

---

# Structure-Preserving Flow Reconstruction from Particle Tracking Data

A Mimetic Spectral Element Approach with Application to Flow  
over a Surface-Mounted Cube

---

Ihab Benyahia

to obtain the degree of Master of Science  
in Aerospace Engineering at Delft University of Technology,  
to be defended publicly on 5<sup>th</sup> June 2026.

Student Number: 4662679

Committee members:	Prof. dr. ir. F. Scarano,	Chair
	Prof. dr. ir. M. Gerritsma,	Supervisor
	Dr. ir. S. Shrestha,	Daily Supervisor
	Dr. ir. T. Lima Pereira,	External Examiner

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.



بِسْمِ اللَّهِ الرَّحْمَنِ الرَّحِيمِ

*In the name of Allah, the Most Gracious, the Most Merciful.*



# Preface

This thesis marks the end of my time as a student at TU Delft. It is the product of long days, longer nights, and the help of many people.

To Marc Gerritsma, thank you for giving me the chance to take on this project and for the trust and freedom that came with it. To Suyash Shrestha, thank you for being there day after day, for the steady guidance, and for taking my questions seriously. I could not have asked for a better daily supervisor.

To my parents: everything I have done as a student has been possible because of you. You raised me with patience, you supported every step I took, and you never made me feel that the path I was choosing was anything but the right one. The quiet sacrifices you made so that I could be here, finishing a degree in Delft, are not lost on me. Thank you for the home you built, for the values you put in me, and for the steady support you have given me through every season. There is no part of this that is not also yours.

To my wife: you never sat with me through equations or read a single line of this thesis, and you did not need to. You were there in a different way. You held the weeks together when they threatened to come apart, you bore my long nights and longer silences with patience, and you accepted, without ever asking me to be otherwise, the version of me that disappears into work for days at a time. I hope, now that this is finished, that you can also see why I did it. Some of this thesis belongs to you.

To my friends: thank you for staying close even when I was hard to reach. Thank you for the late dinners, the long walks, and for showing up exactly when it counted without making a thing of it. You have been a kind of second family to me here, and that has carried me further than you probably know.

Ihab Benyahia  
Delft, June 2026



# Contents

<b>Preface</b>	<b>i</b>
<b>List of Figures</b>	<b>v</b>
<b>List of Tables</b>	<b>xi</b>
<b>Nomenclature</b>	<b>xiii</b>
<b>1 Introduction</b>	<b>1</b>
1.1. Research Objective and Contributions . . . . .	3
1.2. Research Questions . . . . .	4
1.3. Thesis Outline . . . . .	5
<b>2 Background</b>	<b>7</b>
2.1. Particle Image Velocimetry . . . . .	8
2.2. Spectral Element Methods . . . . .	13
2.3. Polynomial Basis Functions and Quadrature Rules . . . . .	16
2.4. Function Spaces and the de Rham Complex . . . . .	21
2.5. Mimetic Spectral Element Discretization . . . . .	26
2.6. Error Analysis for Spectral Element Approximation . . . . .	33
<b>3 Methodology</b>	<b>41</b>
3.1. Problem Statement and Function Space Setting . . . . .	42
3.2. Domain Decomposition and Coordinate Mapping . . . . .	44
3.3. Constrained Least-Squares Formulation for Velocity . . . . .	45
3.4. Reconstruction of Velocity and Streamfunction . . . . .	50
3.5. Vorticity Reconstruction . . . . .	53
3.6. Total and Static Pressure Reconstruction . . . . .	60
3.7. Algorithm Summary . . . . .	64
<b>4 Verification and Validation Setup</b>	<b>69</b>
4.1. Method Verification via Manufactured Solutions . . . . .	70
4.2. Validation Setup: Surface-Mounted Cube Experiment . . . . .	76
<b>5 Results and Discussion</b>	<b>81</b>
5.1. Verification Results: Manufactured Solutions . . . . .	82

5.2. Validation Results: Surface-Mounted Cube . . . . .	114
<b>6 Conclusions and Recommendations</b>	<b>129</b>
6.1. Conclusions . . . . .	130
6.2. Recommendations . . . . .	131
<b>References</b>	<b>133</b>

# List of Figures

2.1	Schematic of a planar 2D-2C PIV setup. The laser beam is shaped into a light sheet using optics and mirrors. Seeded particles scatter light that is captured by a high-speed camera. Cross-correlation of image pairs yields velocity vectors. Adapted from Raffel, Willert, Wereley, and Kompenhans [1]. . . . .	11
2.2	Mapping from the reference element $\hat{\Omega} = [-1, 1]^2$ to a physical element $\Omega_k = [a_x, b_x] \times [a_y, b_y]$ through an affine coordinate transformation $\mathbf{F}_k$ . All basis functions and quadrature rules are defined on $\hat{\Omega}$ and mapped to $\Omega_k$ via the Jacobian of $\mathbf{F}_k$ . . . . .	15
2.3	One-dimensional GLL and GL node distributions for $p = 3$ . . . . .	18
2.4	Two-dimensional tensor-product GLL and GL node distributions on $[-1, 1]^2$ for $p = 3$ . . . . .	19
2.5	Lagrange basis functions $h_i(\xi)$ on GLL nodes for $p = 3$ . . . . .	20
2.6	Edge basis functions $e_i(\xi)$ on GLL nodes for $p = 3$ . . . . .	21
2.7	The discrete de Rham complex with Hodge operators. The primal complex (top) is connected to the dual complex (bottom) by the mass matrices $M^0, M^1$ , and $M^2$ (downward arrows) and their inverses $(M^0)^{-1}, (M^1)^{-1}$ , and $(M^2)^{-1}$ (upward arrows). The primal incidence matrices $E^{1,0}$ (discrete curl) and $E^{2,1}$ (discrete divergence) act left to right; the transposed incidence matrices $(E^{1,0})^T$ (discrete gradient) and $-(E^{2,1})^T$ (discrete curl of a vector) act right to left. . . . .	31
3.1	Multi-element mesh with $K_x = K_y = 2$ and $p_\xi = p_\eta = 2$ . Each element contains a local $3 \times 3$ red node grid with no shared nodes. Light gray lines show internal element structure; thick black lines denote element boundaries. . . . .	44
3.2	A single element with $3 \times 3$ nodes. Blue arrows represent velocity flux DoFs across edges: $u_0$ – $u_5$ (horizontal) and $v_0$ – $v_5$ (vertical). . . . .	47
3.3	$2 \times 2$ element mesh ( $K_x = K_y = 2, p_\xi = p_\eta = 2$ ), each with $3 \times 3$ red nodes. Eight Lagrange multipliers $\lambda_0$ – $\lambda_7$ enforce flux continuity: two per shared edge, drawn exactly between the elements. . . . .	49

3.4	Flux directions and node numbering in a single element. Blue arrows represent flux DoFs $u_i, v_i$ ; green nodes are streamfunction values $\psi_i$ . Counterclockwise arrows indicate the marching orientation used to propagate known $\psi$ -values across the element. . . . .	52
3.5	$2 \times 2$ element mesh ( $K_x = K_y = 2, p_\xi = p_\eta = 2$ ) showing how each element's lower-left node (except the first) is initialised from its neighbour: to the right $\psi_{\text{right},0} = \psi_{\text{left},2}$ , and above $\psi_{\text{top},0} = \psi_{\text{bottom},6}$ . . . . .	53
3.6	Association of vorticity DoFs with the primal mesh. The primal grid is shown in black, dual control volumes in blue dashed lines. Each red node corresponds to a vorticity unknown $\omega_i$ , interpreted as counterclockwise circulation around the surrounding dual cell (green arrows). The dual mesh does not extend beyond the element. . . . .	55
3.7	Assembly of boundary contributions. Orange arrows show tangential velocity components along the dual boundary, oriented upwards or rightwards. Green arrows denote counterclockwise vorticity DoFs. The sign pattern in $\mathcal{B}^1$ arises from reconciling these two conventions. . . . .	57
3.8	Four-element ( $K_x = K_y = 2$ ) layout with continuing vorticity indices $\omega_i$ for each element ( $N = M = 3$ ). Double-headed red arrows $\gamma_j$ indicate interface equalities to be enforced between matching nodal vorticities. One interface in each closed loop is intentionally left unconstrained to avoid redundancy and ensure a nonsingular constraint system. . . . .	59
4.1	Mesh configurations for verification tests. Top left: single element covering entire domain. Top right: $2 \times 2$ multi-element mesh. Bottom: $4 \times 4$ multi-element mesh. . . . .	74
4.2	Schematic of the experimental setup for the surface-mounted cube validation, adapted from [11]. An open-jet wind tunnel with a $60 \times 60 \text{ cm}^2$ exit delivers a freestream velocity $U_\infty = 10 \text{ m s}^{-1}$ over a cube of side length $H = 120 \text{ mm}$ mounted on a flat plate. Seven CMOS cameras (C1–C7) and two LED illumination units record the motion of helium-filled soap bubbles seeded upstream. The dashed rectangle indicates the measurement volume. For the present validation, a two-dimensional velocity field is extracted from a spanwise slab centred at $y = 0$ , yielding a two-dimensional view of the cube's square cross-section in the $x$ - $z$ plane. . . . .	77

4.3	Validation case regions in the $x$ - $z$ plane. The cube cross-section (grey, $120 \times 120$ mm) occupies $-60 \leq x \leq 60$ mm, $0 \leq z \leq 120$ mm. The coordinate origin is at the centre of the cube base. Dashed rectangles indicate the three reconstruction domains: Case 1 ( $-120 \leq x \leq -90$ , $150 \leq z \leq 180$ mm), Case 2 ( $-150 \leq x \leq -60$ , $0 \leq z \leq 60$ mm), and Case 3 ( $-60 \leq x \leq 60$ , $120 \leq z \leq 180$ mm). . . . .	79
5.1	Spatial distribution of measurement points on the reference element $[-1, 1] \times [-1, 1]$ , coloured by velocity magnitude $ u $ . (a) $n_r = 400$ , (b) $n_r = 1600$ , (c) $n_r = 3600$ . Single element, $p = 16$ . . . . .	82
5.2	Analytical (left) and reconstructed (right) $x$ -velocity component $u$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	83
5.3	Analytical (left) and reconstructed (right) $y$ -velocity component $v$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	84
5.4	Analytical (left) and reconstructed (right) divergence $\nabla \cdot \mathbf{u}$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	85
5.5	Analytical (left) and reconstructed (right) streamfunction $\psi$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	86
5.6	Analytical (left) and reconstructed (right) vorticity $\omega$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	87
5.7	Analytical (left) and reconstructed (right) total pressure $\tilde{P}$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	88
5.8	Analytical (left) and reconstructed (right) static pressure $p$ on $[-1, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	89
5.9	Analytical (left) and reconstructed (right) $x$ -velocity component $u$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	91
5.10	Analytical (left) and reconstructed (right) $y$ -velocity component $v$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	92
5.11	Analytical (left) and reconstructed (right) divergence $\nabla \cdot \mathbf{u}$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	93

5.12 Analytical (left) and reconstructed (right) streamfunction $\psi$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	94
5.13 Analytical (left) and reconstructed (right) vorticity $\omega$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	95
5.14 Analytical (left) and reconstructed (right) total pressure $\tilde{P}$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	96
5.15 Analytical (left) and reconstructed (right) static pressure $p$ on $[0, 1]^2$ . Top to bottom: $n_r = 400, 1600, 3600$ . Single element, $p = 16$ . . . . .	97
5.16 Analytical (left) and reconstructed (right) x-velocity component $u$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	98
5.17 Analytical (left) and reconstructed (right) y-velocity component $v$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	98
5.18 Analytical (left) and reconstructed (right) divergence $\nabla \cdot \mathbf{u}$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	99
5.19 Absolute flux mismatch at element interfaces for velocity and pressure gradient. Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	99
5.20 Analytical (left) and reconstructed (right) streamfunction $\psi$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	100
5.21 Analytical (left) and reconstructed (right) vorticity $\omega$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	100
5.22 Analytical (left) and reconstructed (right) total pressure $\tilde{P}$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	100
5.23 Analytical (left) and reconstructed (right) static pressure $p$ . Multi-element mesh $K = 4, p = 4, n_r = 1600$ . . . . .	101
5.24 Error norms for single-element ( $p = 16, K = 1$ ) and multi-element ( $p = 4$ and $p = 8, K = 4$ ) configurations. $n_r = 1600$ . . . . .	101
5.25 $H(\text{div}; \Omega)$ -error versus polynomial order $p$ for $K = 1, K = 4$ , and $K = 7$ . . . . .	103
5.26 Streamfunction $L^2$ -error versus polynomial order $p$ for $K = 1, K = 4$ , and $K = 7$ . . . . .	104
5.27 Vorticity $L^2$ -error versus polynomial order $p$ for $K = 1, K = 4$ , and $K = 7$ . . . . .	104
5.28 Total pressure $L^2$ -error versus polynomial order $p$ for $K = 1, K = 4$ , and $K = 7$ . . . . .	105
5.29 $H(\text{curl}; \Omega)$ -error versus polynomial order $p$ for $K = 1, K = 4$ , and $K = 7$ . . . . .	106

5.30	All five error quantities versus polynomial order $p$ at $K = 7$ and $n_r = 360\,000$ . . . . .	107
5.31	$H(\text{div}; \Omega)$ -error versus element size $h$ for $p = 7$ and three data densities. . . . .	108
5.32	Streamfunction $L^2$ -error versus element size $h$ for $p = 7$ and three data densities. . . . .	109
5.33	Vorticity $L^2$ -error versus element size $h$ for $p = 7$ and three data densities. . . . .	110
5.34	Total pressure $L^2$ -error versus element size $h$ for $p = 7$ and three data densities. . . . .	110
5.35	$H(\text{curl}; \Omega)$ -error versus element size $h$ for $p = 7$ and three data densities. . . . .	111
5.36	All five error quantities versus element size $h$ at $p = 7$ and $n_r = 360\,000$ . . . . .	112
5.37	Measured $h$ -refinement convergence slopes versus polynomial order $p$ for all reconstructed quantities at $n_r = 360\,000$ . . . . .	113
5.38	Velocity magnitude $ \mathbf{u} $ for Case 1. Raw particle data, single-element ( $p = 8, K = 1$ ), multi-element ( $p = 4, K = 4$ ), and Hendriksen reference. $\delta = 3$ mm. . . . .	116
5.39	Vorticity $\omega$ for Case 1. Single-element, multi-element, and Hendriksen reference. $\delta = 3$ mm. . . . .	116
5.40	Divergence $\log_{10}  \nabla \cdot \mathbf{u} $ for Case 1. Single-element, multi-element, and Hendriksen reference. $\delta = 3$ mm. . . . .	117
5.41	Streamfunction $\psi$ for Case 1. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	118
5.42	Total pressure $\tilde{P}$ for Case 1. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	118
5.43	Static pressure $\tilde{p}_{\text{static}}$ for Case 1. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	119
5.44	Velocity magnitude $ \mathbf{u} $ for Case 2. Raw particle data, single-element ( $p = 8, K = 1$ ), multi-element ( $p = 4, K = 4$ ), and Hendriksen reference. $\delta = 3$ mm. . . . .	120
5.45	Vorticity $\omega$ for Case 2. Single-element, multi-element, and Hendriksen reference. $\delta = 3$ mm. . . . .	120
5.46	Divergence $\log_{10}  \nabla \cdot \mathbf{u} $ for Case 2. Single-element, multi-element, and Hendriksen reference. $\delta = 3$ mm. . . . .	121
5.47	Streamfunction $\psi$ for Case 2. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	122
5.48	Total pressure $\tilde{P}$ for Case 2. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	122

5.49	Static pressure $\tilde{p}_{\text{static}}$ for Case 2. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	123
5.50	Velocity magnitude $ \mathbf{u} $ for Case 3. Raw particle data, single-element ( $p = 8, K = 1$ ), multi-element ( $p = 4, K = 4$ ), and Hendriksen reference. $\delta = 3$ mm. . . . .	124
5.51	Vorticity $\omega$ for Case 3. Single-element, multi-element, and Hendriksen reference. $\delta = 3$ mm. . . . .	124
5.52	Divergence $\log_{10}  \nabla \cdot \mathbf{u} $ for Case 3. Single-element, multi-element, and Hendriksen reference. $\delta = 3$ mm. . . . .	125
5.53	Streamfunction $\psi$ for Case 3. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	125
5.54	Total pressure $\tilde{P}$ for Case 3. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	126
5.55	Static pressure $\tilde{p}_{\text{static}}$ for Case 3. Single-element and multi-element reconstruction. $\delta = 3$ mm. . . . .	126

