{R}$ in (8.17) is given by

$$b(u,v) = 2\mu \langle \epsilon(u), \epsilon(v) \rangle. \tag{8.18}$$

Using the no-slip boundary conditions and the divergence-free condition div u = 0, we can simplify this to $b(u, v) = \mu \langle \operatorname{grad} u, \operatorname{grad} v \rangle$, but we prefer not to, because when stress (or other) boundary conditions are imposed on part of the boundary, this is no longer possible.

The formulation (8.17) shares the terms $\langle p, \operatorname{div} v \rangle$ and $\langle \operatorname{div} u, q \rangle$ with the mixed weak formulation (8.15) of the Neumann problem. However, in (8.17), u, v are sought in $\mathring{H}^1(\Omega; \mathbb{R}^n)$, while in (8.15), they are sought in $\mathring{H}(\operatorname{div})$. This suggests a slightly different complex:

$$\mathring{H}^{1}(\Omega; \mathbb{R}^{n}) \xrightarrow{\operatorname{div}} \hat{L}^{2}(\Omega) \to 0.$$
 (8.19)

The complex (8.19) differs from (8.16) in that more smoothness is required in the first space of the complex. To fully define the complex we must specify the base spaces on which div operates as well as the domain spaces shown in (8.19). In

fact, we consider (8.19) as a *bounded* Hilbert complex, so that the base spaces and domain spaces coincide. As the inner product on the first space $\mathring{H}^1(\Omega; \mathbb{R}^n)$ we use the bilinear form *b*, which is equivalent to the H^1 inner product there. With these choices we see that the weak formulation (8.17) is simply the mixed weak formulation of the abstract Hodge Laplacian for the complex (8.19). The wellposedness of this problem therefore follows from the general theory of Chapter 4, as long as the Hilbert complex is closed and has no harmonic forms. In the present case, this means that we require that the divergence operator map $\mathring{H}^1(\Omega; \mathbb{R}^n)$ onto $\hat{L}^2(\Omega)$. This is a stronger statement than that div maps $\mathring{H}(\operatorname{div}, \Omega)$ onto $\hat{L}^2(\Omega)$ which we established easily in Section 3.4. The stronger result is widely used in the analysis of the Stokes equations, with a variety of proofs available, applying to domains with different degrees of generality. For an excellent discussion of the different approaches and results, see [35, Section 2.2].

Applying the theory of Chapter 5, we obtain conditions for accurate Galerkin discretization of (8.19). It is sufficient that the Galerkin subspaces $V_h \subset \mathring{H}^1(\Omega; \mathbb{R}^n)$ and $S_h \subset L^2(\Omega)$ satisfy the subcomplex property

$$\operatorname{div} V_h \subset S_h \tag{8.20}$$

and that there exist bounded cochain projections:

$$\pi_V : \dot{H}^1(\Omega; \mathbb{R}^n) \to V_h, \quad \pi_S : L^2(\Omega) \to S_h, \quad \operatorname{div} \pi_V v = \pi_S \operatorname{div} v.$$

It then follows from the FEEC theory, or by more elementary arguments, that the Galerkin discretization is stable and satisfies the error estimate

$$||u - u_h||_1 + ||p - p_h||_0 \le c \left(\inf_{v \in V_h} ||u - v||_1 + \inf_{p \in S_h} ||p - q||_0 \right).$$
(8.21)

Stokes element pairs that satisfy the subcomplex property $\operatorname{div} V_h \subset S_h$ have an important additional property. The discrete velocity u_h conserves mass, i.e., it *exactly* solves the continuity equation $\operatorname{div} u_h = 0$. This follows immediately from the Galerkin equation

$$\langle \operatorname{div} u_h, q \rangle = 0$$

with $q = \operatorname{div} u_h \in S_h$. Mass conservation is a very desirable property for many applications, particularly when treating nonlinear and/or time-dependent Navier–Stokes equations. See the recent review [55] for an in-depth discussion of this point. Classical Stokes element pairs (the MINI element, \mathcal{P}_2 – P_0 , Taylor–Hood, etc.), such as are discussed in [45] or [19, Chapter 8], do *not* satisfy the subcomplex property. While they do satisfy the estimate (8.21), they are *not* mass conservative.

The development of mass conservative elements has been an active research area for the last decade or so. The primary approach is the development of finite element subcomplexes of the segment (8.19) and of a longer complex of which this forms a portion. In two dimensions, the full complex in question is

$$0 \to H^2(\Omega) \xrightarrow{\operatorname{curl}} H^1(\Omega; \mathbb{R}^n) \xrightarrow{\operatorname{div}} L^2(\Omega) \to 0.$$

where I have simplified by not imposing boundary conditions or suppressing the harmonic forms. This *Stokes complex* is a smoothed version of the L^2 de Rham complex. Although the discretization of the Stokes equations requires only subspaces of the last two nonzero spaces in the complex, in the construction and

analysis of such subspaces it is often helpful to refer to the entire complex and subcomplexes of it. This draws a connection between mass conservative finite elements for the Stokes equations and finite element subspaces of $H^2(\Omega)$, a challenging issue which has been long studied in finite elements. The first finite element discretization of the two-dimensional Stokes complex was given by Mardal, Tai, and Winther in 2002 [57]. However, this was a nonconforming discretization so not a true subcomplex. In 2013, the first conforming finite element discretization on general meshes was constructed by Guzmán and Neilan [49]. They used low degree piecewise polynomials enriched by certain rational functions. The same year Falk and Neilan [42] constructed a family of conforming piecewise polynomial subcomplexes of the Stokes complex. Numerous other possible element choices are possible as well. See [55, Figure 4.1] for more information.