# List of Tables

2.1	Two-dimensional function spaces and their inter-element continuity properties. . . . .	23
2.2	Correspondence between differential forms, function spaces, and the physical quantities they represent in this work. . . . .	24
2.3	Theoretical h-refinement convergence rates at fixed polynomial order $p$ . . . . .	39
4.1	Experimental parameters for the surface-mounted cube measurement [11]. . . . .	77
5.1	Theoretical and measured h-refinement convergence rates at $p = 7$ ( $n_r = 360\,000$ ). . . . .	112



# Nomenclature

## Acronyms

Symbol	Description
CMOS	Complementary metal-oxide-semiconductor
DOF	Degrees of freedom
GL	Gauss–Legendre
GLL	Gauss–Lobatto–Legendre
HFSB	Helium-filled soap bubbles
KKT	Karush–Kuhn–Tucker
LPT	Lagrangian Particle Tracking
MMS	Method of Manufactured Solutions
PIV	Particle Image Velocimetry
PTV	Particle Tracking Velocimetry
STB	Shake-the-Box

## Dimensionless numbers

Symbol	Description	Definition
Re	Reynolds number	$\frac{UL}{\nu}$
Re <sub>H</sub>	Reynolds number based on cube side length H	$\frac{U_{\infty}H}{\nu}$
St	Stokes number (ratio of particle response time to flow time scale)	$\frac{\tau_p}{\tau_f}$

## Latin symbols

Symbol	Description	Dimensions	Units
$\hat{\mathbf{t}}$	Unit tangent vector	–	–
$\mathbf{C}$	Vorticity force (Lamb vector), $\mathbf{u} \times \boldsymbol{\omega}$	$LT^{-2}$	$m s^{-2}$
$\mathbf{F}_k$	Coordinate mapping from $\hat{\Omega}$ to $\Omega_k$	–	–
$\mathbf{g}$	Right-hand side vector	–	–
$\mathbf{n}$	Unit outward normal vector	–	–
$\mathbf{q}$	Pressure gradient vector	$LT^{-2}$	$m s^{-2}$
$\mathbf{u}$	Velocity vector / vector of discrete velocity degrees of freedom	$LT^{-1}$	$m s^{-1}$
$\mathbf{u}^d$	Measured (scattered) veloc- ity data	–	–
$\mathbf{v}$	Test function in $H(\text{div}; \Omega)$	–	–
$\mathcal{C}^k$	Cochain space of k-forms	–	–
$\mathcal{O}$	Order of magnitude (Landau notation)	–	–
$\tilde{\mathbf{C}}$	Discrete vorticity force vector	–	–
$\tilde{\mathbf{P}}$	Discrete total pressure vector (dual representation)	–	–
$\tilde{\mathcal{B}}^1$	Boundary contribution vec- tor	–	–
$\mathbf{A}$	Least-squares data matrix	–	–
$d_p$	Particle diameter	$L$	$m$
$E^{1,0}$	Incidence matrix, discrete curl operator	–	–
$E^{2,1}$	Incidence matrix, discrete divergence operator	–	–
$e_i(\xi)$	Edge basis function (1-form)	–	–
$E_{1D}$	One-dimensional incidence matrix	–	–
$H$	Cube side length	$L$	$mm$
$h$	Element size	$L$	$m$
$H(\text{curl}; \Omega)$	Space of $L^2$ functions with $L^2$ curl	–	–

$H(\text{div}; \Omega)$	Space of $L^2$ vector fields with $L^2$ divergence	–	–
$h_i(\xi)$	Lagrange basis function (0-form)	–	–
$J$	Jacobian matrix / determinant of coordinate mapping	–	–
$K$	Number of elements per direction	–	–
$K_\xi, K_\eta$	Number of elements in $\xi$ and $\eta$ directions	–	–
$L^2(\Omega)$	Space of square-integrable functions on $\Omega$	–	–
$M^0$	Mass matrix for 0-forms (inner product of Lagrange basis)	–	–
$M^1$	Mass matrix for 1-forms (inner product of edge basis)	–	–
$M^2$	Mass matrix for 2-forms (inner product of volume basis)	–	–
$N_{\text{fine}}$	Fine reconstruction grid resolution	–	–
$N_\omega$	Vorticity boundary constraint matrix	–	–
$N_d$	Number of measurement points (validation)	–	–
$N_L$	Lagrange multiplier constraint matrix (flux continuity)	–	–
$n_r$	Number of synthetic measurement points (verification)	–	–
$p$	Polynomial order	–	–
$p_{\text{static}}$	Static pressure	$ML^{-1}T^{-2}$	Pa
$u, v, w$	Velocity components	$LT^{-1}$	$ms^{-1}$
$U_\infty$	Freestream velocity	$LT^{-1}$	$ms^{-1}$
$x, y, z$	Physical coordinates	$L$	m

## Greek symbols

Symbol	Description	Dimensions	Units
$\alpha$	Best approximation constant	–	–
$\beta$	Inf-sup stability constant	–	–
$\gamma$	Lagrange multiplier vector (vorticity system)	–	–
$\lambda$	Lagrange multiplier vector (continuity constraint)	–	–
$\omega$	Vector of discrete vorticity degrees of freedom	–	–
$\psi$	Vector of discrete streamfunc- tion degrees of freedom	–	–
$\zeta$	Lagrange multiplier vector (divergence constraint)	–	–
$\Delta$	Difference or increment	–	–
$\delta$	Spanwise slab half-thickness	L	mm
$\eta$	Reference element coordi- nate	–	–
$\hat{\Omega}$	Reference element domain $[-1, 1]^2$	–	–
$\mu$	Dynamic viscosity	$ML^{-1}T^{-1}$	Pa s
$\nu$	Kinematic viscosity	$L^2T^{-1}$	$m^2 s^{-1}$
$\Omega$	Computational domain	–	–
$\omega$	Vorticity (scalar, 2D)	$T^{-1}$	$s^{-1}$
$\Omega_k$	Physical element domain	–	–
$\phi$	Scalar test function	–	–
$\pi$	Projection operator / mathe- matical constant	–	–
$\psi$	Streamfunction	$L^2T^{-1}$	$m^2 s^{-1}$
$\Psi^{(0)}, \Psi^{(1)}, \Psi^{(2)}$	Basis function matrices for 0-forms, 1-forms, 2-forms	–	–
$\rho$	Fluid density	$ML^{-3}$	$kg m^{-3}$
$\sigma$	Exponential convergence rate constant	–	–
$\tau_f$	Characteristic flow time scale	T	s

$\tau_p$	Particle response time	T	s
$\epsilon^h$	Discrete error function	–	–
$\xi$	Reference element coordinate	–	–

## Superscripts

Symbol Description

(0), (1), (2) 0-form, 1-form, 2-form

d Measurement data

h Discrete (finite-dimensional) approximation

k Element index

## Subscripts

Symbol Description

$\infty$  Freestream condition

$\xi, \eta$  Reference element directions

f Fluid quantity

h Discrete approximation

k Element index

p Particle quantity

x, y, z Physical coordinate directions

# Chapter 1

## Introduction

## 1

Particle Image Velocimetry (PIV) [1] and Lagrangian Particle Tracking (LPT) [2] have become the standard tools for obtaining spatially resolved velocity fields in experimental fluid mechanics. These techniques capture the motion of tracer particles seeded into the flow, producing velocity data at thousands or millions of measurement locations within the flow domain. The resulting velocity fields resolve phenomena such as separation, recirculation, vortex dynamics, and turbulent transport.

These techniques exist to provide quantitative flow visualisations, spatially resolved velocity fields that can serve as inputs for downstream analysis such as vortex identification, statistical averaging, and pressure or vorticity recovery. They have proven extraordinarily useful, and most of what is known experimentally about complex separated flows in the laboratory has been measured by one of these techniques. At the same time, the velocity fields they produce inherit several limitations that become important whenever derived quantities are computed by differentiation: measurement uncertainty introduces errors in individual velocity vectors; data gaps arise near solid boundaries where optical access is restricted or where tracer particles are absent; and the measured velocity field does not exactly satisfy the incompressibility constraint, so its divergence is nonzero even in flows that are physically divergence-free. These shortcomings are amplified when derived quantities are computed from the velocity: the vorticity, obtained by differentiation, amplifies measurement errors, and the pressure, computed from the velocity through the momentum equation, is sensitive to errors in the velocity gradients. Without a reconstruction step that enforces physical constraints, errors in the raw velocity propagate directly into the vorticity and pressure.

Several approaches have been developed to address these limitations. Smoothing and filtering methods [3] reduce measurement noise but do not enforce physical constraints. The divergence-free condition remains violated, and the derived quantities still suffer from error amplification through spatial differentiation. Physics-informed methods such as VIC+ [4] and FlowFit [5] incorporate the divergence-free condition, but typically enforce it approximately through penalty terms or iterative projection. The constraint is not satisfied exactly at the discrete level, leaving residual mass conservation errors that propagate into derived quantities. In particular, an energy-balance argument shows why exact satisfaction of  $\nabla \cdot \mathbf{u} = 0$  matters: when  $\nabla \cdot \mathbf{u} \neq 0$ , the pressure does spurious work on the flow through a non-vanishing term  $(\nabla \cdot \mathbf{u}, p)$  that is absent in the continuous problem, which contaminates derived quantities such as the pressure field

itself. Pressure-from-PIV methods [6, 7] compute pressure by integrating the momentum equation using finite differences on the measurement grid. These methods treat velocity processing and pressure computation as separate steps, with no guarantee that the discrete operators used at each step are mutually consistent.<sup>1</sup>

What is missing is a reconstruction framework with three properties at once: the incompressibility constraint is satisfied exactly at the discrete level; vorticity, streamfunction, and pressure are recovered from the velocity by a single, internally consistent procedure; and the reconstruction error of each quantity is theoretically predictable under mesh refinement. The thesis closes this gap by applying a class of high-order discretisation techniques that preserves these algebraic relationships of the continuous problem by construction; the formal objective, the structure of the method, and the contributions of this work are stated in Section 1.1..

### 1.1. Research Objective and Contributions

The objective of this thesis is to develop, verify, and validate a reconstruction methodology that, from scattered experimental velocity data, produces a velocity field satisfying  $\nabla \cdot \mathbf{u} = 0$  exactly at the discrete level together with vorticity, streamfunction, and pressure fields that are recovered through the same discrete operators as the velocity. The method is built on the Mimetic Spectral Element Method [8, 9, 10], a class of high-order discretisation techniques in which the discrete approximation spaces and the discrete differential operators reproduce the algebraic structure of the continuous problem by construction. The technical content (function spaces, the de Rham complex, incidence and mass matrices) is developed in Chapter 2; the reconstruction method itself is developed in Chapter 3.

To the best of the author's knowledge, Mimetic Spectral Element Methods have not previously been applied to the reconstruction of flow fields from scattered experimental particle tracking data. The contributions of this work are:

1. **An exactly divergence-free reconstruction from scattered experimental data.** Existing physics-informed methods (e.g., VIC+ [4], FlowFit [5]) enforce  $\nabla \cdot \mathbf{u} = 0$  approximately through penalty terms or iterative projection; the present method enforces it to machine

---

<sup>1</sup>Throughout this work, *consistent* refers to the property that derived quantities, vorticity, streamfunction, and pressure, are computed using the same set of discrete operators that yield the velocity, so that algebraic identities of the continuous problem (notably  $\nabla \cdot (\nabla \times \psi) = 0$ ) hold exactly at the discrete level.

precision through an algebraic identity, eliminating the spurious pressure-work term  $(\nabla \cdot \mathbf{u}, p)$  that otherwise contaminates derived fields.

2. **Vorticity, streamfunction, and pressure obtained from the *same* discrete operators as the velocity.** Conventional pipelines treat velocity processing, vorticity computation by finite differences, and pressure recovery by momentum integration as separate steps with their own discretisations. The present formulation derives all four fields from a single set of mimetic operators, so identities such as  $\nabla \cdot (\nabla \times \psi) = 0$  hold at the discrete level too.
3. **Theoretically predictable convergence under *h*- and *p*-refinement.** The approximation error of each reconstructed quantity inherits a rate from the polynomial-approximation theory of the underlying function spaces. These rates are derived for the velocity, streamfunction, vorticity, and pressure in Chapter 2 and verified numerically against a manufactured solution in Chapter 5.
4. **Validation on a separated-flow benchmark.** The method is applied to 3D Lagrangian-particle-tracking data of flow over a surface-mounted cube [11], projected to a two-dimensional spanwise slab. Three regions of increasing complexity are studied: a freestream region, the upstream base of the cube, and the recirculation region above the cube.

The scope of this thesis is deliberately restricted to two dimensions, to rectangular axis-aligned meshes with an affine reference-to-physical mapping, and to the steady (time-averaged) reconstruction of an incompressible flow from a fixed dataset; three-dimensional reconstructions, curvilinear elements, time-resolved tracking, and compressible flows are left to future work.

## 1.2. Research Questions

The main research question is:

Can a Mimetic Spectral Element Method reconstruct velocity, vorticity, and pressure fields from scattered particle tracking data while satisfying the incompressibility constraint exactly, and does the reconstruction converge at theoretically predictable rates?

This question is addressed through three sub-questions that follow the development of the method from formulation through verification to application.

**SQ1.** How should the reconstruction be formulated so that the velocity is divergence-free, the vorticity and pressure are recovered through the same discrete operators as the velocity, and all inter-element continuity conditions are enforced exactly?

**SQ2.** What convergence rates does the method achieve under p-refinement and h-refinement, and how do these rates differ between the velocity, streamfunction, vorticity, and pressure?

**SQ3.** How does the method perform when applied to experimental particle tracking data of separated flow over a surface-mounted cube?

### 1.3. Thesis Outline

Chapter 2 introduces Particle Image Velocimetry and its limitations, followed by Spectral Element Methods with their polynomial basis functions and quadrature rules. It then presents the function spaces, differential forms, and the de Rham complex underlying the mimetic framework, the Mimetic Spectral Element discretisation with its incidence matrices and mass matrices, and the theoretical error analysis for h-refinement and p-refinement.

Chapter 3 formulates the reconstruction problem in the function space setting  $\mathbf{u} \in H(\text{div}; \Omega)$ . The domain is decomposed into a multi-element mesh of non-overlapping rectangular elements, each related to a reference element by an affine coordinate mapping. The constrained least-squares system is then developed with exact divergence-free and continuity constraints, and reconstruction procedures are derived for velocity, streamfunction, vorticity, and pressure.

Chapter 4 describes the manufactured solution test case used for code verification and the surface-mounted cube experiment configuration used for validation.

Chapter 5 presents verification results demonstrating constraint satisfaction and convergence rates, followed by validation results on the surface-mounted cube.

Chapter 6 answers each sub-question, states the main conclusions, and identifies directions for future work.



# Chapter 2

## Background

This chapter introduces the theoretical foundations required for the mimetic spectral element reconstruction method developed in this thesis. Section 2.1. describes Particle Image Velocimetry and Lagrangian Particle Tracking, including the measurement principle, the experimental configuration relevant to this work, and the inherent limitations of the raw velocity data. Section 2.2. introduces spectral element methods, the reference element, coordinate mapping, and the distinction between  $h$ -refinement and  $p$ -refinement. Section 2.3. defines the one-dimensional polynomial building blocks (Legendre polynomials, quadrature rules, Lagrange and edge basis functions) from which all two-dimensional operators are constructed. Section 2.4. introduces the Sobolev spaces  $H(\text{curl}; \Omega)$ ,  $H(\text{div}; \Omega)$ , and  $L^2(\Omega)$ , the continuous de Rham complex that relates them, and the consequences of placing the velocity in  $H(\text{div}; \Omega)$ . Section 2.5. shows how the continuous spaces and operators are discretised through incidence matrices, mass matrices, and Piola transformations while preserving the algebraic structure of the de Rham complex exactly. Section 2.6. derives the theoretical convergence rates under  $h$ -refinement and  $p$ -refinement for all reconstructed quantities.

## 2.1. Particle Image Velocimetry

Particle Image Velocimetry is a non-intrusive, optical measurement technique that provides spatially resolved velocity fields in fluid flows. The technique captures the motion of small tracer particles suspended in the flow by recording successive images illuminated by a pulsed laser sheet. Statistical cross-correlation of particle patterns between image pairs yields local displacement vectors, which are converted to velocity through division by the known time interval between exposures. Since its development in the 1980s, PIV has become one of the most widely used experimental methods in fluid mechanics, offering full-field velocity information that point-measurement techniques such as hot-wire anemometry and laser Doppler velocimetry cannot provide [1, 12, 13].

The validation data used in this thesis are not from PIV but from a closely related technique, three-dimensional Particle Tracking Velocimetry (PTV) [11], in which individual particle trajectories are reconstructed by multi-camera triangulation rather than by cross-correlation. The hardware components (seeding, illumination, cameras, calibration) and the downstream data-quality limitations described below are common to both techniques; the specific 3D PTV configuration is described in Chapter 4.

This section describes the operating principle of planar PIV, the main components of a typical experimental setup, and the data-quality limitations

that motivate the physics-based reconstruction method developed in this thesis.

### 2.1.1. Principle of Operation

In a planar PIV measurement, the velocity field is acquired in a two-dimensional slice of the flow. Depending on the optical configuration, either two in-plane velocity components are measured (2D-2C) or, with stereoscopic arrangements, all three components within the measurement plane (2D-3C). The PTV validation data used in this thesis are reduced to 2D-2C form before being passed to the reconstruction (Chapter 4); the description below therefore applies.

The fundamental measurement relation is

$$\underline{\mathbf{u}}(x, y) \approx \frac{\Delta \underline{\mathbf{x}}}{\Delta t}, \quad (2.1)$$

where  $\Delta \underline{\mathbf{x}}$  is the particle displacement vector and  $\Delta t$  is the time separation between the two laser pulses. The relation is exact only in the limit  $\Delta t \rightarrow 0$ . Even in that limit, the recovered velocity corresponds to some point along the particle trajectory between  $\underline{\mathbf{x}}$  and  $\underline{\mathbf{x}} + \Delta \underline{\mathbf{x}}$ , at some time between  $t$  and  $t + \Delta t$ . Neither the exact point nor the exact time is known. For small  $\Delta t$  this ambiguity is negligible, but in principle the result of (2.1) is not the velocity at any single specified location and instant. To determine  $\Delta \underline{\mathbf{x}}$ , each recorded image is divided into small subregions called interrogation windows. For each window in the first image, a matching region is identified in the second image using a Fast Fourier Transform (FFT)-based cross-correlation algorithm. The location of the correlation peak indicates the most probable displacement of the particle pattern within that window. Dividing this displacement by  $\Delta t$  yields the local velocity vector [12].

The spatial resolution of the resulting velocity field is governed by the size and overlap of the interrogation windows. Smaller windows provide finer resolution but require higher seeding density to maintain sufficient particle image pairs for a reliable correlation peak. In practice, adaptive multi-pass algorithms with decreasing window size and window offset are employed to balance resolution and accuracy [1, 14].

### 2.1.2. Experimental Setup

The experimental hardware used in both planar PIV and Particle Tracking Velocimetry (PTV) consists of several synchronised components that together enable the acquisition of velocity fields [1]. The validation data of this thesis come from a three-dimensional PTV campaign [11], but the

components described below (seeding, illumination, cameras, calibration) are the same for both techniques. The specific 3D PTV configuration is reported in Chapter 4.


The flow is seeded with tracer particles that are sufficiently small to follow the local fluid motion faithfully. Typical seeding materials include oil droplets generated by Laskin nozzles, di-ethyl-hexyl-sebacate (DEHS) aerosol, or hollow glass microspheres. The particle response time must be small relative to the smallest time scales of the flow to ensure that the measured particle displacement accurately represents the local fluid velocity. The Stokes number, defined as the ratio of the particle response time to a characteristic flow time scale,

$$St = \frac{\tau_p}{\tau_f}, \quad \tau_p = \frac{\rho_p d_p^2}{18 \mu_f}, \quad (2.2)$$

with  $\rho_p$  the particle density,  $d_p$  the particle diameter, and  $\mu_f$  the fluid dynamic viscosity, should satisfy  $St \ll 1$  for accurate flow tracing [15].

Illumination is provided by a double-pulsed Nd:YAG laser operating at 532 nm wavelength. Each pulse has a duration of approximately 5–10 ns, which effectively freezes the particle motion during exposure. A system of spherical and cylindrical lenses shapes the laser beam into a thin light sheet, typically less than 1 mm in thickness, that illuminates a planar cross-section of the flow. The thickness of the light sheet determines the out-of-plane depth over which particles contribute to the recorded images and should be minimised to reduce perspective errors in 2D measurements.

A CCD or CMOS camera with high quantum efficiency and a global shutter records the scattered light from the illuminated particles. The camera is aligned perpendicular to the light sheet to minimise optical distortion. In double-frame mode, two exposures separated by  $\Delta t$  are captured in rapid succession. A programmable timing unit synchronises the laser pulses and camera exposures with sub-microsecond precision, ensuring that the two images correspond to the intended time separation. Before measurements commence, a calibration target placed in the light sheet plane establishes the mapping between image pixels and physical coordinates, accounting for lens distortion and magnification [1].



**FIGURE 2.1** Schematic of a planar 2D-2C PIV setup. The laser beam is shaped into a light sheet using optics and mirrors. Seeded particles scatter light that is captured by a high-speed camera. Cross-correlation of image pairs yields velocity vectors. Adapted from Raffel, Willert, Wereley, and Kompenhans [1].

Figure 2.1 illustrates the standard optical configuration of a planar PIV system.

### 2.1.3. Challenges and Limitations of PIV Data

While PIV provides dense spatial velocity information, several factors compromise data quality, particularly in complex flow configurations involving separation, recirculation, and regions near solid boundaries.

Measurement noise arises from multiple sources, including image noise from the camera sensor, background scattering from surfaces near the measurement plane, and sub-optimal seeding density that produces weak or ambiguous correlation peaks. In regions of high velocity gradients, the in-plane displacement variation within a single interrogation window reduces the sharpness of the correlation peak, introducing a bias towards lower gradient magnitudes. Out-of-plane particle motion causes loss of particle pairs between the two exposures, further degrading the correlation signal [1].

Spurious vectors (outliers) are an unavoidable feature of PIV processing. These arise when the correlation algorithm locks onto a secondary peak rather than the true displacement peak, or when insufficient particles are present in the interrogation window. Standard post-processing applies median-based filters and local consistency checks to identify and remove

## 2

outliers, replacing them with interpolated values. However, in regions where errors are spatially correlated, such as near walls or in shadow zones, these filters may be insufficient [16].

Data gaps occur in regions where the laser sheet is blocked by solid surfaces or where strong reflections from walls saturate the camera sensor. In flows over bluff bodies, optical access limitations create shadow regions immediately downstream of the obstacle where no reliable velocity data can be obtained. This effect has been documented for separated flows over wall-mounted obstacles, including periodic hill configurations [17], and is equally present in the surface-mounted cube geometry considered in this work.

A fundamental limitation of conventional PIV processing is that no physical constraints are imposed on the measured velocity field. The cross-correlation algorithm treats each interrogation window independently, and the resulting velocity field is not required to satisfy conservation laws. In particular, the divergence-free condition that characterises incompressible flows is not enforced, meaning that the raw PIV data may exhibit non-physical mass sources and sinks. This lack of physical consistency has concrete consequences for derived quantities. Vorticity is obtained from spatial derivatives of the velocity field. Spatial differentiation acts as a high-pass filter. A velocity error of order  $\varepsilon$  at a single grid point therefore produces a vorticity error of order  $\varepsilon/\Delta x$  in the adjacent stencil, so high-wavenumber measurement noise is amplified by a factor proportional to the inverse of the grid spacing. The pressure is recovered by integrating a Poisson equation whose source term contains the velocity gradient and the convective term  $\mathbf{u} \cdot \nabla \mathbf{u}$ . Velocity errors therefore enter the pressure source quadratically and through derivatives that further amplify high-wavenumber noise. In addition, a velocity field with  $\nabla \cdot \mathbf{u} \neq 0$  violates the energy balance of the continuous problem, producing a spurious work term ( $\nabla \cdot \mathbf{u}, p$ ) that contaminates the computed pressure even when the velocity itself is locally accurate [7].

These limitations motivate post-processing methods that enforce physical constraints on the reconstructed velocity field. One such approach is constrained data fitting. Among all velocity fields in a chosen approximation space that satisfy the physical constraint  $\nabla \cdot \mathbf{u} = 0$ , the one that best matches the measurements is selected. The specific choice of “best fit” used in this thesis is a least-squares fit, with the divergence-free condition enforced exactly as a constraint. This combination of an  $\ell^2$ -data-misfit and an exact constraint is one option among several. Alternative choices

include weighted, robust, or regularised norms. The  $\ell^2$  version is selected here because it leads to a linear KKT system that is straightforward to assemble and solve, and because it is compatible with the mimetic discretisation introduced later. The resulting method reduces sensitivity to measurement noise, fills data gaps through physically informed interpolation, and produces velocity fields suitable for computing derived quantities such as vorticity and pressure.

#### 2.1.4. Application to the Surface-Mounted Cube

Flow over a surface-mounted cube [11] presents a canonical configuration of bluff-body separated flow, featuring boundary layer development, separation at the upstream edge, a recirculation zone above and downstream of the cube, shear layer development, and eventual reattachment. Periodic-hill configurations [17, 18] have been studied extensively as benchmarks for separated-flow simulation, but they differ fundamentally from the surface-mounted cube: they use smooth contours, they are planar with no in/outflow boundaries, and the bulk-flow setup is not directly comparable to the present geometry. The cube is used in this thesis only as a representative validation case for the reconstruction method.

Applying PIV to this geometry provides high-resolution experimental velocity data but also exposes the challenges described above. The sharp corners of the square profile create strong velocity gradients and flow reversal that challenge the cross-correlation algorithm. Shadow zones behind the cube edges produce data gaps precisely in the regions of greatest physical interest. The recirculation zone contains low velocities and reversed flow, where the signal-to-noise ratio of the PIV measurement is reduced.

These conditions make the surface-mounted cube an ideal validation case for the constrained least-squares reconstruction method developed in this thesis. The method must accurately capture the separation and reattachment locations, reconstruct smooth and divergence-free velocity fields in regions of strong gradients, and provide self-consistent pressure and vorticity fields from the same underlying reconstruction [19, 20].

## 2.2. Spectral Element Methods

The reconstruction method developed in this thesis approximates the velocity field using high-order polynomial basis functions defined on a decomposition of the domain into non-overlapping quadrilateral elements. This spectral element approach combines the geometric flexibility of finite element methods with the rapid convergence of spectral methods. This

section provides an overview of the spectral element framework and introduces the key concepts underlying the polynomial discretisation detailed in the subsequent sections.

## 2

### 2.2.1. From Finite Elements to Spectral Methods

Classical low-order finite element methods approximate the solution using a piecewise linear basis on each element, while finite volume methods instead store cell-averaged values and reconstruct fluxes across cell faces. Finite element methods are geometrically flexible and naturally extend to complex domains. Finite volume methods are less flexible and typically require special treatments for curved geometries. Both share an algebraic convergence under mesh refinement. The error decreases with  $h$  at a rate set by the local polynomial degree, so achieving high accuracy requires very small  $h$  with correspondingly large numbers of degrees of freedom.

Spectral methods take the opposite approach, representing the solution as a single global expansion in high-degree polynomials or trigonometric functions over the entire domain. For problems with smooth solutions, the approximation error decreases faster than any algebraic power of the number of modes, a property known as spectral or exponential convergence. However, global spectral methods are restricted to simple geometries (rectangles, circles, spheres) and cannot easily accommodate complex domain shapes or localised features [21].

Spectral element methods bridge these two approaches. The computational domain is partitioned into a number of non-overlapping elements, and within each element the solution is represented by a high-degree polynomial. The polynomial degree  $p$  can be chosen independently of the mesh, and inter-element continuity is enforced through appropriate constraints on the shared degrees of freedom at element interfaces. This combination retains the geometric flexibility of finite elements, since the domain can be decomposed into arbitrary quadrilateral (2D) or hexahedral (3D) patches, while achieving spectral convergence rates within each element when the solution is smooth [22, 23].


### 2.2.2. The Reference Element and Coordinate Mapping

A central concept in the spectral element framework is the reference element. All polynomial basis functions and quadrature rules are defined on a canonical domain, taken here as  $[-1, 1]^2$  in two dimensions. Each physical element in the mesh is related to this reference element through a coordinate mapping. For the rectangular axis-aligned meshes used through-

out this work, an affine mapping suffices; more general curvilinear elements would require a non-constant Jacobian, but that generalisation is not needed here and is therefore not pursued. The affine map reads:

$$x = \frac{b_x - a_x}{2} \xi + \frac{b_x + a_x}{2}, \quad y = \frac{b_y - a_y}{2} \eta + \frac{b_y + a_y}{2}, \quad (2.3)$$

where  $(a_x, b_x)$  and  $(a_y, b_y)$  are the physical extents of the element and  $(\xi, \eta) \in [-1, 1]^2$  are the reference coordinates. The Jacobian of this mapping captures the scaling and orientation between reference and physical space, and enters the transformation of basis functions, integrals, and differential operators from one coordinate system to the other. The specific form of this Jacobian and its role in the Piola transformation of vector-valued basis functions are discussed in Section 2.5.5.



**FIGURE 2.2** Mapping from the reference element  $\hat{\Omega} = [-1, 1]^2$  to a physical element  $\Omega_k = [a_x, b_x] \times [a_y, b_y]$  through an affine coordinate transformation  $F_k$ . All basis functions and quadrature rules are defined on  $\hat{\Omega}$  and mapped to  $\Omega_k$  via the Jacobian of  $F_k$ .

Because all basis functions and quadrature weights are defined on the reference element, they need to be computed only once and can be reused in all elements, regardless of their size or position in the physical domain. This is a major computational advantage of the spectral element approach [23, 24].

### 2.2.3. $h$ -Refinement and $p$ -Refinement

Spectral element methods improve accuracy in two ways. The first,  $h$ -refinement, reduces the element size at fixed polynomial degree and produces algebraic convergence at rate  $\mathcal{O}(h^{p+1})$  in the  $L^2$ -norm for sufficiently smooth solutions [25, 26]. The second,  $p$ -refinement, increases the polynomial degree at fixed mesh and produces exponential (spectral) convergence  $\mathcal{O}(e^{-cp})$  for analytic solutions, with  $c > 0$  depending on the smoothness of the solution [21].

Both strategies are studied in this thesis, where the convergence behaviour under  $h$ - and  $p$ -refinement is the central diagnostic used to verify the implementation. On the verification side, the manufactured solution is analytic, so  $p$ -refinement is expected to exhibit spectral convergence until the error saturates at a level determined by the data density. The experimental validation data, in contrast, are scattered samples corrupted by measurement noise and are not pointwise smooth, so spectral convergence under  $p$ -refinement is not expected on the experimental side. There the role of  $h$ - and  $p$ -refinement is to confirm that each reconstructed quantity approaches the measurement-noise floor at the algebraic rate predicted by the error analysis of Section 2.6., and reported in Chapter 5.

The polynomial basis functions and quadrature rules used to construct the spectral element approximation on the reference element are described in the next section.

### 2.3. Polynomial Basis Functions and Quadrature Rules

The constrained least-squares mimetic method requires polynomial basis functions defined on a reference element  $[-1, 1]$ . The following subsections define the one-dimensional building blocks from which all two-dimensional operators in the mimetic framework are constructed via tensor products.

#### 2.3.1. Legendre Polynomials and Derivatives

Legendre polynomials  $P_n(x)$  form a classical  $L^2$ -orthogonal family on  $[-1, 1]$ , satisfying

$$\int_{-1}^1 P_n(x) P_m(x) dx = 0 \quad \text{for } n \neq m. \quad (2.4)$$

The use of an  $L^2$ -orthogonal family here is not arbitrary: orthogonality with respect to the same inner product that defines the function spaces of Section 2.4.1. is what makes mass matrices well-conditioned and lets Gauss–Legendre quadrature integrate the relevant inner products exactly up to a known polynomial degree (Section 2.3.2.), so coupling between modes is contributed only by the discrete differential operators rather than by the basis itself. The Legendre polynomials satisfy:

$$\frac{d}{dx} \left[ (1 - x^2) \frac{dP_n(x)}{dx} \right] + n(n + 1) P_n(x) = 0. \quad (2.5)$$

The first and second derivatives are also needed:

$$P'_n(x) = \frac{d}{dx} P_n(x), \quad (2.6)$$

$$P''_n(x) = \frac{d^2}{dx^2} P_n(x), \quad (2.7)$$

which are essential for computing quadrature nodes and basis-function derivatives [27].

2

### 2.3.2. Gauss-Legendre and Gauss-Lobatto-Legendre Quadrature

Numerical integration over  $[-1, 1]$  is performed using Gaussian quadrature rules derived from Legendre polynomials, which enable exact integration of polynomials up to a prescribed degree. Two rules are used: Gauss-Legendre (GL) and Gauss-Lobatto-Legendre (GLL), both highly efficient for smooth integrands [28].

#### 2.3.2.1. Gauss-Legendre Quadrature

This rule employs  $p$  interior nodes and is exact for polynomials of degree up to  $2p - 1$ :

$$\int_{-1}^1 f(x) dx \approx \sum_{i=1}^p w_i f(x_i), \quad (2.8)$$

where

$$P_p(x_i) = 0, \quad w_i = \frac{2}{(1 - x_i^2) [P'_p(x_i)]^2}. \quad (2.9)$$

The accuracy and efficiency of Gauss quadrature for smooth integrands is discussed in [28].