The development of mass conservative Stokes elements in three dimensions can similarly be based on the *three-dimensional Stokes complex*

$$0 \to H^2 \xrightarrow{\text{grad}} H^1(\text{curl}) \xrightarrow{\text{curl}} H^1(\Omega; \mathbb{R}^3) \xrightarrow{\text{div}} L^2 \to 0,$$

where $H^1(\text{curl})$ consists of all H^1 vector fields whose curls are again H^1 . Neilan constructs a finite element subcomplex in [60]. Another approach, for special mesh geometries, was recently developed by Guzmán, Fu, and Neilan [48].

8.8 • Elasticity

Among the notable successes of the FEEC viewpoint has been the development of stable mixed finite element methods for the equations of elasticity. In this case, the relevant Hilbert complex is not the de Rham complex but rather some version of the *elasticity complex*.

Let Ω be a domain in \mathbb{R}^n occupied by an elastic body. The primary unknowns in elasticity are the displacement field $u : \Omega \to \mathbb{R}^n$ and the stress field $\sigma : \Omega \to \mathbb{S} := \mathbb{R}^{n \times n}_{symm}$. In steady state, the stress field must satisfy the equilibrium condition

div
$$\sigma = f$$
,

where $f: \Omega \to \mathbb{R}^n$ is the load vector field giving the imposed body force per unit volume. Furthermore, the displacement and stress are related by the constitutive equation

$$A\sigma = \epsilon u,$$

where ϵ denotes the infinitesimal strain operator, i.e., the symmetric part of the gradient, and the compliance tensor A is a symmetric positive definite map $\mathbb{S} \to \mathbb{S}$, possibly variable. The elastic boundary value problem is completed with suitable boundary conditions, such as the displacement condition u = 0 on $\partial\Omega$.

The divergence operator, div, acting rowwise, defines a closed operator from L^2 symmetric tensor fields (i.e., symmetric $n \times n$ matrix-valued functions) to L^2 vector fields. Its domain is the space $H(\text{div}, \Omega; \mathbb{S})$ of L^2 symmetric tensor fields with divergence in L^2 . On the space of L^2 symmetric tensor fields, we weight the L^2 inner product by the compliance tensor A, and so we obtain a segment of a Hilbert complex

$$L^2_A(\Omega; \mathbb{S}) \xrightarrow{(\operatorname{div}, H(\operatorname{div}; \mathbb{S}))} L^2(\Omega, \mathbb{R}^n) \to 0.$$
 (8.22)

This is altogether analogous to the complex (8.2), the difference being that here the divergence operator maps symmetric tensor to vectors, rather than vectors to scalars. The constraint of symmetry makes the elasticity system much more difficult to discretize.

As we did for (8.2), we immediately obtain the weak form of the abstract Hodge Laplace problem for this complex. It seeks $\sigma \in H(\text{div}; \mathbb{S}), u \in L^2(\Omega; \mathbb{R}^n)$ such that

$$\begin{split} \langle A\sigma, \tau \rangle - \langle u, \operatorname{div} \tau \rangle &= 0, \qquad \tau \in H(\operatorname{div}, \Omega; \mathbb{S}), \\ \langle \operatorname{div} \sigma, v \rangle &= \langle f, v \rangle, \quad v \in L^2(\Omega; \mathbb{R}^n). \end{split}$$

Clearly, the first equation is a weak form of the constitutive equation and the displacement boundary condition, while the second equation directly translates the equilibrium condition. The next theorem asserts that the divergence operator acting on symmetric tensors is in fact surjective. This implies that the Hilbert complex (8.22) is closed and that there are no harmonic forms. It then follows from the general theory of Chapter 4 that the mixed elasticity system is well-posed.

Theorem 8.1. The divergence operator (acting rowwise) maps $H(\text{div}, \Omega; \mathbb{S})$ onto $L^2(\Omega, \mathbb{R}^n)$.

In fact, we shall prove a stronger result. Analogous to the space \mathbb{S} of symmetric $n \times n$ tensor fields, let \mathbb{K} denote the space of skew-symmetric $n \times n$ tensor fields. The dimensions are dim $\mathbb{S} = \binom{n+1}{2}$, dim $\mathbb{K} = \binom{n}{2}$, respectively. As usual, for any matrix $\tau \in \mathbb{R}^{n \times n}$, skw $\tau \in \mathbb{K}$ denotes its skew-symmetric part.

Theorem 8.2. Given $f \in L^2(\Omega, \mathbb{R}^n)$ and $g \in L^2(\Omega, \mathbb{K})$, there exists $\tau \in H(\operatorname{div}, \Omega; \mathbb{R}^{n \times n})$ such that

$$\operatorname{div} \tau = f, \quad \operatorname{skw} \tau = g.$$

Theorem 8.1 follows from Theorem 8.2 by taking g = 0.

Proof. We shall only prove the result in n = 2 or 3 dimensions, with some remarks on the general case given at the end of the proof.

The construction is based on two ingredients. First is the surjectivity of the divergence acting on $H^1(\Omega; \mathbb{R}^n)$: given $q \in L^2(\Omega)$ we can find $v \in H^1(\Omega; \mathbb{R}^n)$ such that div v = q. In the preceding section, on the Stokes equations, we discussed the more difficult result of surjectivity of the divergence acting on $\mathring{H}^1(\Omega; \mathbb{R}^n)$. Without boundary conditions, the result is much easier. We can, e.g., extend q by zero to a ball containing Ω , then solve the Dirichlet problem with forcing term q on the ball to get an H^2 function, and finally take v to be the restriction to Ω of the gradient of the solution.

The second ingredient is an elementary calculus identity. Primarily for reasons of notation, we present it separately in n = 2 dimensions and then in three dimensions. For n = 2, the space \mathbb{K} is one-dimensional. Thus we may state the theorem in terms of a scalar function g belonging to $L^2(\Omega)$, and we must show that there exists τ such that div $\tau = f$ and asym $\tau = g$, where asym $\tau := \tau_{12} - \tau_{21}$. In the two-dimensional case, the elementary identity we require is that the asymmetric part of the matrix curl of a two-dimensional vector field is the negative of

its divergence. That is,

asym curl
$$\psi = -\operatorname{div} \psi, \quad \psi \in H^1(\Omega; \mathbb{R}^2).$$
 (8.23)

(Here the curl operator is the two-dimensional vector curl applied to each scalar component of ψ to give the rows of a 2 × 2 matrix.) The identity is expressed as well by the commutativity of the following diagram:

$$\begin{array}{ccc} H^1(\Omega; \mathbb{R}^2) & \xrightarrow{\operatorname{curl}} & L^2(\Omega; \mathbb{R}^{2 \times 2}) \\ & & & \downarrow^{id} & & \downarrow^{-\operatorname{asym}} \\ H^1(\Omega; \mathbb{R}^2) & \xrightarrow{\operatorname{div}} & L^2(\Omega) \end{array}$$

With these two ingredients in hand, we proceed to the construction. For arbitrary $f \in L^2(\Omega; \mathbb{R}^2)$, we can find $\rho \in H^1(\Omega; \mathbb{R}^{2 \times 2})$ such that

$$\operatorname{div} \rho = f. \tag{8.24}$$

Indeed, since no constraint on the symmetry of f is imposed, we can simply define the two rows of ρ independently using the surjectivity of the divergence on $H^1(\Omega; \mathbb{R}^2)$. Next we form the scalar function $-g + \operatorname{asym} \rho$ and again apply the surjectivity of the divergence to get a vector field $\psi \in H^1(\Omega; \mathbb{R}^2)$ satisfying

$$\operatorname{div}\psi = -g + \operatorname{asym}\rho. \tag{8.25}$$

Note that the matrix curl of ψ belongs to L^2 and is divergence-free and so belongs to $H(\operatorname{div}, \Omega; \mathbb{R}^{2 \times 2})$. Thus

$$\tau := \rho + \operatorname{curl} \psi \in H(\operatorname{div}, \Omega; \mathbb{R}^{2 \times 2}).$$

Finally, from (8.24), (8.23), and (8.25), we get

$$\operatorname{div} \tau = \operatorname{div} \rho = f, \quad \operatorname{asym} \tau = \operatorname{asym} \rho + \operatorname{asym} \operatorname{curl} \psi = \operatorname{asym} \rho - \operatorname{div} \psi = g,$$

as desired. This proves Theorem 8.2 in the two-dimensional case.

In three dimensions, the space \mathbb{K} is itself three-dimensional, i.e., $\mathbb{K} \cong \mathbb{R}^3$. Thus we take g to be a 3-vector field and impose the constraint asym $\tau = g$, where now

$$\operatorname{asym} \tau = (\tau_{23} - \tau_{32}, \tau_{31} - \tau_{13}, \tau_{12} - \tau_{12}).$$

We also need the operator $S : \mathbb{R}^{3 \times 3} \to \mathbb{R}^{3 \times 3}$ given by

$$S\rho = \rho^T - (\operatorname{Tr} \rho)I,$$

where Tr denotes the matrix trace and I the 3×3 identity matrix. This operator is invertible with

$$S^{-1}\mu = \mu^T - \frac{1}{2}(\operatorname{Tr}\mu)I.$$

The relevant identity in three dimensions can now be expressed as commutativity of the diagram

$$\begin{array}{ccc} H^1(\Omega; \mathbb{R}^{3\times 3}) & \stackrel{\text{curl}}{\longrightarrow} L^2(\Omega; \mathbb{R}^{3\times 3}) \\ & & \downarrow^S & & \downarrow^{-\operatorname{asym}} \\ H^1(\Omega; \mathbb{R}^{3\times 3}) & \stackrel{\text{div}}{\longrightarrow} L^2(\Omega, \mathbb{R}^3) \end{array}$$

The argument then proceeds as in two dimensions. We define first $\rho \in H^1(\Omega; \mathbb{R}^{3\times 3})$ then $\psi \in H^1(\Omega; \mathbb{R}^{3\times 3})$ such that

div
$$\rho = f$$
, div $\psi = -g + \operatorname{asym} \rho$,

and then set

$$\tau = \rho + \operatorname{curl} S^{-1} \psi. \qquad \Box$$

Remark 8.3. To prove the above theorem in general dimension n, it is preferable to view the spaces that enter the statement of the theorem and the proof, namely, spaces of scalar-, vector-, and matrix-valued functions, as spaces of differential forms. The differential forms which so arise are not scalar-valued but take values in $\operatorname{Alt}^m \mathbb{R}^n$ for m = n - 1 or n - 2. Thus, for example, the space $L^2(\Omega; \mathbb{R}^n)$, in which f is given, should actually be thought of as $L^2\Lambda^n(\Omega) \otimes \operatorname{Alt}^{n-1} \mathbb{R}^n$, the space of L^2 differential n-forms on Ω taking values in $\operatorname{Alt}^{n-1} \mathbb{R}^n$. Recall that $\operatorname{Alt}^{n-1} \mathbb{R}^n$ is canonically isomorphic to \mathbb{R}^n (vector proxy) and that differential n-forms can be viewed as simply scalar-valued functions (scalar proxy), so we can indeed canonically identify an element of $L^2\Lambda^n(\Omega) \otimes \operatorname{Alt}^{n-1} \mathbb{R}^n$ with an \mathbb{R}^n -valued L^2 function on Ω . In a similar way we can make identifications of the spaces in which g is given and in which ρ , ψ , and τ are defined. In this context, it is possible to define analogues of the operators asym and S for general n and establish the commuting diagram needed for the proof.