#### 2.3.2.2. Gauss-Lobatto-Legendre Quadrature

Including the interval endpoints, this rule uses  $p + 1$  points and is exact up to degree  $2p - 1$ . Nodes satisfy:

$$(1 - x^2) P'_p(x) = 0 \quad \Rightarrow \quad x_0 = -1, x_p = 1, \quad (2.10)$$

with interior nodes from  $P'_p(x) = 0$ . Weights are:

$$w_i = \frac{2}{p(p+1) [P_p(x_i)]^2}, \quad i = 0, \dots, p. \quad (2.11)$$

Gauss-Lobatto-Legendre quadrature is widely used in spectral element methods for two reasons [29]. First, the inclusion of the endpoints  $x = \pm 1$  is essential for inter-element continuity: because neighbouring elements

share boundary nodes, degrees of freedom can be matched directly at interfaces without additional interpolation. Second, when the mass matrix of the nodal (Lagrange) basis is assembled using GLL quadrature at the same nodes as the interpolation, the result is a diagonal matrix, which greatly reduces the computational cost of its inversion.

### 2.3.2.3. Node Distributions

Figure 2.3 illustrates the one-dimensional GL and GLL node distributions for polynomial order  $p = 3$ . The  $N = p + 1 = 4$  GLL nodes include the endpoints  $\xi = \pm 1$  and serve as the interpolation points for the Lagrange basis and as the endpoints of the subintervals on which the edge basis is defined. The  $p = 3$  GL nodes lie strictly in the interior of  $[-1, 1]$ . They are introduced here only as quadrature points for evaluating integrals where Gauss-Legendre integration is used. No separate dual mesh is constructed.


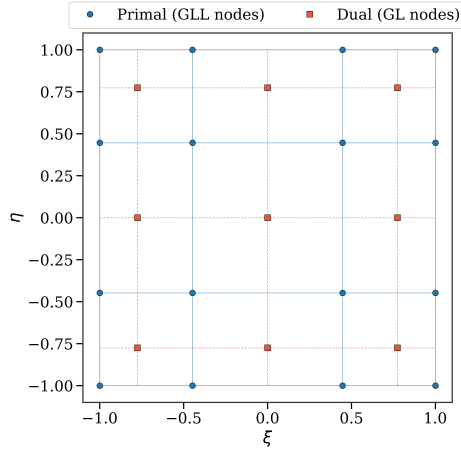


FIGURE 2.3 One-dimensional GLL and GL node distributions for  $p = 3$ .

In two dimensions, the node distributions are formed as tensor products of the one-dimensional GLL and GL nodes. Figure 2.4 shows the resulting  $4 \times 4$  grid of GLL nodes and the  $3 \times 3$  grid of GL nodes on the reference element  $[-1, 1]^2$ . The GLL grid carries the nodal degrees of freedom, the edge fluxes along the GLL edges, and the cell-integrated values on the cells bounded by those edges. The GL grid is shown only to indicate the locations of the quadrature points used to evaluate integrals. The grid lines connecting the nodes illustrate the tensor-product structure that allows all two-dimensional operators to be constructed from one-dimensional building blocks.

### 2.3.3. Lagrange Basis Functions

Nodal interpolation on the GLL nodes is performed using Lagrange basis functions  $h_i(\xi)$ , constructed at nodes  $\{\xi_0, \dots, \xi_N\}$ . They satisfy the



**FIGURE 2.4** Two-dimensional tensor-product GLL and GL node distributions on  $[-1, 1]^2$  for  $p = 3$ .

Kronecker-delta property:

$$h_i(\xi_j) = \delta_{ij}, \quad (2.12)$$

and are defined by:

$$h_i(\xi) = \prod_{\substack{j=0 \\ j \neq i}}^N \frac{\xi - \xi_j}{\xi_i - \xi_j}. \quad (2.13)$$

These basis functions are used to represent scalar fields in spectral elements [30].

In this work, the Lagrange basis is used to represent pointwise scalar quantities such as the streamfunction and the vorticity, whose degrees of freedom are values at nodes. Any scalar field  $f(\xi)$  is approximated as  $f_h(\xi) = \sum_{i=0}^N f_i h_i(\xi)$ , where  $f_i = f(\xi_i)$  are the nodal values.

Figure 2.5 shows the four Lagrange basis functions  $h_0(\xi)$  through  $h_3(\xi)$  for polynomial order  $p = 3$ , constructed on the GLL nodes  $\{\xi_0, \xi_1, \xi_2, \xi_3\}$ . Each basis function equals unity at its associated node and vanishes at all other nodes, in accordance with the Kronecker-delta property (2.12). The black dots mark the GLL node locations; the coloured dots at  $h_i(\xi_i) = 1$  confirm the interpolation property. The basis functions associated with the interior nodes ( $h_1$  and  $h_2$ ) are smooth and mirrored along the  $\xi = 0$  line. Those at the boundary nodes ( $h_0$  and  $h_3$ ) are asymmetric and exhibit larger excursions in the element interior, a characteristic feature of endpoint interpolation on non-uniformly spaced nodes.




FIGURE 2.5 Lagrange basis functions  $h_i(\xi)$  on GLL nodes for  $p = 3$ .

### 2.3.4. Edge Basis Functions

To capture flux-like properties in the mimetic framework, edge basis functions are used. The edge basis functions  $e_i(\xi)$  are derived from the derivatives of the Lagrange polynomials. The derivative of  $h_j(\xi)$  is:

$$h_j'(\xi) = \frac{d}{d\xi} h_j(\xi), \quad (2.14)$$

from which the edge basis functions are defined as:

$$e_i(\xi) = -\sum_{j=0}^{i-1} h_j'(\xi), \quad i = 1, \dots, N. \quad (2.15)$$

By construction, the edge basis [8, 31] forms a partition of unity,

$$\sum_{i=1}^N e_i(\xi) = 1, \quad \xi \in [-1, 1], \quad (2.16)$$

and satisfy the integral identity

$$\int_{\xi_{j-1}}^{\xi_j} e_i(\xi) d\xi = \delta_{ij}, \quad (2.17)$$

so the degree of freedom associated with  $e_i$  is the integral of the represented function over the subinterval  $[\xi_{i-1}, \xi_i]$ . The edge degrees of freedom are integral values along subintervals, not fluxes in the 2D/3D sense; the flux interpretation arises only in the 2D/3D outer-oriented setting of Section 2.4.2.. Together, equations (2.16)–(2.17) make the discrete derivative the exact analogue of the fundamental theorem of calculus,

$$\int_{\xi_{j-1}}^{\xi_j} \sum_{i=1}^N (f_i - f_{i-1}) e_i(\xi) d\xi = f_j - f_{j-1}, \quad (2.18)$$

for any nodal coefficient vector  $\mathbf{f} = (f_0, \dots, f_N)^T$ . The local discrete derivative is therefore exact, and summing over all subintervals gives  $f_N - f_0$  exactly with no internal residual.

Whereas the Lagrange basis functions represent pointwise values at nodes, the edge basis functions represent integrated quantities over the intervals between nodes. In the discretisation built in Section 2.5., the edge basis functions will discretise quantities whose degrees of freedom are integral values along subintervals, for instance the components of the velocity field that appear in the integral form of the continuity equation.

Figure 2.6 shows the three edge basis functions  $e_1(\xi)$ ,  $e_2(\xi)$ , and  $e_3(\xi)$  for polynomial order  $p = 3$ . Each edge function is concentrated over its associated subinterval  $[\xi_{i-1}, \xi_i]$ , indicated by the shaded regions. The integral of each edge function over its own subinterval equals unity (2.17), while the integral over any other subinterval vanishes. The black dots mark the GLL node positions that define the subinterval boundaries. The two-dimensional tensor products of edge and Lagrange functions that arise in the discretisation are constructed in Section 2.5.3..




FIGURE 2.6 Edge basis functions  $e_i(\xi)$  on GLL nodes for  $p = 3$ .

## 2.4. Function Spaces and the de Rham Complex

The Mimetic Spectral Element Method starts from a choice of de Rham complex, a sequence of Sobolev spaces linked by differential operators, and constructs discrete approximation spaces that preserve the algebraic structure of the complex exactly at the discrete level. The complex relevant to (but not limited to) two-dimensional incompressible-flow reconstruction consists of the Sobolev spaces  $H(\text{curl}; \Omega)$ ,  $H(\text{div}; \Omega)$ , and  $L^2(\Omega)$ ,

whose regularity and continuity conditions match the reconstruction problem addressed in Chapter 3. This section introduces those Sobolev spaces, defines the continuous de Rham complex that links them, and draws the consequences for the reconstruction problem.

## 2

### 2.4.1. Sobolev Spaces

Three function spaces appear in the de Rham complex on a bounded domain  $\Omega \subset \mathbb{R}^2$  with Lipschitz boundary  $\partial\Omega$  [32]. The Lipschitz assumption ensures that trace operators on  $\partial\Omega$  are well-defined. The base space is  $L^2(\Omega)$ , the space of square-integrable real functions on  $\Omega$ ,

$$L^2(\Omega) = \left\{ f : \Omega \rightarrow \mathbb{R} \mid \int_{\Omega} |f|^2 \, d\Omega < \infty \right\}, \quad (2.19)$$

$$(f, g)_{L^2} = \int_{\Omega} f g \, d\Omega, \quad \|f\|_{L^2}^2 = (f, f)_{L^2}.$$

The two function spaces below,  $H(\text{curl}; \Omega)$  and  $H(\text{div}; \Omega)$ , are then defined as  $L^2$ -functions whose specified differential operator also lies in  $L^2$ .

The space  $H(\text{curl}; \Omega)$  consists of square-integrable scalar fields (in two dimensions) whose curl is also square-integrable:

$$H(\text{curl}; \Omega) = \{q \in L^2(\Omega) \mid \nabla \times q \in [L^2(\Omega)]^2\}, \quad (2.20)$$

$$\|q\|_{H(\text{curl}; \Omega)}^2 = \|q\|_{L^2}^2 + \|\nabla \times q\|_{L^2}^2.$$

Here, the two-dimensional curl of a scalar field  $q$  is defined as  $\nabla \times q = (-\partial_y q, \partial_x q)^T$ , which produces a vector field. In three dimensions, the  $H(\text{curl}; \Omega)$  space contains vector fields; in this two-dimensional setting, it contains fields that point solely in the out-of-plane direction. For example, the vorticity vector is  $\boldsymbol{\omega} = (0, 0, \omega)^T$ , and we work with the single out-of-plane component  $\omega$  as a scalar whose curl (i.e. perpendicular gradient) is square-integrable. The same applies to the streamfunction  $\psi$ , which is also a scalar in 2D. Functions in  $H(\text{curl}; \Omega)$  are continuous across element interfaces: their values must match at inter-element boundaries.

The space  $H(\text{div}; \Omega)$  consists of square-integrable vector fields whose divergence is also square-integrable:

$$H(\text{div}; \Omega) = \{\mathbf{v} \in [L^2(\Omega)]^2 \mid \nabla \cdot \mathbf{v} \in L^2(\Omega)\}, \quad (2.21)$$

$$\|\mathbf{v}\|_{H(\text{div}; \Omega)}^2 = \|\mathbf{v}\|_{L^2}^2 + \|\nabla \cdot \mathbf{v}\|_{L^2}^2.$$

Functions in  $H(\text{div}; \Omega)$  have continuous normal components across element interfaces, but their tangential components may be discontinuous.

This is the appropriate regularity for the velocity field  $\mathbf{u}$ : the normal flux must be continuous to define the divergence element-by-element, while the tangential component is left free.

The pressure of the formulation lives in  $L^2(\Omega)$ . It is computed as a cell-averaged quantity on each element, and neighbouring elements are allowed to have different pressure values at their shared interface. The pressure also acts as a Lagrange multiplier for the divergence-free constraint, but its  $L^2$  regularity (rather than continuity across elements) follows from the choice of placement in the de Rham complex, not from the Lagrange-multiplier role itself.

The key distinction among these spaces is the inter-element continuity they enforce. Table 2.1 summarises this relationship in two dimensions: as one moves from  $H(\text{curl}; \Omega)$  to  $L^2$ , progressively less continuity is required across element interfaces.

**TABLE 2.1** Two-dimensional function spaces and their inter-element continuity properties.

Space	Interface continuity	Physical quantity
$H(\text{curl}; \Omega)$	Full (value)	$\psi, \omega$
$H(\text{div}; \Omega)$	Normal component	$\mathbf{u}$
$L^2$	None	P

### 2.4.2. Differential Forms in Two Dimensions

Differential forms provide a unified language for scalars, vectors, and integrated quantities, and they underpin the mimetic discretisation of Section 2.5. [9, 33]. In two dimensions there are three kinds, distinguished by the geometric object each is paired with.

- A **0-form** is paired with a point. It is a scalar field, with pointwise values as degrees of freedom.
- A **1-form** is paired with a curve. Its degrees of freedom are line integrals along curves, and in 2D it can be identified with a vector field via the standard Euclidean inner product.
- A **2-form** is paired with a surface. Its degrees of freedom are integrals over oriented areas, and in 2D it can be identified with a scalar density at cells.

Table 2.2 lists which Sobolev space each form lives in and which physical quantity each represents in this thesis. The discrete approximation spaces built in Section 2.5. reproduce the same geometric pairing at the discrete level.

2

**TABLE 2.2** Correspondence between differential forms, function spaces, and the physical quantities they represent in this work.

Form	Geometric pairing	Function space	Physical quantity
0-form	node (point value)	$H(\text{curl}; \Omega)$	$\psi, \omega$
1-form	edge (line integral)	$H(\text{div}; \Omega)$	$\mathbf{u}$
2-form	cell (area integral)	$L^2(\Omega)$	$\tilde{P}, \nabla \cdot \mathbf{u}$

### 2.4.3. The Continuous de Rham Complex

The function spaces defined in Section 2.4.1. are connected by differential operators that form the  $L^2$  de Rham complex [33]. In two dimensions, the primal complex reads:

$$\mathbb{R} \hookrightarrow H(\text{curl}; \Omega) \xrightarrow{\nabla \times} H(\text{div}; \Omega) \xrightarrow{\nabla \cdot} L^2 \longrightarrow 0, \quad (2.22)$$

where the first arrow indicates that constants form the null space of the curl on  $H(\text{curl}; \Omega)$  and the trailing zero terminates the sequence in 2D. On a contractible domain (such as the rectangles used in this thesis), the complex is exact, meaning that the kernel of each operator equals the image of the previous one. The range of the divergence operator itself depends on the boundary conditions.

The primal complex (2.22) admits the strong curl of a scalar and the strong divergence of a vector field directly, because  $H(\text{curl}; \Omega)$  and  $H(\text{div}; \Omega)$  have sufficient regularity for those operations to produce square-integrable results. The reverse-direction operators, the curl of a vector field (needed for vorticity) and the gradient of a scalar (needed for pressure), cannot be applied directly to primal-space fields. The Hodge-dual complex,

$$0 \longleftarrow H(\text{curl}; \Omega) \xleftarrow{(\nabla \times)^*} H(\text{div}; \Omega) \xleftarrow{(\nabla \cdot)^*} L^2(\Omega) \longleftarrow \mathbb{R}, \quad (2.23)$$

encodes these weak operators. The Hodge-conjugated arrows are the  $L^2$ -adjoints of the primal operators, obtained by integration by parts. The first dual operator is

$$(\nabla \cdot)^* = -\nabla, \quad (2.24)$$

the negative gradient acting from  $L^2(\Omega)$  on a test function in  $H(\text{div}; \Omega)$ . The second dual operator is

$$(\nabla \times)^* = \text{rot}, \quad \text{rot}(\mathbf{u}) = \partial_x v - \partial_y u, \quad (2.25)$$

the scalar curl of a vector field, acting from  $H(\text{div}; \Omega)$  on a test function in  $H(\text{curl}; \Omega)$ . Each Hodge-conjugated arrow is weak in the sense that it is defined through the integration-by-parts identity rather than by point-wise application, and the boundary integral produced by the integration by parts is part of the operator.

This complex is better suited for the problems considered in this thesis, because the divergence-free condition  $\nabla \cdot \mathbf{u} = 0$  is a strong constraint that must be enforced exactly. Placing the velocity in  $H(\text{div}; \Omega)$  and using the strong divergence operator allows the constraint to be enforced exactly as an algebraic equation. Vorticity and pressure, derived quantities, are recovered afterwards through the weak operators of the Hodge-dual complex.

The primal complex (2.22) has two fundamental properties [34]. The first is the complex property: the composition of two consecutive operators in the sequence vanishes identically,

$$\nabla \cdot (\nabla \times \psi) = 0 \quad \forall \psi \in H(\text{curl}; \Omega). \quad (2.26)$$

This can be verified directly:  $\nabla \cdot (\nabla \times \psi) = \nabla \cdot (-\partial_y \psi, \partial_x \psi)^T = -\partial_{xy} \psi + \partial_{xy} \psi = 0$ . Physically, this means that any velocity field derived from a streamfunction is automatically divergence-free.