We now turn to discretization of the elasticity system, which, as we have seen, is just the abstract Hodge Laplace problem for the complex (8.22). Thus, we need to develop a finite element subcomplex of the spaces arising in (8.22), to be used to discretize the stress and displacement, respectively. As mentioned above, the symmetry constraint makes this quite difficult. It turns out to be useful to extend the sequence (8.22) to a longer complex, just as (8.2) extends to become the de Rham complex. In two dimensions this longer *elasticity complex* can be written

$$0 \to L^{2}(\Omega) \xrightarrow{(\operatorname{curl}\operatorname{curl},H^{2})} L^{2}(\Omega; \mathbb{S}) \xrightarrow{(\operatorname{div},H(\operatorname{div},\mathbb{S}))} L^{2}(\Omega,\mathbb{R}^{2}) \to 0, \quad (8.26)$$

where, for simplicity, we drop the weighting coefficient A. Here the operator curl curl is the composition of the vector-valued curl of a scalar function, followed by the matrix-valued curl of a vector function. The resulting tensor field,

$$\operatorname{curl}\operatorname{curl}\phi = \begin{pmatrix} \partial^2 \phi/\partial x_2^2 & -\partial^2 \phi/\partial x_2^2 \\ -\partial^2 \phi/\partial x_2^2 & \partial^2 \phi/\partial x_1^2 \end{pmatrix},$$

is known in elasticity theory as the Airy stress function associated to a scalar potential function ϕ .

In three dimensions, the elasticity complex is

$$0 \to L^{2}(\Omega; \mathbb{R}^{3}) \xrightarrow{(\epsilon, H^{1}(\Omega; \mathbb{R}^{3}))} L^{2}(\Omega; \mathbb{S})$$
$$\xrightarrow{\operatorname{curl} T \operatorname{curl}} L^{2}(\Omega; \mathbb{S}) \xrightarrow{(\operatorname{div}, H(\operatorname{div}, \mathbb{S}))} L^{2}(\Omega, \mathbb{R}^{3}) \to 0.$$

In this case the operator $\operatorname{curl} T \operatorname{curl}$, called the St. Venant operator or the incompatibility operator, is obtained by applying the curl to the rows of a symmetric 3×3 matrix field, transposing, and again applying the curl (or, equivalently, by applying the curl first to the rows and then to the columns). Its domain consists of those L^2 symmetric matrix fields with curl T curl also in L^2 . Note that the matrix fields in its range are symmetric. An extensive analysis of this incompatibility operator is given in [2].

If we interchange the diagonal elements and reverse the signs of the offdiagonal elements of the Airy stress function (8.8), we obtain the Hessian of ϕ . This indicates that discretization of the two-dimensional elasticity complex is strongly related to H^2 scalar finite elements. This observation helped guide the way to the discovery of the first stable mixed finite elements for elasticity using polynomial trial functions in two dimensions. These were proposed in [15] in 2002. Together with the classical Hermite quintic, or Argyris, finite element subspace of H^2 , the stress and displacement elements form a finite element subcomplex of (8.26). Since that work many alternatives and improvements have been proposed. The elements of [15] were generalized to three dimensions, first in the lowest degree case [1] and then for general degree [8]. More efficient elements, applicable in any dimension, were developed in [53] and [54]. See the latter reference for a more extensive survey of the relevant literature.

We mention briefly another approach to developing efficient mixed finite elements for elasticity, namely, by imposing symmetry weakly (an idea that goes back to the early days of finite elements [37]). In this approach we obtain a different mixed weak formulation, in which the stress field is sought in a space on which symmetry is not imposed, but instead the symmetry is weakly imposed by orthogonality of the skew-symmetric part of the stress to a Lagrange multiplier. The resulting system may again be viewed as the abstract Hodge Laplacian of a complex, but the relevant complex is now, instead of (8.2),

$$L^2_A(\Omega; \mathbb{R}^{n \times n}) \xrightarrow{((\operatorname{div}, \operatorname{skw}), H(\operatorname{div}; \mathbb{R}^{2 \times 2}))} L^2(\Omega; \mathbb{R}^n) \times L^2(\Omega; \mathbb{K}) \to 0.$$

Note that the differential is now composed of a pair of operators: it associates to a matrix field both its divergence and its skew-symmetric part. We proved above in Theorem 8.2 that this differential is surjective, and so again we obtain a well-posed mixed formulation. This segment can again be extended to a longer complex (the *elasticity complex with weak symmetry*). The discretization of this complex can be accomplished with simpler elements than those needed for the standard elasticity complex. Such elements were developed in three dimensions in [12] and a variety of alternative elements have been proposed as well. See, for example, [32], [46], and the work referenced in [47].

Bibliography

- Scot Adams and Bernardo Cockburn, A mixed finite element method for elasticity in three dimensions, J. Sci. Comput. 25 (2005), no. 3, 515–521. MR 2221175 (Cited on p. 111)
- [2] Samuel Amstutz and Nicolas Van Goethem, Analysis of the incompatibility operator and application in intrinsic elasticity with dislocations, SIAM J. Math. Anal. 48 (2016), no. 1, 320–348. MR 3452250 (Cited on p. 111)
- [3] Todd Arbogast and Maicon R. Correa, *Two families of* H(div) *mixed finite elements on quadrilaterals of minimal dimension*, SIAM J. Numer. Anal. **54** (2016), no. 6, 3332–3356. MR 3574588 (Cited on p. 93)
- [4] Todd Arbogast and Zhen Tao, *Construction of H*(div)-*Conforming Mixed Finite Elements on Cuboidal Hexahedra*, ICES Report 17-17, The Institute for Computational Engineering and Sciences, The University of Texas at Austin, July 2017. (Cited on p. 93)
- [5] Douglas N. Arnold, *Differential complexes and numerical stability*, Proceedings of the International Congress of Mathematicians, Vol. I (Beijing, 2002) (L. Tatsien, ed.), Higher Education Press, Beijing, 2002, pp. 137–157. MR 1989182 (Cited on p. x)
- [6] _____, Spaces of finite element differential forms, Analysis and Numerics of Partial Differential Equations (U. Gianazza, F. Brezzi, P. Colli Franzone, and G. Gilardi, eds.), Springer, New York, 2013, 117–140. MR 3051399 (Cited on pp. 87, 88)
- [7] Douglas N. Arnold and Gerard Awanou, *Finite element differential forms on cubical meshes*, Math. Comp. 83 (2014), no. 288, 1551–1570. MR 3194121 (Cited on pp. 92, 93)
- [8] Douglas N. Arnold, Gerard Awanou, and Ragnar Winther, *Finite elements for symmetric tensors in three dimensions*, Math. Comp. 77 (2008), 1229–1251. MR 2398766 (Cited on p. 111)
- [9] Douglas N. Arnold, Daniele Boffi, and Francesca Bonizzoni, *Finite element differential forms on curvilinear cubic meshes and their approximation properties*, Numer. Math. **129** (2015), no. 1, 1–20. MR 3296150 (Cited on pp. 92, 93)
- [10] Douglas N. Arnold and Hongtao Chen, *Finite element exterior calculus for parabolic problems*, ESAIM Math. Model. Numer. Anal. **51** (2017), no. 1, 17–34. MR 3600999 (Cited on p. 101)
- [11] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther, *Finite element exterior calculus, homological techniques, and applications*, Acta Numer. **15** (2006), 1–155. MR 2269741 (Cited on pp. x, 59, 68, 85, 86, 88, 90, 94, 98)

- [12] _____, Mixed finite element methods for linear elasticity with weakly imposed symmetry, Math. Comp. 76 (2007), 1699–1723. MR 2336264 (Cited on p. 111)
- [13] _____, Finite element exterior calculus: From Hodge theory to numerical stability, Bull. Amer. Math. Soc. (N.S.) 47 (2010), no. 2, 281–354. MR 2594630 (Cited on pp. x, 33, 59, 61, 62, 86, 98)
- [14] Douglas N. Arnold and Lizao Li, *Finite element exterior calculus with lower-order terms*, Math. Comp. 86 (2017), no. 307, 2193–2212. MR 3647955 (Cited on p. 100)
- [15] Douglas N. Arnold and Ragnar Winther, *Mixed finite elements for elasticity*, Numer. Math. 92 (2002), no. 3, 401–419. MR 1930384 (Cited on p. 111)
- [16] I. Babuška and J. Osborn, *Eigenvalue problems*, Handbook of Numerical Analysis, vol. II, North-Holland, Amsterdam, 1991, 641–787. MR 1115240 (Cited on pp. 3, 98)
- [17] Garth A. Baker, *Combinatorial Laplacians and Sullivan-Whitney forms*, Differential Geometry (College Park, MD, 1981/1982), Progress in Mathematics, vol. 32, Birkhäuser Boston, Cambridge, MA, 1983, 1–33. MR 702525 (Cited on p. 94)
- [18] Daniele Boffi, *Finite element approximation of eigenvalue problems*, Acta Numer. 19 (2010), 1–120. MR 2652780 (Cited on pp. 3, 4, 98)
- [19] Daniele Boffi, Franco Brezzi, and Michel Fortin, *Mixed Finite Element Methods and Applications*, Springer Series in Computational Mathematics, vol. 44, Springer, Heidelberg, 2013. MR 3097958 (Cited on pp. 74, 106)
- [20] Alain Bossavit, Whitney forms: A class of finite elements for three-dimensional computations in electromagnetism, IEE Trans. Mag. 135, Part A (1988), 493–500. (Cited on p. 94)
- [21] Dietrich Braess, *Finite Elements: Theory, Fast Solvers, and Applications in Elasticity Theory*, third ed., Cambridge University Press, Cambridge, UK, 2007. Translated from the German by Larry L. Schumaker. MR 2322235 (Cited on p. 3)
- [22] Susanne C. Brenner and L. Ridgway Scott, *The Mathematical Theory of Finite Element Methods*, third ed., Texts in Applied Mathematics, vol. 15, Springer, New York, 2008. MR 2373954 (Cited on p. 3)
- [23] Haim Brezis, Functional Analysis, Sobolev Spaces and Partial Differential Equations, Universitext, Springer, New York, 2011. MR 2759829 (Cited on p. 21)
- [24] Jochen Brüning and Matthias Lesch, *Hilbert complexes*, J. Funct. Anal. **108** (1992), no. 1, 88–132. MR 1174159 (Cited on pp. 33, 91)
- [25] Annalisa Buffa, *Trace theorems on non-smooth boundaries for functional spaces related to Maxwell equations: An overview*, Computational Electromagnetics (Kiel, 2001), Lecture Notes in Computational Science and Engineering, vol. 28, Springer, Berlin, 2003, 23–34. MR 1986130 (Cited on p. 31)
- [26] Annalisa Buffa, Martin Costabel, and Dongwoo Sheen, *On traces for* $\mathbf{H}(\mathbf{curl}, \Omega)$ *in Lipschitz domains*, J. Math. Anal. Appl. **276** (2002), no. 2, 845–867. MR 1944792 (Cited on p. 31)
- [27] Élie Cartan, Sur certaines expressions différentielles et le problème de Pfaff, Ann. Sci. École Norm. Sup. (3) 16 (1899), 239–332. MR 1508969 (Cited on p. 16)