The second property is exactness: on a contractable domain, every divergence-free vector field  $\mathbf{u} \in H(\text{div}; \Omega)$  can be written as  $\mathbf{u} = \nabla \times \psi$  for some streamfunction  $\psi \in H(\text{curl}; \Omega)$ . This guarantees that enforcing the discrete divergence-free condition is equivalent to requiring the velocity to be the curl of a streamfunction.

#### 2.4.4. Consequences for the Reconstruction Problem

The choice  $\mathbf{u} \in H(\text{div}; \Omega)$  has three direct consequences for the reconstruction method developed in Chapter 3.

First, the inter-element continuity condition enforced on the velocity is normal flux continuity. At a shared interface between two elements, the normal component of the velocity from one element must equal the normal component from the neighbouring element. The tangential component, however, is not required to match. This is the minimum regularity needed to define  $\nabla \cdot \mathbf{u}$  element-by-element and to enforce the divergence-free constraint exactly at the discrete level (Section 2.5.2.).

Second, the strong curl of  $\mathbf{u}$  is not well-defined when  $\mathbf{u} \in H(\text{div}; \Omega)$ , because the tangential component may be discontinuous across element in-

interfaces and the strong curl is not guaranteed to be square-integrable. Consequently, the vorticity  $\omega = \nabla \times \mathbf{u}$  cannot be computed by direct differentiation. Instead, it must be recovered through a weak formulation obtained by integration by parts: find  $\omega \in H(\text{curl}; \Omega)$  such that

$$\int_{\Omega} \omega q \, d\Omega = - \int_{\Omega} \mathbf{u} \cdot (\nabla \times q) \, d\Omega + \oint_{\partial\Omega} (\mathbf{u} \cdot \hat{\mathbf{t}}) q \, ds, \quad \forall q \in H(\text{curl}; \Omega), \quad (2.27)$$

where  $\hat{\mathbf{t}}$  is the unit tangent vector on  $\partial\Omega$ . Integration by parts transfers the curl from  $\mathbf{u}$ , which lacks the required regularity, onto the test function  $q \in H(\text{curl}; \Omega)$ , whose curl is square-integrable by definition. The boundary integral that this produces accounts for the tangential trace of  $\mathbf{u}$  on  $\partial\Omega$ .

Third, the pressure  $\tilde{P} \in L^2(\Omega)$  has no strong gradient, so the gradient acts on test functions  $\mathbf{v} \in H(\text{div}; \Omega)$  only via integration by parts,

$$\int_{\Omega} \nabla \tilde{P} \cdot \mathbf{v} \, d\Omega = - \int_{\Omega} \tilde{P} \nabla \cdot \mathbf{v} \, d\Omega + \oint_{\partial\Omega} \tilde{P} \mathbf{v} \cdot \mathbf{n} \, ds.$$

The pressure is recovered from the momentum equation in this weak form.

In summary, the de Rham complex determines the regularity of each field and the weak formulations used to recover the derived quantities. The concrete reconstruction procedures are developed in Chapter 3.

## 2.5. Mimetic Spectral Element Discretization

This section shows how the continuous function spaces and operators from Section 2.4. are discretized using the polynomial building blocks from Section 2.3.. The construction proceeds in three steps. First, each space in the de Rham complex is approximated by a finite-dimensional subspace built from tensor products of the one-dimensional Lagrange and edge basis functions (Section 2.5.1.). Second, the differential operators between these spaces are represented by sparse incidence matrices that depend only on the connectivity of the basis (Section 2.5.2.). Third, the continuous de Rham complex is shown to carry over to the discrete spaces as an exact algebraic identity (Section 2.5.3.). Inner products between the discrete spaces are encoded by mass matrices (Section 2.5.4.), which provide the metric information needed to express weak operators via integration by parts.

### 2.5.1. Discrete Approximation Spaces

One family of basis functions is associated with each form degree introduced in Section 2.4.2..

### 2.5.1.1. 0-Form Basis Functions

The discrete subspace of  $H(\text{curl}; \Omega)$  is spanned by tensor products of Lagrange basis functions:

$$\Psi_{ij}^{(0)}(\xi, \eta) = h_i(\xi) h_j(\eta), \quad i = 0, \dots, p, \quad j = 0, \dots, p. \quad (2.28)$$

This space has dimension  $(p + 1)^2$ . Each basis function is associated with a node in the  $N \times N$  GLL grid, and the corresponding degrees of freedom are pointwise values. Because the Lagrange basis functions include the endpoints  $\xi = \pm 1$ , neighbouring elements share nodes on their common boundary, which enforces the full continuity required by  $H(\text{curl}; \Omega)$ . In this work, the 0-form basis functions represent the streamfunction  $\psi$  and the vorticity  $\omega$ .

### 2.5.1.2. 1-Form Basis Functions

The discrete subspace of  $H(\text{div}; \Omega)$  is spanned by two families of tensor products, one for each velocity component:

$$\Psi_{\xi, ij}^{(1)}(\xi, \eta) = h_i(\xi) e_j(\eta), \quad i = 0, \dots, p, \quad j = 1, \dots, p, \quad (2.29)$$

$$\Psi_{\eta, ij}^{(1)}(\xi, \eta) = e_i(\xi) h_j(\eta), \quad i = 1, \dots, p, \quad j = 0, \dots, p. \quad (2.30)$$

The  $\xi$ -component uses Lagrange in  $\xi$  (continuous across vertical interfaces) and edge in  $\eta$ , while the  $\eta$ -component uses edge in  $\xi$  and Lagrange in  $\eta$  (continuous across horizontal interfaces). This construction automatically enforces normal flux continuity across element interfaces: the Lagrange direction is continuous at the boundary, while the edge direction is not. The total number of 1-form degrees of freedom per element is  $2p(p + 1)$ . The 1-form basis functions represent the velocity  $\mathbf{u}$ .

### 2.5.1.3. 2-Form Basis Functions

The discrete subspace of  $L^2$  is spanned by tensor products of edge basis functions:

$$\Psi_{ij}^{(2)}(\xi, \eta) = e_i(\xi) e_j(\eta), \quad i = 1, \dots, p, \quad j = 1, \dots, p. \quad (2.31)$$

This space has dimension  $p^2$ . Since edge basis functions do not include the endpoints, no inter-element continuity is enforced, consistent with  $L^2$ . The degrees of freedom represent cell-integrated values. The 2-form basis functions represent the pressure  $P$  and the divergence  $\nabla \cdot \mathbf{u}$ .

## 2.5.2. Incidence Matrices

The incidence matrices encode the topological relations between the discrete spaces. They are the discrete analogues of the differential operators in the de Rham complex and contain only entries from  $\{-1, 0, +1\}$ , independent of the element geometry [31].

### 2.5.2.1. One-Dimensional Incidence Matrix

In one dimension, the incidence matrix  $E_{1D}$  maps nodal values (0-form degrees of freedom) to edge values (1-form degrees of freedom) by taking consecutive differences. For polynomial order  $p$  with  $N = p + 1$  nodes,  $E_{1D}$  is a  $p \times (p + 1)$  matrix:

$$(E_{1D})_{ij} = \begin{cases} -1 & \text{if } j = i, \\ +1 & \text{if } j = i + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (2.32)$$

For  $p = 3$ , this gives:

$$E_{1D} = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{pmatrix}. \quad (2.33)$$

The action  $E_{1D} \mathbf{f} = (f_1 - f_0, f_2 - f_1, f_3 - f_2)^T$  computes the discrete derivative: the difference of nodal values across each edge. This is the discrete analogue of the fundamental theorem of calculus:  $\int_{\xi_i}^{\xi_{i+1}} f'(\xi) d\xi = f(\xi_{i+1}) - f(\xi_i)$ .

### 2.5.2.2. Two-Dimensional Incidence Matrices

In two dimensions, the incidence matrices are constructed from tensor products of  $E_{1D}$  and the identity matrix  $I$  [31].

The matrix  $E^{1,0}$  is the discrete curl, mapping 0-forms to 1-forms. It acts on the  $(p + 1)^2$  nodal coefficients and produces the  $2p(p + 1)$  edge coefficients:

$$E^{1,0} = \begin{pmatrix} I_{p+1} \otimes E_{1D} \\ E_{1D} \otimes I_{p+1} \end{pmatrix}, \quad (2.34)$$

where  $I_{p+1}$  is the  $(p + 1) \times (p + 1)$  identity matrix. The top block computes differences in the  $\eta$ -direction (producing  $\xi$ -edge fluxes) and the bottom block computes differences in the  $\xi$ -direction (producing  $\eta$ -edge fluxes). This corresponds to the discrete version of  $\mathbf{u} = \nabla \times \psi = (-\partial_\eta \psi, \partial_\xi \psi)^T$ .

The matrix  $E^{2,1}$  is the discrete divergence, mapping 1-forms to 2-forms. It acts on the  $2p(p+1)$  edge coefficients and produces the  $p^2$  cell coefficients:

$$E^{2,1} = \begin{pmatrix} E_{1D} \otimes I_p & -I_p \otimes E_{1D} \end{pmatrix}, \quad (2.35)$$

where  $I_p$  is the  $p \times p$  identity matrix. This computes the net flux through the edges of each cell, which is the discrete divergence.

The fundamental property  $E^{2,1}E^{1,0} = 0$  can be verified by direct substitution:

$$\begin{aligned} E^{2,1}E^{1,0} &= (E_{1D} \otimes I_p)(I_{p+1} \otimes E_{1D}) - (I_p \otimes E_{1D})(E_{1D} \otimes I_{p+1}) \\ &= E_{1D} \otimes E_{1D} - E_{1D} \otimes E_{1D} = 0. \end{aligned}$$

This identity is exact and purely algebraic: it holds for any  $p$  and does not depend on the node positions, quadrature weights, or element geometry. Its significance is taken up in the next subsection.

### 2.5.3. The Discrete de Rham Complex

With the discrete approximation spaces from Section 2.5.1. and the incidence matrices from Section 2.5.2. in place, the continuous de Rham complex of Section 2.4.3. carries over to the discrete setting. Denoting by  $\mathcal{C}^0 \subset H(\text{curl}; \Omega)$ ,  $\mathcal{C}^1 \subset H(\text{div}; \Omega)$ , and  $\mathcal{C}^2 \subset L^2(\Omega)$  the discrete subspaces of dimensions  $(p+1)^2$ ,  $2p(p+1)$ , and  $p^2$  respectively, the discrete primal complex reads

$$\mathcal{C}^0 \xrightarrow{E^{1,0}} \mathcal{C}^1 \xrightarrow{E^{2,1}} \mathcal{C}^2, \quad (2.36)$$

where  $E^{1,0}$  is the discrete curl that maps 0-form coefficients to 1-form coefficients, and  $E^{2,1}$  is the discrete divergence that maps 1-form coefficients to 2-form coefficients. The complex property

$$E^{2,1}E^{1,0} = 0 \quad (2.37)$$

holds as an exact algebraic identity, which is the discrete analogue of  $\nabla \cdot (\nabla \times \psi) = 0$ . It ensures that any velocity coefficients derived from a streamfunction through  $\mathbf{u} = E^{1,0}\psi$  are exactly divergence-free in the sense that  $E^{2,1}\mathbf{u} = 0$ , and conversely that the divergence-free constraint  $E^{2,1}\mathbf{u} = 0$  can be enforced algebraically without any quadrature approximation.

The corresponding dual discrete complex reads

$$(\mathcal{C}^2)^* \xleftarrow{(E^{1,0})^T} (\mathcal{C}^1)^* \xleftarrow{-(E^{2,1})^T} (\mathcal{C}^0)^*, \quad (2.38)$$

with transposed incidence matrices acting on the dual (Riesz-paired) coefficient vectors. The negative sign on  $-(E^{2,1})^T$  reflects that the  $L^2$ -adjoint of the divergence operator carries a sign through integration by parts (it corresponds to the negative gradient). No separate dual mesh is constructed; the dual coefficients  $(\mathcal{C}^k)^*$  are algebraic companions of the primal ones, and the weak operators of Chapter 3 are assembled from  $E^{1,0}$ ,  $E^{2,1}$ , their transposes, and the mass matrices introduced next, all acting on the primal degrees of freedom.

#### 2.5.4. Mass Matrices as Hodge Operators

Whereas the incidence matrices encode topology, the mass matrices encode metric information (geometry and inner products). They establish the correspondence between the primal coefficients and their dual (Riesz-paired) representations, and play the role of discrete Hodge operators in the mimetic framework [31, 35]. For two discrete fields  $\mathbf{a}, \mathbf{b} \in \mathcal{C}^k$  with coefficient vectors  $\mathbf{a}$  and  $\mathbf{b}$ , the  $L^2$  inner product is  $\mathbf{a}^T M^k \mathbf{b}$ , and the dual representation of a primal coefficient vector  $\mathbf{a}$  is  $\tilde{\mathbf{a}} = M^k \mathbf{a}$ . The three mass matrices are defined as

$$M_{ij}^0 = \int_{-1}^1 \int_{-1}^1 \Psi_i^{(0)}(\xi, \eta) \Psi_j^{(0)}(\xi, \eta) d\xi d\eta, \quad (2.39)$$

where  $\Psi_i^{(0)} = h_k(\xi) h_l(\eta)$  with a suitable ordering of the double index  $(k, l)$  into a single index  $i$ ;


$$M_{ij}^1 = \int_{-1}^1 \int_{-1}^1 \Psi_i^{(1)}(\xi, \eta) \cdot \Psi_j^{(1)}(\xi, \eta) d\xi d\eta, \quad (2.40)$$

where  $\Psi_i^{(1)}$  is one of the 1-form basis functions from (2.29)–(2.30) (the tensor-product structure makes  $M^1$  block-diagonal with separate blocks for the  $\xi$ - and  $\eta$ -components); and

$$M_{ij}^2 = \int_{-1}^1 \int_{-1}^1 \Psi_i^{(2)}(\xi, \eta) \Psi_j^{(2)}(\xi, \eta) d\xi d\eta. \quad (2.41)$$

Whereas  $E^{1,0}$  and  $E^{2,1}$  are purely topological and have entries only in  $\{-1, 0, +1\}$ , the mass matrices depend on the element geometry through the choice of basis and on the inner product. This separation between topological and metric information is a defining feature of the mimetic approach: identities such as  $E^{2,1} E^{1,0} = 0$  hold exactly regardless of the element shape, while the geometric content is confined to the mass matrices and adapts to the element through the Piola transformations of Section 2.5.5.

The Hodge operators  $M^0$ ,  $M^1$ , and  $M^2$  together with the incidence matrices  $E^{1,0}$  and  $E^{2,1}$  form the algebraic structure of the discrete de Rham complex. Figure 2.7 summarises this structure: the primal incidence matrices act left to right between  $\mathcal{C}^0$ ,  $\mathcal{C}^1$ ,  $\mathcal{C}^2$ , the mass matrices link primal coefficients to their dual representations, and the transposed incidence matrices act on the dual side. Two practical points are worth noting. First, the dual representations  $(\mathcal{C}^k)^*$  are algebraic Riesz-paired companions of the primal coefficients, not a separately discretised dual mesh. Second, some adjoint operators differ from the corresponding vector-calculus operator by a sign: for example, the  $L^2$ -adjoint of the divergence is the negative gradient,  $(\nabla \cdot)^* = -\nabla$ , and at the discrete level this sign carries over to the corresponding combination of  $(E^{2,1})^T$  and  $M^2$ .



**FIGURE 2.7** The discrete de Rham complex with Hodge operators. The primal complex (top) is connected to the dual complex (bottom) by the mass matrices  $M^0$ ,  $M^1$ , and  $M^2$  (downward arrows) and their inverses  $(M^0)^{-1}$ ,  $(M^1)^{-1}$ , and  $(M^2)^{-1}$  (upward arrows). The primal incidence matrices  $E^{1,0}$  (discrete curl) and  $E^{2,1}$  (discrete divergence) act left to right; the transposed incidence matrices  $(E^{1,0})^T$  (discrete gradient) and  $-(E^{2,1})^T$  (discrete curl of a vector) act right to left.

### 2.5.5. Piola Transformations for $k$ -Forms

The basis functions and operators defined in the previous subsections are formulated on the reference element  $[-1, 1]^2$ . To map them to a physical element with coordinates  $(x, y)$ , Piola transformations are used. These transformations ensure that the inter-element continuity properties and integral identities of each  $k$ -form are preserved under the coordinate mapping [27].

Let  $J$  denote the Jacobian matrix of the mapping from reference to physical coordinates,

$$J = \begin{pmatrix} \partial x / \partial \xi & \partial x / \partial \eta \\ \partial y / \partial \xi & \partial y / \partial \eta \end{pmatrix}, \quad (2.42)$$

with determinant  $\det J$ . In this work, we restrict to affine (linear) maps for which  $J$  is constant over each element. The Piola transformations are valid

for general diffeomorphisms, but the expressions simplify considerably for affine maps.