- [28] Snorre H. Christiansen, Stability of Hodge decompositions in finite element spaces of differential forms in arbitrary dimension, Numer. Math. 107 (2007), no. 1, 87–106. MR 2317829 (Cited on p. 90)
- [29] Snorre H. Christiansen, Hans Z. Munthe-Kaas, and Owren Brynjulf, *Topics in structure-preserving discretization*, Acta Numer. **20** (2011), 1–119. MR 2805152 (Cited on p. 82)
- [30] Snorre H. Christiansen and Ragnar Winther, Smoothed projections in finite element exterior calculus, Math. Comp. 77 (2008), no. 262, 813–829. MR 2373181 (Cited on p. 90)
- [31] Philippe G. Ciarlet, *The Finite Element Method for Elliptic Problems*, Classics in Applied Mathematics, vol. 40, SIAM, Philadelphia, 2002. Reprint of the 1978 original (North-Holland, Amsterdam; MR 0520174). MR 1930132 (Cited on p. 3)
- [32] Bernardo Cockburn, Jayadeep Gopalakrishnan, and Johnny Guzmán, A new elasticity element made for enforcing weak stress symmetry, Math. Comp. 79 (2010), no. 271, 1331–1349. MR 2629995 (Cited on p. 111)
- [33] Bernardo Cockburn and Weifeng Qiu, Commuting diagrams for the TNT elements on cubes, Math. Comp. 83 (2014), no. 286, 603–633. MR 3143686 (Cited on p. 93)
- [34] Martin Costabel, A coercive bilinear form for Maxwell's equations, J. Math. Anal. Appl. 157 (1991), 527–541. MR 1112332 (Cited on p. 53)
- [35] Martin Costabel and Monique Dauge, On the inequalities of Babuška-Aziz, Friedrichs and Horgan-Payne, Arch. Ration. Mech. Anal. 217 (2015), no. 3, 873– 898. MR 3356990 (Cited on p. 106)
- [36] Richard Courant, Variational methods for the solution of problems of equilibrium and vibrations, Bull. Amer. Math. Soc. 49 (1943), 1–23. MR 0007838 (Cited on p. 93)
- [37] B. M. Fraejis de Veubeke, *Stress function approach*, Proceedings of the World Congress on Finite Element Methods in Structural Mechanics, Robinson and Associates, Verwood, England, 1975, pp. J.1–J.51. (Cited on p. 111)
- [38] Dario Del Pasqua, Su una nozione di varietà lineari disgiunte di uno spazio di Banach, Rend. Mat. e Appl. (5) 13 (1955), 406–422. MR 0077083 (Cited on p. 57)
- [39] Gerhard Dziuk and Charles M. Elliott, *Finite element methods for surface PDEs*, Acta Numer. 22 (2013), 289–396. MR 3038698 (Cited on p. 101)
- [40] Alexandre Ern and Jean-Luc Guermond, *Theory and Practice of Finite Elements*, Springer, New York, 2004. MR 2050138 (Cited on p. 3)
- [41] Lawrence C. Evans, *Partial Differential Equations*, second ed., Graduate Studies in Mathematics, vol. 19, AMS, Providence, RI, 2010. MR 2597943 (Cited on pp. 28, 75, 77)
- [42] Richard S. Falk and Michael Neilan, Stokes complexes and the construction of stable finite elements with pointwise mass conservation, SIAM J. Numer. Anal. 51 (2013), no. 2, 1308–1326. MR 3045658 (Cited on p. 107)
- [43] Richard S. Falk and Ragnar Winther, *The bubble transform: A new tool for analysis of finite element methods*, Found. Comput. Math. **16** (2016), no. 1, 297–328. MR 3451429 (Cited on p. 90)

- [44] Andrew Gillette and Tyler Kloefkorn, *Trimmed serendipity finite element differential forms*, Math. Comp., published electronically May 18, 2018. (Cited on p. 93)
- [45] Vivette Girault and Pierre-Arnaud Raviart, Finite Element Approximation of the Navier–Stokes Equations, Springer, New York, 1986. MR 0851383 (Cited on p. 106)
- [46] J. Gopalakrishnan and J. Guzmán, A second elasticity element using the matrix bubble, IMA J. Numer. Anal. 32 (2012), no. 1, 352–372. MR 2875255 (Cited on p. 111)
- [47] J. Guzmán, A unified analysis of several mixed methods for elasticity with weak stress symmetry, J. Sci. Comput. 44 (2010), no. 2, 156–169. MR 2659794 (Cited on p. 111)
- [48] Guosheng Fu, Johnny Guzmán, and Michael Neilan, *Exact Smooth Piecewise Polynomial Sequences on Alfeld Splits*, preprint, arXiv:1807.05883, 2018. (Cited on p. 107)
- [49] Johnny Guzmán and Michael Neilan, Conforming and divergence-free Stokes elements on general triangular meshes, Math. Comp. 83 (2014), no. 285, 15–36. MR 3120580 (Cited on p. 107)
- [50] Ralf Hiptmair, Canonical construction of finite elements, Math. Comp. 68 (1999), 1325–1346. MR 1665954 (Cited on p. 94)
- [51] Michael Holst and Ari Stern, Geometric variational crimes: Hilbert complexes, finite element exterior calculus, and problems on hypersurfaces, Found. Comput. Math. 12 (2012), no. 3, 263–293. MR 2915563 (Cited on p. 101)
- [52] Lars Hörmander, *The Analysis of Linear Partial Differential Operators*. III: *Pseudo-differential Operators*, Classics in Mathematics, Springer, Berlin, 2007. Reprint of the 1994 edition. MR 2304165 (Cited on p. 26)
- [53] Jun Hu, Finite element approximations of symmetric tensors on simplicial grids in \mathbb{R}^n : The higher order case, J. Comput. Math. **33** (2015), no. 3, 283–296. MR 3352360 (Cited on p. 111)
- [54] Jun Hu and Shangyou Zhang, Finite element approximations of symmetric tensors on simplicial grids in ℝⁿ: The lower order case, Math. Models Methods Appl. Sci. 26 (2016), no. 9, 1649–1669. MR 3529252 (Cited on p. 111)
- [55] Volker John, Alexander Linke, Christian Merdon, Michael Neilan, and Leo G. Rebholz, On the divergence constraint in mixed finite element methods for incompressible flows, SIAM Review 59 (2017), no. 3, 492–544. MR 3683678 (Cited on pp. 106, 107)
- [56] Tosio Kato, *Perturbation Theory for Linear Operators*, Classics in Mathematics, Springer, Berlin, 1995. Reprint of the 1980 edition. MR 1335452 (Cited on pp. 21, 58)
- [57] Kent Andre Mardal, Xue-Cheng Tai, and Ragnar Winther, A robust finite element method for Darcy–Stokes flow, SIAM J. Numer. Anal. 40 (2002), no. 5, 1605–1631. MR 1950614 (Cited on p. 107)
- [58] James Clerk Maxwell, *Treatise on Electricity and Magnetism*, Vol. I, Oxford University Press, Oxford, UK, 1873. (Cited on p. 69)
- [59] J.-C. Nédélec, Mixed finite elements in R³, Numer. Math. 35 (1980), 315–341. MR 592160 (Cited on p. 94)
- [60] Michael Neilan, Discrete and conforming smooth de Rham complexes in three dimensions, Math. Comp. 84 (2015), no. 295, 2059–2081. MR 3356019 (Cited on p. 107)

- [61] Jindřich Nečas, Sur une méthode pour résoudre les équations aux dérivées partielles du type elliptique, voisine de la variationnelle, Ann. Sc. Norm. Sup. Pisa Cl. Sci. 3 16 (1962), no. 4, 305–326. MR 0163054 (Cited on p. 42)
- [62] Rainer Picard, An elementary proof for a compact imbedding result in generalized electromagnetic theory, Math. Z. 187 (1984), 151–164. MR 753428 (Cited on p. 38)
- [63] Vincent Quenneville-Belair, A New Approach to Finite Element Simulations of General Relativity, Ph.D. thesis, Department of Mathematics, University of Minnesota, 2015. (Cited on p. 102)
- [64] P.-A. Raviart and J. M. Thomas, A mixed finite element method for 2nd order elliptic problems, Mathematical Aspects of Finite Element Methods (Proc. Conf., Consiglio Naz. delle Ricerche (C.N.R.), Rome, 1975), Lecture Notes in Mathematics, vol. 606, Springer, Berlin, 1977, 292–315. MR 0483555 (Cited on p. 94)
- [65] W. H. Reed and T. R. Hill, *Triangular Mesh Methods for the Neutron Transport Equation*, Tech. report LA-UR-73-479, Los Alamos National Lab, Los Alamos, NM, 1973. (Cited on p. 94)
- [66] Joachim Schöberl, A multilevel decomposition result in H(curl), EMG 2005: 8th European Multigrid Conference on Multigrid, Multilevel and Multiscale Methods (P. Wesseling, C.W. Oosterlee, and P. Hemker, eds.), EMG 2005 CD, ISBN 90-9020969-7, 2005. (Cited on p. 90)
- [67] Dennis Sullivan, Differential forms and the topology of manifolds, Proceedings of Manifolds—Tokyo 1973, University of Tokyo Press, 1975, 37–49. MR 0370611 (Cited on p. 94)
- [68] _____, Infinitesimal computations in topology, Inst. Hautes Études Sci. Publ. Math. (1977), no. 47, 269–331 (1978). MR 0646078 (Cited on p. 94)
- [69] Daniel B. Szyld, *The many proofs of an identity on the norm of oblique projections*, Numer. Algorithms **42** (2006), nos. 3–4, 309–323. MR 2279449 (Cited on p. 57)
- [70] Hassler Whitney, *Geometric Integration Theory*, Princeton University Press, Princeton, NJ, 1957. MR 0087148 (Cited on pp. 90, 94)
- [71] Kosaku Yosida, Functional Analysis, Classics in Mathematics, Springer, Berlin, 1995. MR 1336382 (Cited on pp. 21, 25)

Index

adjoint, *see* operator, adjoint Airy stress function, 111 algebra, homological, 11 anticommutativity, 65 assembly, finite element, 82

Betti number, 14 boundary, 12 boundary conditions electric, 47 essential, 5 magnetic, 5 natural, 5 no-slip, 105 perfect conductor, 103

chain

boundary of, 13 simplicial, 13 chain projection, 15 coboundary, 16 cocycle, 16 codifferential, 75 coefficients, variable, 98 cohomology, 16 reduced, 34 compactness property, 36 complex chain. 12 cochain. 16 de Rham. 16 dual, 16 elasticity, 107, 110 exact. 12 Hilbert, 33 approximation of, 51 closed, 34 dual, 34 Fredholm, 34