### 2.5.5.1. 0-Forms

Scalar fields require no transformation beyond the coordinate substitution:

$$\Psi^{(0)}(x, y) = \Psi^{(0)}(\xi(x, y), \eta(x, y)). \quad (2.43)$$

Pointwise values are preserved: if  $\psi$  takes the value  $\psi_i$  at node  $(\xi_i, \eta_j)$  on the reference element, it takes the same value at the corresponding physical node.

### 2.5.5.2. 1-Forms

1-form basis functions are mapped using the contravariant Piola transformation:

$$\Psi^{(1)}(x, y) = \frac{1}{\det J} J \Psi^{(1)}(\xi(x, y), \eta(x, y)). \quad (2.44)$$

This transformation preserves normal flux continuity across element interfaces: if two adjacent elements share an edge, the normal component of  $\Psi^{(1)}$  computed from either element agrees at the interface. For affine maps with a diagonal Jacobian ( $J_{01} = J_{10} = 0$ ), this reduces to scaling each component independently:

$$\Psi_x^{(1)} = \frac{J_{00}}{\det J} \Psi_\xi^{(1)}, \quad \Psi_y^{(1)} = \frac{J_{11}}{\det J} \Psi_\eta^{(1)}. \quad (2.45)$$

### 2.5.5.3. 2-Forms

Scalar fields representing cell-integrated quantities are mapped by dividing by the determinant:

$$\Psi^{(2)}(x, y) = \frac{1}{\det J} \Psi^{(2)}(\xi(x, y), \eta(x, y)). \quad (2.46)$$

This ensures that the integral of a 2-form over a physical cell equals the integral of the reference 2-form over the corresponding reference cell, up to the area scaling encoded in  $\det J$ .

Under affine maps, the mass matrices transform to the physical domain as [31]:

$$M^0(x, y) = M^0(\xi, \eta) \det J, \quad (2.47)$$

$$M_x^1(x, y) = M^1(\xi, \eta) \frac{J_{00}^2}{\det J}, \quad M_y^1(x, y) = M^1(\xi, \eta) \frac{J_{11}^2}{\det J}, \quad (2.48)$$

$$M^2(x, y) = M^2(\xi, \eta) \frac{1}{\det J}. \quad (2.49)$$

The incidence matrices  $E^{1,0}$  and  $E^{2,1}$  are unchanged by the mapping, because they encode topology rather than geometry. This separation is what makes the mimetic framework robust: the topological identities (such as  $E^{2,1}E^{1,0} = 0$ ) hold on any mesh, while the metric information is confined to the mass matrices and adapts to the element shape through the Jacobian.

## 2.6. Error Analysis for Spectral Element Approximation

The mimetic spectral element discretization introduced in Section 2.5. approximates the continuous function spaces with finite-dimensional polynomial subspaces. This section establishes the theoretical convergence rates under  $p$ - and  $h$ -refinement that govern how the reconstruction error decreases as the polynomial order is increased or the mesh is refined. These rates provide the benchmarks against which the numerical results in Chapter 5 are measured.

### 2.6.1. Polynomial Approximation Theory

The starting point is the classical approximation result for spectral element methods. Let  $u$  be a function in the Sobolev space  $H^s(\Omega)$  for some regularity index  $s \geq 1$ , and let  $u_h$  denote its best approximation in a polynomial space of degree  $p$  on a mesh with element size  $h$ . Then the approximation error satisfies [25, 26]:

$$\|u - u_h\|_{H^m} \leq C h^{\min(p+1, s) - m} \|u\|_{H^s}, \quad (2.50)$$

where  $m$  is the order of the Sobolev norm in which the error is measured ( $m = 0$  for  $L^2$ ,  $m = 1$  for  $H^1$ , etc.),  $s$  is the Sobolev regularity of the exact solution (the number of square-integrable derivatives it possesses), and  $C$  is a constant independent of  $h$  and  $p$ . The rate is  $\min(p + 1, s) - m$  because it is limited by whichever is smaller: the polynomial order  $p$  or the regularity  $s$  of the solution. Each derivative included in the error norm ( $m$ ) costs one algebraic order.

A key tool for relating the convergence rates of derived quantities is the inverse estimate [26, 36]:

$$\|\nabla v_h\|_{L^2} \leq C \frac{p^2}{h} \|v_h\|_{L^2}, \quad (2.51)$$

which holds for any polynomial  $v_h$  of degree  $p$  in the discrete space. The factor  $p^2$  comes from a polynomial inverse estimate, the Markov brothers' inequality bounds the sup-norm of a derivative by a  $p^2$ -multiple of the sup-norm of the polynomial, and the corresponding  $L^2$  estimate carries

the same  $p^2$  scaling up to an inessential constant on compact intervals [26]. The factor  $h^{-1}$  comes from the scaling to a physical element of size  $h$ . Under  $h$ -refinement at fixed  $p$ , the  $p^2$  is absorbed into the constant and only the  $h^{-1}$  affects the convergence rate. Under  $p$ -refinement at fixed  $h$ , the  $p^2$  introduces an algebraic prefactor that shifts the error level but does not destroy exponential convergence.

### 2.6.2. Convergence Under $p$ -Refinement

Under  $p$ -refinement, the mesh is held fixed and the polynomial order  $p$  is increased. The convergence behaviour depends on the smoothness of the exact solution. If  $u$  is analytic (infinitely differentiable with bounded derivatives), the error decreases exponentially with  $p$  [21]:

$$\|u - u_h\|_{H^m} \leq C e^{-\sigma p}, \quad (2.52)$$

for some constant  $\sigma > 0$  that depends on the domain of analyticity of the solution. On a semi-log plot (logarithmic error versus polynomial order), this produces a straight line with slope  $-\sigma$ . For solutions with finite Sobolev regularity  $s$ , the convergence under  $p$ -refinement is algebraic ( $\mathcal{O}(p^{-s})$ ) rather than exponential. In this work, the manufactured solutions used for verification are analytic, so exponential convergence is expected.

For the derived quantities (vorticity, pressure), the inverse estimate (2.51) introduces algebraic prefactors of the form  $p^2/h$  under  $p$ -refinement, but these do not destroy the exponential decay. The result is that on a semi-log plot, the convergence curves for all reconstructed quantities are approximately parallel, separated by constant vertical offsets. Integration (streamfunction from velocity) shifts the curve downward, while the weak curl (vorticity from velocity) shifts it upward.

### 2.6.3. Convergence Under $h$ -Refinement

Under  $h$ -refinement, the polynomial order  $p$  is held fixed and the element size  $h$  is decreased. This yields algebraic convergence: the error decreases as  $\mathcal{O}(h^r)$  for some rate  $r$  that depends on  $p$  and the quantity being measured. On a log-log plot of error versus element size, this produces a straight line with slope  $r$ .

In the mimetic reconstruction, the velocity is the primary unknown obtained by least-squares fitting, and all other quantities are derived from it. Each derivation step changes the convergence rate by one order. The following derivations establish the rate for each quantity.

### 2.6.3.1. Velocity

The velocity  $\mathbf{u}$  is represented by 1-form basis functions, which are tensor products of Lagrange functions (degree  $p$ ) and edge functions (degree  $p-1$ ). The approximation power is limited by the lower-degree direction. By the tensor-product approximation result applied to (2.50) with  $m=0$ , for an analytic solution ( $s \rightarrow \infty$ ):

$$\|\mathbf{u} - \pi_h \mathbf{u}\|_{L^2} \leq C h^p |\mathbf{u}|_{H^p}, \quad (2.53)$$

where  $\pi_h \mathbf{u}$  denotes the best approximation in the 1-form space. The  $H(\text{div}; \Omega)$ -norm of the error is:

$$\|\mathbf{u} - \mathbf{u}_h\|_{H(\text{div}; \Omega)}^2 = \|\mathbf{u} - \mathbf{u}_h\|_{L^2}^2 + \|\nabla \cdot (\mathbf{u} - \mathbf{u}_h)\|_{L^2}^2.$$

For a general  $\mathbf{u}$  this simplification would not apply, and the divergence term would limit the  $H(\text{div}; \Omega)$  convergence rate by one order. In the present setting, however,  $\nabla \cdot \mathbf{u} = 0$  exactly (incompressible flow) and  $E^{2,1} \mathbf{u}_h = 0$  exactly (enforced as a hard constraint in the saddle-point system), so the divergence term vanishes to machine precision. Therefore:

$$\|\mathbf{u} - \mathbf{u}_h\|_{H(\text{div}; \Omega)} = \|\mathbf{u} - \mathbf{u}_h\|_{L^2} = \mathcal{O}(h^p). \quad (2.54)$$

### 2.6.3.2. Vorticity

The vorticity  $\omega$  is not fitted to data directly; it is derived from the velocity through the weak curl formulation (2.27). If the vorticity were fitted by least squares using the 0-form basis functions (degree  $p$  in both directions), the full approximation power of the 0-form space would give  $\|\omega - \omega_h\|_{L^2} = \mathcal{O}(h^{p+1})$ . However, because  $\omega$  is obtained from  $\mathbf{u}$  through the weak curl, the actual convergence rate is lower. To quantify this, we derive both a lower and an upper bound on the vorticity error.

**Error equation.** The exact vorticity satisfies:

$$(\omega, q) = -(\mathbf{u}, \nabla \times q) + \oint_{\partial\Omega} (\mathbf{u} \cdot \hat{\mathbf{t}}) q \, ds, \quad \forall q,$$

and the discrete vorticity satisfies the same weak form with discrete quantities:

$$(\omega_h, q_h) = -(\mathbf{u}_h, \nabla \times q_h) + \oint_{\partial\Omega} (\mathbf{u}_h \cdot \hat{\mathbf{t}}) q_h \, ds, \quad \forall q_h,$$

where  $(\cdot, \cdot)$  denotes the  $L^2$ -inner product. Subtracting and dropping the boundary terms (which are treated analogously) gives the error equation:

$$(q_h, \omega - \omega_h) = -(\nabla \times q_h, \mathbf{u} - \mathbf{u}_h), \quad \forall q_h \in \mathcal{C}^0, \quad (2.55)$$

where  $q_h$  is an arbitrary test function in the 0-form space.

2

**Lower bound.** The Cauchy–Schwarz inequality applied to the left-hand side of (2.55) gives:

$$|(q_h, \omega - \omega_h)| \leq \|q_h\| \|\omega - \omega_h\|,$$

since  $(a, b) = \|a\| \|b\| \cos \theta$  and  $|\cos \theta| \leq 1$ . Substituting the error equation:

$$|(\nabla \times q_h, \mathbf{u} - \mathbf{u}_h)| \leq \|q_h\| \|\omega - \omega_h\|.$$

Dividing both sides by  $\|q_h\|$ :

$$\frac{|(\nabla \times q_h, \mathbf{u} - \mathbf{u}_h)|}{\|q_h\|} \leq \|\omega - \omega_h\|.$$

Since this holds for every  $q_h \neq 0$ , it also holds for the  $q_h$  that maximises the left-hand side:

$$\sup_{q_h \neq 0} \frac{|(\nabla \times q_h, \mathbf{u} - \mathbf{u}_h)|}{\|q_h\|} \leq \|\omega - \omega_h\|. \quad (2.56)$$

The supremum on the left is the dual norm of  $\mathbf{u} - \mathbf{u}_h$ , which is bounded below by the  $H(\text{div}; \Omega)$ -norm of the velocity error. This gives the lower bound:

$$\|\mathbf{u} - \mathbf{u}_h\|_{H(\text{div}; \Omega)} \leq \|\omega - \omega_h\|_{L^2}. \quad (2.57)$$

Since the velocity error is  $\mathcal{O}(h^p)$  from (2.54), the vorticity error is bounded below by  $\mathcal{O}(h^p)$ .

**Upper bound.** To obtain the upper bound, the vorticity error is decomposed using a C ea’s lemma-style splitting. Let  $\omega^I$  denote the canonical interpolant of the exact vorticity at the GLL nodes. It commutes with the discrete curl, so  $\omega^I$  and  $\omega_h$  are compared in the same approximation space and the resulting bound carries the rate dictated by the projection error. The total error is split as:

$$\omega - \omega_h = \underbrace{(\omega - \omega^I)}_{\eta} + \underbrace{(\omega^I - \omega_h)}_{e^h},$$

where  $\eta = \omega - \omega^I$  is the projection error (the part of  $\omega$  that cannot be represented in  $\mathcal{C}^0$ ) and  $e^h = \omega^I - \omega_h$  is the in-space error (both  $\omega^I$  and  $\omega_h$  live in  $\mathcal{C}^0$ ).

Substituting this decomposition into the error equation (2.55):

$$(q_h, \eta + e^h) = -(\nabla \times q_h, \mathbf{u} - \mathbf{u}_h), \quad \forall q_h \in \mathcal{C}^0.$$

Rearranging isolates the discrete error:

$$(q_h, e^h) = -(\nabla \times q_h, \mathbf{u} - \mathbf{u}_h) - (q_h, \eta).$$

Since this holds for all  $q_h$  in the discrete space, choosing  $q_h = e^h$  gives:

$$\|e^h\|^2 = -(\nabla \times e^h, \mathbf{u} - \mathbf{u}_h) - (e^h, \eta).$$

Applying the triangle inequality and then Cauchy–Schwarz to each term on the right-hand side:

$$\|e^h\|^2 \leq |(\nabla \times e^h, \mathbf{u} - \mathbf{u}_h)| + |(e^h, \eta)| \leq \|\nabla \times e^h\| \|\mathbf{u} - \mathbf{u}_h\| + \|e^h\| \|\eta\|.$$

The inverse estimate (2.51) gives  $\|\nabla \times e^h\| \leq C p^2 h^{-1} \|e^h\|$ , so:

$$\|e^h\|^2 \leq C p^2 h^{-1} \|e^h\| \|\mathbf{u} - \mathbf{u}_h\| + \|e^h\| \|\eta\|.$$

Dividing both sides by  $\|e^h\|$ :

$$\|e^h\| \leq C p^2 h^{-1} \|\mathbf{u} - \mathbf{u}_h\| + \|\eta\|.$$

Finally, the triangle inequality  $\|\omega - \omega_h\| = \|e^h + \eta\| \leq \|e^h\| + \|\eta\|$  gives:

$$\|\omega - \omega_h\| \leq C p^2 h^{-1} \|\mathbf{u} - \mathbf{u}_h\| + 2\|\eta\|.$$

The velocity error is  $\|\mathbf{u} - \mathbf{u}_h\| = \mathcal{O}(h^p)$  from (2.54), and the interpolation error  $\|\eta\| = \|\omega - \omega^I\| = \mathcal{O}(h^{p+1})$  since  $\omega^I$  is the 0-form interpolant of degree  $p$ . Therefore:

$$\|\omega - \omega_h\| \leq C \mathcal{O}(p^2 h^{p-1}) + 2 \mathcal{O}(h^{p+1}) = \mathcal{O}(p^2 h^{p-1}),$$

where the first term dominates. The  $p^2$  prefactor is absorbed into the constant for  $h$ -refinement at fixed  $p$ , but is a visible algebraic prefactor under  $p$ -refinement at fixed  $h$ .

**Combined bound.** Together, the lower and upper bounds show that the vorticity error satisfies:

$$\mathcal{O}(h^p) \leq \|\omega - \omega_h\|_{L^2} \leq \mathcal{O}(h^{p-1}). \quad (2.58)$$

For the  $H(\text{curl}; \Omega)$ -norm, the curl of the vorticity error must also be bounded. Since the curl of the vorticity is effectively a second derivative of the velocity, applying the inverse estimate (2.51) a second time gives:

$$\|\omega - \omega_h\|_{H(\text{curl}; \Omega)} = \mathcal{O}(h^{p-2}). \quad (2.59)$$

### 2.6.3.3. Streamfunction

The streamfunction  $\psi$  is a 0-form recovered by integrating the velocity fluxes through the incidence matrix:  $E^{1,0}\psi = \mathbf{u}$ . Integration gains one algebraic order relative to the velocity. The streamfunction error therefore satisfies:

$$\|\psi - \psi_h\|_{L^2} \leq Ch \|\mathbf{u} - \mathbf{u}_h\|_{L^2} = \mathcal{O}(h^{p+1}). \quad (2.60)$$

This is consistent with the full approximation power of the 0-form space ( $h^{p+1}$  for degree- $p$  Lagrange polynomials), even though  $\psi$  is not fitted directly to data.

### 2.6.3.4. Total Pressure

The total pressure  $\tilde{P}$  is a 2-form computed from a discrete system whose right-hand side involves the vorticity force term  $\mathbf{C} = \mathbf{u} \times \omega$ . The pressure error is bounded from below by the best approximation in the 2-form space, which uses edge functions of degree  $p-1$  in both directions, giving  $\mathcal{O}(h^p)$ . For the upper bound, the error in the discrete vorticity force term  $\mathbf{C}_h = \mathbf{u}_h \times \omega_h$  must be quantified. Adding and subtracting:

$$\mathbf{C} - \mathbf{C}_h = \mathbf{u} \times (\omega - \omega_h) + (\mathbf{u} - \mathbf{u}_h) \times \omega_h.$$

By Hölder's inequality, which states that for scalar functions  $f \in L^\infty$  and  $g \in L^2$ :

$$\|f \cdot g\|_{L^2} \leq \|f\|_{L^\infty} \|g\|_{L^2}, \quad (2.61)$$

applied component-wise (noting that in two dimensions  $\mathbf{u} \times \omega$  reduces to the scalar-vector product  $(v\omega, -u\omega)^T$ , so each component is a scalar product to which (2.61) applies), the first term satisfies  $\|\mathbf{u} \times (\omega - \omega_h)\|_{L^2} \leq \|\mathbf{u}\|_{L^\infty} \|\omega - \omega_h\|_{L^2}$ . Since  $\|\mathbf{u}\|_{L^\infty}$  is bounded independently of  $h$ , this is  $\mathcal{O}(h^{p-1})$  from (2.58). The second term satisfies  $\|(\mathbf{u} - \mathbf{u}_h) \times \omega_h\|_{L^2} \leq \|\omega_h\|_{L^\infty} \|\mathbf{u} - \mathbf{u}_h\|_{L^2} = \mathcal{O}(h^p)$  from (2.54). The dominant contribution is the first, so:

$$\|\mathbf{C} - \mathbf{C}_h\|_{L^2} = \mathcal{O}(h^{p-1}). \quad (2.62)$$

By the stability theory for saddle-point problems [32], the pressure error depends continuously on the right-hand side error: if the data is perturbed by  $\mathcal{O}(h^{p-1})$ , the pressure error is at most  $\mathcal{O}(h^{p-1})$ . The total pressure error is therefore bounded by:

$$\mathcal{O}(h^p) \leq \|\tilde{P} - \tilde{P}_h\|_{L^2} \leq \mathcal{O}(h^{p-1}). \quad (2.63)$$

### 2.6.3.5. Summary of Rates

Table 2.3 collects the theoretical convergence rates under  $h$ -refinement. The ordering  $\psi \succ \mathbf{u} \succ \omega$  reflects the relationship between each quantity and the primary unknown: the streamfunction is obtained by integration of the velocity (gaining one order), while the vorticity is obtained through the weak curl of the velocity (costing one order through the inverse estimate).

**TABLE 2.3** Theoretical  $h$ -refinement convergence rates at fixed polynomial order  $p$ .

Quantity	Norm	Rate	Reason
$\psi$	$L^2$	$\mathcal{O}(h^{p+1})$	Primitive of $\mathbf{u}$
$\mathbf{u}$	$H(\text{div}; \Omega)$	$\mathcal{O}(h^p)$	Primary unknown
$\tilde{P}$	$L^2$	$\mathcal{O}(h^{p-1})$ to $\mathcal{O}(h^p)$	Saddle-point data error
$\omega$	$L^2$	$\mathcal{O}(h^{p-1})$ to $\mathcal{O}(h^p)$	Weak curl of $\mathbf{u}$
$\omega$	$H(\text{curl}; \Omega)$	$\mathcal{O}(h^{p-2})$	Curl of weak curl of $\mathbf{u}$

These theoretical rates will be verified numerically in Chapter 5: under  $h$ -refinement at fixed  $p$ , the measured slopes on log-log plots should match the rates in Table 2.3, while under  $p$ -refinement with analytic solutions, all quantities should exhibit exponential convergence with approximately parallel curves on a semi-log plot.



# Chapter 3

## Methodology

This chapter develops the reconstruction methodology. From scattered particle tracking data, the method produces an exactly divergence-free velocity field and recovers vorticity, streamfunction, and pressure from it through the same set of discrete mimetic operators. The velocity is sought in  $H(\text{div}; \Omega)$ , which means normal flux continuity is enforced across element interfaces while the tangential component may be discontinuous. Because  $\mathbf{u} \notin H(\text{curl}; \Omega)$ , the strong curl is not guaranteed to be square-integrable, and the vorticity  $\omega \in H(\text{curl}; \Omega)$  must be recovered via a weak formulation. The total pressure  $\tilde{P} \in L^2$  is obtained from a second saddle-point system involving the vorticity force term.

The three function spaces and their associated quantities are connected through the de Rham complex (2.22):

$$H(\text{curl}; \Omega) \xrightarrow{\nabla \times} H(\text{div}; \Omega) \xrightarrow{\nabla \cdot} L^2.$$

The discrete complex (Section 2.5.3.) preserves this structure exactly through the incidence matrices  $E^{1,0}$  and  $E^{2,1}$  (Section 2.5.2.). In particular, the complex property  $E^{2,1}E^{1,0} = 0$  (2.37) guarantees that the divergence-free constraint is satisfied exactly at the algebraic level, not merely approximately. The theoretical convergence rates for the reconstructed quantities were established in Section 2.6. and are summarised in Table 2.3.

The chapter proceeds as follows. Section 3.1. states the reconstruction problem and the function space setting. Section 3.2. introduces the multi-element mesh and coordinate mapping. Section 3.3. develops the constrained least-squares formulation for the velocity. Sections 3.4. and 3.5. describe the reconstruction of velocity, streamfunction, and vorticity. Section 3.6. derives the pressure computation. Section 3.7. summarises the complete algorithm.

### 3.1. Problem Statement and Function Space Setting

The goal of this work is to reconstruct an exactly divergence-free velocity field from scattered particle tracking velocimetry (PTV) measurements, and to recover the vorticity, streamfunction, and pressure from that velocity through the same set of discrete operators. Let  $\Omega \subset \mathbb{R}^2$  be the measurement domain, and let  $\mathbf{u}^d = \{(x_i, y_i, u_i, v_i)\}_{i=1}^{n_r}$  denote the set of  $n_r$  measured velocity vectors at scattered locations  $(x_i, y_i) \in \Omega$ . The reconstruction problem is:

*Given scattered particle tracking data  $\mathbf{u}^d$  on  $\Omega$ , find a velocity field  $\mathbf{u}$  that best fits the data in a least-squares sense while satisfying the incompressibility constraint  $\nabla \cdot \mathbf{u} = 0$  exactly. From*

this velocity, *derive* the vorticity  $\omega$ , streamfunction  $\psi$ , and total pressure  $\tilde{P}$ .

The function space setting determines the regularity of each reconstructed quantity. The velocity is sought in  $H(\text{div}; \Omega)$ , as defined in (2.21):

$$\mathbf{u} \in H(\text{div}; \Omega), \quad \omega \in H(\text{curl}; \Omega), \quad \tilde{P} \in L^2(\Omega).$$

The Mimetic Spectral Element discretization (Section 2.5.) represents  $\mathbf{u}$  as a 1-form using the basis functions (2.29)–(2.30), which are mapped to the physical domain by the contravariant Piola transformation (2.44). This transformation preserves normal flux continuity across element interfaces: on a shared edge, the normal component of  $\mathbf{u}$  is continuous, while the tangential component may be discontinuous. This is precisely  $H(\text{div}; \Omega)$ -regularity (Table 2.1).

The velocity is placed in  $H(\text{div}; \Omega)$  rather than  $H(\text{curl}; \Omega)$  primarily so that  $\nabla \cdot \mathbf{u}$  is well-defined element-by-element and the incompressibility condition  $\nabla \cdot \mathbf{u} = 0$  can be enforced exactly at the discrete level through the discrete divergence operator. A secondary reason concerns the experimental data: the tangential velocity across a shear layer or near a wall may exhibit sharp gradients or even discontinuities due to measurement noise or insufficient spatial resolution. Enforcing tangential continuity (i.e.  $H(\text{curl}; \Omega)$ -regularity) would impose a smoothness constraint that the data does not support, whereas enforcing only normal flux continuity preserves mass conservation across element interfaces while allowing the tangential component the freedom to represent sharp gradients faithfully.

A consequence of  $\mathbf{u} \in H(\text{div}; \Omega)$  is that the strong curl  $\nabla \times \mathbf{u}$  is not guaranteed to be square-integrable, since  $H(\text{div}; \Omega)$  does not guarantee sufficient regularity for the tangential derivatives. The vorticity must therefore be recovered through the weak curl formulation (2.27), as derived in Section 2.4.4.. The resulting vorticity  $\omega$  belongs to  $H(\text{curl}; \Omega)$  (2.20): its tangential trace is continuous across element interfaces, but its normal derivative may jump. This is the natural dual regularity dictated by the de Rham complex (Section 2.4.3.).

The total pressure  $\tilde{P} = p + \frac{1}{2}|\mathbf{u}|^2$  is a 2-form and belongs to  $L^2(\Omega)$  (2.19). It is computed from a second saddle-point system involving the vorticity force term  $\mathbf{u} \times \omega$  (Section 3.6.).

### 3.2. Domain Decomposition and Coordinate Mapping

The Mimetic Spectral Element Method requires the global measurement domain to be partitioned into a structured mesh of elements, each mapped to a common reference element on which the basis functions are defined (Section 2.3.). This section defines the multi-element mesh and the affine coordinate mapping between the physical and reference domains.

#### 3.2.1. Multi-Element Mesh

The global domain  $[x_L, x_R] \times [y_L, y_R]$  is partitioned into  $K_x \times K_y$  rectangular elements arranged in a Cartesian grid. The element boundaries in the  $x$ -direction are defined by their left and right edges:

$$a_x^{(k)} = x_L + \frac{k}{K_x}(x_R - x_L), \quad (3.1)$$

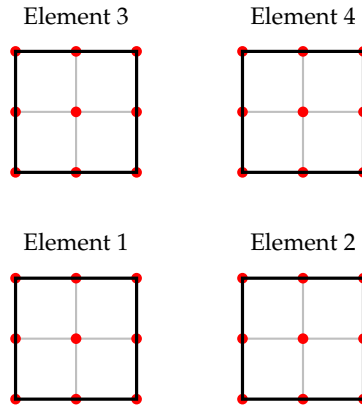
$$b_x^{(k)} = x_L + \frac{k+1}{K_x}(x_R - x_L), \quad (3.2)$$

for  $k = 0, \dots, K_x - 1$ , where  $a_x^{(k)}$  is the left boundary and  $b_x^{(k)}$  is the right boundary of element  $k$ . Similarly, in the  $y$ -direction:

$$a_y^{(\ell)} = y_L + \frac{\ell}{K_y}(y_R - y_L), \quad (3.3)$$

$$b_y^{(\ell)} = y_L + \frac{\ell+1}{K_y}(y_R - y_L), \quad (3.4)$$

for  $\ell = 0, \dots, K_y - 1$ . Each element is defined by the rectangle  $[a_x^{(k)}, b_x^{(k)}] \times [a_y^{(\ell)}, b_y^{(\ell)}]$ , with element size  $h_x = (x_R - x_L)/K_x$  in the  $x$ -direction and  $h_y = (y_R - y_L)/K_y$  in the  $y$ -direction. Figure 3.1 illustrates the multi-element mesh for  $K_x = K_y = 2$ .



**FIGURE 3.1** Multi-element mesh with  $K_x = K_y = 2$  and  $p_x = p_y = 2$ . Each element contains a local  $3 \times 3$  red node grid with no shared nodes. Light gray lines show internal element structure; thick black lines denote element boundaries.

### 3.2.2. Affine Mapping to the Reference Element

All basis functions are defined on the reference element  $[-1, 1]^2$  (Section 2.3.). To evaluate these basis functions at physical measurement locations, the coordinates  $(x, y)$  of each data point are mapped to reference coordinates  $(\xi, \eta)$  using the inverse affine transformation:

$$\xi = \frac{2}{b_x^{(k)} - a_x^{(k)}} x - \frac{b_x^{(k)} + a_x^{(k)}}{b_x^{(k)} - a_x^{(k)}}, \quad (3.5)$$

$$\eta = \frac{2}{b_y^{(\ell)} - a_y^{(\ell)}} y - \frac{b_y^{(\ell)} + a_y^{(\ell)}}{b_y^{(\ell)} - a_y^{(\ell)}}. \quad (3.6)$$

The Jacobian matrix of the forward mapping from reference to physical coordinates is:

$$J = \begin{pmatrix} \frac{\partial x}{\partial \xi} & \frac{\partial x}{\partial \eta} \\ \frac{\partial y}{\partial \xi} & \frac{\partial y}{\partial \eta} \end{pmatrix} = \begin{pmatrix} \frac{b_x^{(k)} - a_x^{(k)}}{2} & 0 \\ 0 & \frac{b_y^{(\ell)} - a_y^{(\ell)}}{2} \end{pmatrix}, \quad (3.7)$$

with determinant:

$$\det J = \frac{(b_x^{(k)} - a_x^{(k)})(b_y^{(\ell)} - a_y^{(\ell)})}{4}. \quad (3.8)$$

For a uniform mesh,  $b_x^{(k)} - a_x^{(k)} = h_x$  and  $b_y^{(\ell)} - a_y^{(\ell)} = h_y$  for all elements, so the Jacobian is constant across the mesh:  $J = \text{diag}(h_x/2, h_y/2)$  with  $\det J = h_x h_y / 4$ .

This work restricts attention to affine maps between the reference and physical elements. For curvilinear elements, the Jacobian would vary within each element and would need to be evaluated at every quadrature point, requiring the general Piola transformation framework described in Section 2.5.5.. The affine restriction simplifies the mass matrix transformations (2.47)–(2.49) to constant scalings per element.

### 3.3. Constrained Least-Squares Formulation for Velocity

This section builds the constrained least-squares system for the velocity reconstruction. Starting from local shape function evaluation, the formulation proceeds through global assembly, inter-element continuity constraints, and the divergence-free constraint, culminating in the complete saddle-point system.

Throughout this work, we denote  $N = p_\xi + 1$  and  $M = p_\eta + 1$  as the number of Gauss–Lobatto–Legendre nodes in the  $\xi$ - and  $\eta$ -directions, respectively.

### 3.3.1. Local Shape Function Evaluation

Within each element  $k$ , the velocity field is represented using the 1-form basis functions defined in Section 2.5.1.. The shape function matrix for element  $k$  is:

$$\Psi_k^{(1)} = \begin{pmatrix} \Psi_u^{(1)} \\ \Psi_v^{(1)} \end{pmatrix}, \quad (3.9)$$

where the components are tensor products of Lagrange and edge basis functions (Section 2.3.3. and 2.3.4.):

$$\Psi_u^{(1)} = h_i(\xi) e_j(\eta), \quad (3.10)$$

$$\Psi_v^{(1)} = e_i(\xi) h_j(\eta). \quad (3.11)$$

These basis functions are defined on the reference element  $[-1, 1]^2$  and are evaluated at the reference coordinates  $(\xi_m, \eta_m)$  of each measurement point, obtained from the inverse affine mapping (3.5)–(3.6).

To relate the reference-domain basis functions to physical-domain velocities, the contravariant Piola transformation (2.44) is applied. For an affine mapping with Jacobian (3.7), this acts component-wise on the two-component velocity vector at each evaluation point:

$$\begin{pmatrix} u(x, y) \\ v(x, y) \end{pmatrix} = \frac{1}{\det J} \begin{pmatrix} \frac{\partial x}{\partial \xi} & \frac{\partial x}{\partial \eta} \\ \frac{\partial y}{\partial \xi} & \frac{\partial y}{\partial \eta} \end{pmatrix} \begin{pmatrix} \Psi_u^{(1)}(\xi, \eta) \\ \Psi_v^{(1)}(\xi, \eta) \end{pmatrix} \mathbf{u}_k, \quad (3.12)$$

where  $\mathbf{u}_k$  is the vector of edge-based degrees of freedom for element  $k$ . The Jacobian matrix  $J$  acts on the two-component basis function vector at each evaluation point, not on individual rows of the shape function matrix. For the diagonal Jacobian of an affine map, this simplifies to scaling  $\Psi_u^{(1)}$  by  $J_{00}/\det J$  and  $\Psi_v^{(1)}$  by  $J_{11}/\det J$ .

The resulting physical-domain shape function matrix  $\Psi_k^{(1)}(x, y)$  has dimensions  $2n_k \times n_{\text{dof}}^{(k)}$ , where  $n_k$  is the number of measurement points in element  $k$  (the factor 2 accounts for both velocity components) and  $n_{\text{dof}}^{(k)}$  is the number of local degrees of freedom.