Koszul, 83 L^2 de Rham, 37, 74 simplicial, 12 pure, 12 simplicial chain, 12 simplicial cochain, 16 Stokes, 106 tensor product, 92 compliance tensor, 107 contraction, 65, 70 cycle, 12

degree of freedom (DOF), 82 differential, 12 Koszul, 83 displacement, 107 DOF, *see* degree of freedom duality, Poincaré–Lefschetz, 49

eigenvalue problem Hodge Laplace, 97 Laplace, 3 Maxwell, 7 vector Laplace, 4 eigenvalue, spurious, 9 elasticity, 107 compliance tensor, 107 discretization, 110 displacement, 107 strain, 107 stress, 107 weak formulation, 108 weak symmetry, 111 equation Hodge heat, 101 Hodge Laplace, 78 Hodge wave, 102 Laplace-Beltrami, 101 Maxwell's, 103 Stokes, 104 exterior algebra, 63

exterior calculus, 68 exterior derivative, 71 adjoint of, 78 exterior product, 64, 70

finite element method, 2 finite element space complete polynomial, 82 cubical, 91 serendipity, 92 trimmed polynomial, 85 finite elements, Lagrange, 2, 3 form alternating, 64 differential, 69 finite element, 86 L^2 theory, 74 multilinear, 63 Whitney, 90

Galerkin method, 2 gap, 58 grading, *see* vector space, graded graph norm, 22

harmonic form, 35 discrete, 54 Hille–Yosida–Phillips theory, 101, 102 Hodge decomposition, 36 Hodge star, 66, 70 homological algebra, 11 homology, 12 homotopy formula, 84

incompatibility operator, 110 integration by parts, 74 integration of differential forms, 72 Laplacian abstract Hodge, 39 well-posedness of, 41 Hodge, 39, 81 vector, 4 Leibniz rule, 72 for Koszul differential, 83 load vector, 2 lower order terms, 99

manifold piecewise flat, 101 polyhedral, 101 manifolds, PDEs on, 101 map chain, 15 de Rham, 17 graded, 11 mass conservation, 106 matrix, incidence, 16 Maxwell's equations, 103 weak formulation, 104

operator adjoint, 23 boundary, 12 closed, 22 densely defined, 21 domain of, 21 graph, 22 null space, 22 range, 22 unbounded, 21 Piola transform, 74 Poincaré inequality, 2, 37 discrete, 59 projection bounded cochain, 56 canonical, 88 proxy, 66, 70 pullback, 67, 70

shape functions, 82 simplex, 12 boundary of, 13 face, 12 orientation, 13 St. Venant operator, 110 stiffness matrix, 2 Stokes equations, 104 weak formulation, 105 strain, 107 stress, 107 subcomplex, 12 subcomplex property, 55 subspace, graded, 12

theorem closed graph, 22 closed range, 25 de Rham, 18, 91 Stokes, 73 trace operator on differential forms, 70 on H(curl), 31 on H(div), 29 on H^1 , 28 on $H\Lambda$, 77 triangulation, 13

unisolvence, 87

vector space, graded, 11 volume form, 66

wedge product, 64, 70

Computational methods to approximate the solution of differential equations play a crucial role in science, engineering, mathematics, and technology. The key processes that govern the physical world—wave propagation, thermodynamics, fluid flow, solid deformation, electricity and magnetism, quantum mechanics, general relativity, and many more—are described by differential equations. We depend on numerical methods for the ability to simulate, explore, predict, and control systems involving these processes.

The finite element exterior calculus, or FEEC, is a powerful new theoretical approach to the design and understanding of numerical methods to solve partial differential equations (PDEs). The methods derived with FEEC preserve crucial geometric and topological structures underlying the equations and are among the most successful examples of structure-preserving methods in numerical PDEs. *Finite Element Exterior Calculus* aims to help numerical analysts master the fundamentals of FEEC, including the geometrical and functional analysis preliminaries, quickly and in one place. It is also accessible to mathematicians and students of mathematics from areas other than numerical analysis who are interested in understanding how techniques from geometry and topology play a role in numerical PDEs.



Douglas N. Arnold is McKnight Presidential Professor of Mathematics at the University of Minnesota. He is a mathematician and educator whose research focuses on numerical analysis, PDEs, mechanics, and the interplay among these fields. He is known as the originator of FEEC, first presented in his plenary lecture at the International Congress of Mathematicians in 2002. Professor Arnold's many accomplishments include the award of a Guggenheim Fellowship; election as a foreign member of the Norwegian Academy of Science and Letters; Fellowship in the Society for Industrial and Applied Mathematics (SIAM),

the American Association for the Advancement of Science (AAAS), and the American Mathematical Society (AMS); and serving as president of SIAM and director of the Institute of Mathematics and Its Applications (IMA). He received the SIAM Prize for Distinguished Service to the Profession and the J. Tinsley Oden Medal of the U.S. Association for Computational Mechanics. In 2007 he coauthored an award-winning video, "Möbius Transformations Revealed," which has garnered over two million views on YouTube.

For more information about SIAM books, journals, conferences, memberships, or activities, contact:



Society for Industrial and Applied Mathematics 3600 Market Street, 6th Floor Philadelphia, PA 19104-2688 USA +1-215-382-9800 • Fax: +1-215-386-7999 siam@siam.org • www.siam.org



CB93