### 3.3.2. Edge-Based Degrees of Freedom

Within each element, the velocity degrees of freedom correspond to flux integrals along element edges (Section 2.5.1.). The horizontal fluxes  $u_i$  are associated with vertical edges, and the vertical fluxes  $v_j$  with horizontal


edges. For polynomial orders  $p_\varepsilon$  and  $p_\eta$ , each element contains  $N(M-1)$  horizontal flux DoFs and  $M(N-1)$  vertical flux DoFs, giving a total of

$$n_{\text{dof}}^{(k)} = N(M-1) + M(N-1) \quad (3.13)$$

degrees of freedom per element. For example, with  $N = M = 3$  ( $p = 2$ ), each element has  $3 \cdot 2 + 3 \cdot 2 = 12$  DoFs: six horizontal fluxes  $u_0$ – $u_5$  and six vertical fluxes  $v_0$ – $v_5$ . The local DoF vector is ordered as:

$$\mathbf{u}_k = (u_0, u_1, \dots, u_{N(M-1)-1}, v_0, v_1, \dots, v_{M(N-1)-1})^T. \quad (3.14)$$

This ordering is consistent with the structure of  $\Psi_k^{(1)}$ , where the first  $N(M-1)$  columns correspond to horizontal edge basis functions and the remaining  $M(N-1)$  columns to vertical edge basis functions. Figure 3.2 illustrates the flux-based degrees of freedom for a single element with  $N = M = 3$ .



**FIGURE 3.2** A single element with  $3 \times 3$  nodes. Blue arrows represent velocity flux DoFs across edges:  $u_0$ – $u_5$  (horizontal) and  $v_0$ – $v_5$  (vertical).

### 3.3.3. Global Least-Squares Assembly

Each element  $k$  contributes a local shape function matrix  $\Psi_k^{(1)}$  and a local data vector  $\mathbf{u}_k^d$  (the subset of measurement points falling within element  $k$ ). These are assembled into a global block-diagonal system:

$$\min_{\mathbf{u}} \left\| \underbrace{\begin{pmatrix} \Psi_0^{(1)} & & \\ & \ddots & \\ & & \Psi_{K-1}^{(1)} \end{pmatrix}}_A \begin{pmatrix} \mathbf{u}_0 \\ \vdots \\ \mathbf{u}_{K-1} \end{pmatrix} - \mathbf{u}^d \right\|^2, \quad (3.15)$$

where  $K = K_x \cdot K_y$  is the total number of elements. The global matrix  $A \in \mathbb{R}^{n_{\text{data}} \times n_{\text{dof}}}$  is block-diagonal, with  $n_{\text{data}}$  the total number of measurement points (counting both velocity components) and  $n_{\text{dof}} = \sum_{k=0}^{K-1} n_{\text{dof}}^{(k)}$  the total number of flux DoFs. The normal-equation solution of the unconstrained problem is:

$$\hat{\mathbf{u}} = (A^T A)^{-1} A^T \mathbf{u}^d. \quad (3.16)$$

3

Because  $A$  is block-diagonal, each element is solved independently at this stage. The velocity field obtained from (3.16) is therefore discontinuous across element interfaces: the fluxes on shared edges are determined independently by each element's local data, with no communication between neighbours.

### 3.3.4. Normal Flux Continuity Constraints

To enforce  $H(\text{div}; \Omega)$ -regularity (Section 3.1.), the normal flux must be continuous across element interfaces. This is achieved by introducing a continuity constraint matrix  $N_L$  that equates matching flux DoFs on shared edges. For each pair of adjacent elements sharing an edge, the constraint requires:


$$\mathbf{u}_i^{(L)} - \mathbf{u}_j^{(R)} = 0 \quad \text{or} \quad \mathbf{v}_i^{(B)} - \mathbf{v}_j^{(T)} = 0, \quad (3.17)$$

where the superscripts denote the left/right or bottom/top element at the shared edge. Each shared edge contributes  $M - 1$  constraints (for vertical interfaces) or  $N - 1$  constraints (for horizontal interfaces), one per interior flux DoF on that edge. These constraints are collected into the linear system:

$$N_L \mathbf{u} = \mathbf{0}, \quad (3.18)$$

where  $N_L \in \mathbb{R}^{n_c \times n_{\text{dof}}}$  is a sparse matrix with  $n_c$  rows (one per constraint). Each row of  $N_L$  contains exactly two nonzero entries:  $+1$  at the column corresponding to the flux on one side of the interface and  $-1$  at the column corresponding to the matching flux on the other side.

For example, with  $K_x = K_y = 2$  and  $N = M = 3$ , each element has 12 DoFs, giving  $n_{\text{dof}} = 48$ . There are 4 shared edges with 2 constraints each, resulting in  $N_L \in \mathbb{R}^{8 \times 48}$ . Figure 3.3 illustrates the 8 Lagrange multipliers  $\lambda_0 - \lambda_7$  enforcing flux continuity at the element interfaces.



**FIGURE 3.3**  $2 \times 2$  element mesh ( $K_x = K_y = 2$ ,  $p_\xi = p_\eta = 2$ ), each with  $3 \times 3$  red nodes. Eight Lagrange multipliers  $\lambda_0$ – $\lambda_7$  enforce flux continuity: two per shared edge, drawn exactly between the elements.

In addition to internal continuity, the matrix  $N_L$  can be extended to enforce boundary conditions. For no-penetration conditions on solid walls, extra rows are added with a single  $+1$  entry at the column of the boundary flux DoF and zero elsewhere, enforcing that the normal flux vanishes at the wall.

### 3.3.5. Divergence-Free Constraint

The continuity constraints ensure matching fluxes across element interfaces but do not by themselves enforce that the net flux into each cell vanishes. To impose the divergence-free condition  $\nabla \cdot \mathbf{u} = 0$  exactly, the incidence matrix  $E^{2,1}$  (Section 2.5.2.) is used. This matrix maps edge-based flux DoFs to cell divergences, with dimensions:

$$E^{2,1} \in \mathbb{R}^{(N-1)(M-1) \times [N(M-1) + M(N-1)]}. \quad (3.19)$$

The number of rows equals the number of 2-form cells (divergence control volumes), and the number of columns equals the number of 1-form edge DoFs. Each row contains entries  $\{-1, 0, +1\}$  corresponding to the incoming and outgoing fluxes through the edges surrounding that cell, as illustrated for  $N = M = 3$  in (2.35).

The global divergence operator is constructed by replicating the local incidence matrix across all elements using the Kronecker product:

$$E_t^{2,1} = I_K \otimes E^{2,1}, \quad (3.20)$$

where  $I_K$  is the  $K \times K$  identity matrix. The divergence-free constraint is then:

$$E_t^{2,1} \mathbf{u} = \mathbf{0}, \quad (3.21)$$

which enforces zero net divergence in every cell of every element. This constraint is exact at the algebraic level: it holds to machine precision in the discrete solution, regardless of mesh resolution or polynomial order. The Lagrange multipliers  $\zeta$  associated with this constraint can be interpreted as cell-wise potential-like quantities that enforce local mass conservation.

## 3

### 3.3.6. Complete Saddle-Point System

Combining the least-squares objective (3.15) with the continuity constraints (3.18) and the divergence-free constraint (3.21), the constrained minimisation problem is:

$$\min_{\mathbf{u}} \|\mathbf{A}\mathbf{u} - \mathbf{u}^d\|^2 \quad \text{subject to} \quad \mathbf{N}_L \mathbf{u} = \mathbf{0}, \quad \mathbf{E}_t^{2,1} \mathbf{u} = \mathbf{0}. \quad (3.22)$$

Introducing Lagrange multipliers  $\lambda$  for the continuity constraints and  $\zeta$  for the divergence constraints, the first-order optimality conditions yield the saddle-point system (KKT system):

$$\begin{pmatrix} \mathbf{A}^T \mathbf{A} & \mathbf{N}_L^T & (\mathbf{E}_t^{2,1})^T \\ \mathbf{N}_L & \mathbf{0} & \mathbf{0} \\ \mathbf{E}_t^{2,1} & \mathbf{0} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \lambda \\ \zeta \end{pmatrix} = \begin{pmatrix} \mathbf{A}^T \mathbf{u}^d \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \quad (3.23)$$

This is the algebraic system arising from a constrained least-squares formulation with  $\mathbf{u} \in H(\text{div}; \Omega)$ . The Lagrange multiplier  $\lambda$  enforces inter-element normal flux continuity, ensuring  $H(\text{div}; \Omega)$ -regularity of the reconstructed velocity field. The Lagrange multiplier  $\zeta$  enforces the divergence-free constraint in every cell, ensuring exact incompressibility at the discrete level. Solving (3.23) yields the velocity flux DoFs  $\mathbf{u}$  that best fit the particle tracking data while satisfying both constraints exactly.

## 3.4. Reconstruction of Velocity and Streamfunction

Once the velocity flux DoFs  $\mathbf{u}$  have been obtained from the saddle-point system (3.23), the continuous velocity field, its divergence, and the streamfunction are reconstructed on a fine evaluation grid. For each element, this grid consists of  $N_{\text{fine}} \times N_{\text{fine}}$  Gauss–Lobatto–Legendre nodes, yielding a global grid of  $K_x N_{\text{fine}} \times K_y N_{\text{fine}}$  points.

### 3.4.1. Velocity Reconstruction on Fine Grid

The velocity field is reconstructed by evaluating the 1-form basis functions (2.29)–(2.30) at the fine-grid Gauss–Lobatto–Legendre nodes and applying the Piola transformation (2.44). For each element  $k$ , the velocity at

a physical point  $(x, y)$  is:

$$\begin{pmatrix} u(x, y) \\ v(x, y) \end{pmatrix} = \frac{1}{\det J} J \Psi_k^{(1)}(\xi, \eta) \mathbf{u}_k, \quad (3.24)$$

where  $(\xi, \eta)$  are the reference coordinates corresponding to  $(x, y)$  via the inverse affine mapping (3.5)–(3.6), and  $\mathbf{u}_k$  is the local vector of flux DoFs for element  $k$ . The reconstruction is performed independently on each element. Because the continuity constraints (3.18) enforce matching normal fluxes at shared edges, the resulting global velocity field is continuous in the normal direction across element interfaces ( $H(\text{div}; \Omega)$ -regularity).

3

### 3.4.2. Divergence Verification

Although the divergence-free constraint (3.21) is enforced exactly at the algebraic level, it is useful to verify this on the fine reconstruction grid. The divergence field is computed using the 2-form basis functions (2.31) and the incidence matrix  $E^{2,1}$  (2.35):

$$(\nabla \cdot \mathbf{u})(x, y) = \Psi_k^{(2)}(x, y) E^{2,1} \mathbf{u}_k, \quad (3.25)$$

where  $\Psi_k^{(2)}(x, y) = e_i(\xi) e_j(\eta) / \det J$  is the 2-form basis mapped to physical coordinates via the Piola transformation for 2-forms (2.46). Since  $E^{2,1} \mathbf{u}_k = \mathbf{0}$  is enforced exactly, the reconstructed divergence should be  $\mathcal{O}(\varepsilon_{\text{machine}})$  everywhere, confirming that incompressibility holds at the discrete level.

### 3.4.3. Streamfunction Reconstruction

The streamfunction  $\psi$  is a 0-form (scalar field on nodes) related to the velocity fluxes through the incidence matrix  $E^{1,0}$  (2.34):

$$E^{1,0} \boldsymbol{\psi}_k = \mathbf{u}_k, \quad (3.26)$$

where  $\boldsymbol{\psi}_k$  is the vector of nodal streamfunction values in element  $k$ . This relation is exact: it follows from the de Rham complex property  $E^{2,1} E^{1,0} = 0$  (2.37) combined with the divergence-free constraint  $E^{2,1} \mathbf{u}_k = \mathbf{0}$ , which together guarantee that  $\mathbf{u}_k$  lies in the image of  $E^{1,0}$ .

#### 3.4.3.1. Single-Element Marching Update


Relation (3.26) is rectangular ( $E^{1,0}$  has more rows than columns), but it is consistent for divergence-free  $\mathbf{u}_k$ , so no linear system needs to be solved. The streamfunction is determined up to a constant, which is fixed by setting  $\psi_0 = 0$  at one corner node. From this seed, every remaining nodal

value is then propagated by a marching update: each unknown node is obtained directly from a single row of (3.26) that involves one already-known and one unknown nodal value,

$$\psi_j = \frac{u_e - \sigma_i \psi_i}{\sigma_j},$$

where  $u_e$  is the flux on edge  $e$  connecting nodes  $i$  and  $j$ , and  $\sigma_i, \sigma_j$  are the corresponding sign entries in  $E^{1,0}$ . Figure 3.4 illustrates the node and flux numbering, with the counterclockwise marching orientation.

3



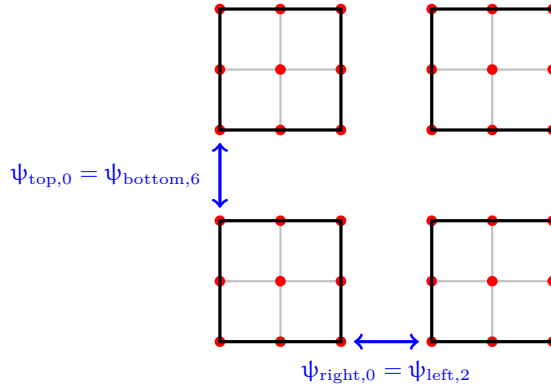
**FIGURE 3.4** Flux directions and node numbering in a single element. Blue arrows represent flux DoFs  $u_i, v_i$ ; green nodes are streamfunction values  $\psi_i$ . Counterclockwise arrows indicate the marching orientation used to propagate known  $\psi$ -values across the element.

### 3.4.3.2. Multi-Element Update

In the multi-element case, the same marching update is applied independently on each element; no global linear system is assembled or solved. The elements are visited in a deterministic order: starting from the bottom-left element, the algorithm marches left to right along each row, then proceeds to the row above. Each element receives a single seed value from a previously processed neighbour:

- The first element uses  $\psi_0 = 0$  at the bottom-left corner of the domain.
- The element to the right uses the bottom-right corner node of the current element as its seed:  $\psi_0^{\text{right}} = \psi_{N-1}^{\text{left}}$ .
- The element above uses the top-left corner node of the element below as its seed:  $\psi_0^{\text{top}} = \psi_{(M-1)N}^{\text{bottom}}$ .

Coupling between elements therefore reduces to passing a single corner value across the shared boundary; no global system is needed. Figure 3.5 illustrates the multi-element marching strategy.



**FIGURE 3.5**  $2 \times 2$  element mesh ( $K_x = K_y = 2$ ,  $p_\xi = p_\eta = 2$ ) showing how each element's lower-left node (except the first) is initialised from its neighbour: to the right  $\psi_{right,0} = \psi_{left,2}$ , and above  $\psi_{top,0} = \psi_{bottom,6}$ .

### 3.4.3.3. Fine-Grid Reconstruction

Once the nodal streamfunction values  $\psi_k$  are known, the continuous streamfunction field is obtained by evaluating the 0-form basis functions (2.28):

$$\psi(x, y) = \Psi_k^{(0)}(\xi, \eta) \psi_k, \quad \Psi_k^{(0)}(\xi, \eta) = h_i(\xi) h_j(\eta). \quad (3.27)$$

Since the 0-form represents a scalar field, no Jacobian scaling is required:  $\Psi_k^{(0)}(x, y) = \Psi_k^{(0)}(\xi(x, y), \eta(x, y))$  (2.43). The reconstruction is performed on the same  $N_{\text{fine}} \times N_{\text{fine}}$  Gauss–Lobatto–Legendre grid as the velocity.

## 3.5. Vorticity Reconstruction

This section derives the discrete vorticity field from the reconstructed velocity. Because  $\mathbf{u} \in H(\text{div}; \Omega)$ , the strong curl is not guaranteed to be square-integrable (Section 3.1.); instead, the vorticity is recovered via a weak formulation. The section proceeds from the continuous weak form through the mimetic discretisation, boundary treatment, multi-element assembly, and fine-grid reconstruction.

### 3.5.1. Weak Formulation Requirement

The velocity field obtained from the constrained least-squares solve (Section 3.3.6.) belongs to  $H(\text{div}; \Omega)$ : its normal component is continuous across element interfaces, but the tangential component may jump. If  $\mathbf{u}$

were in  $H(\text{curl}; \Omega)$ , the strong  $\text{curl } \nabla \times \mathbf{u}$  would be well-defined. However, the primal 1-form discretisation places  $\mathbf{u}$  in  $H(\text{div}; \Omega)$ , so the classical curl cannot be evaluated directly. Instead, the vorticity must be recovered through a weak formulation that transfers the derivative onto a test function via integration by parts. The resulting vorticity  $\omega$  is a 0-form (scalar field on nodes) with the regularity expected from the discrete de Rham complex (2.22).

## 3

### 3.5.2. Weak Formulation

We begin from the continuous vorticity definition  $\omega = \nabla \times \mathbf{u}$ . Multiplying by a scalar test function  $q \in H(\text{curl}; \Omega)$  and integrating over  $\Omega$ :

$$\int_{\Omega} q \omega \, d\Omega = \int_{\Omega} q (\nabla \times \mathbf{u}) \, d\Omega. \quad (3.28)$$

Integration by parts of the right-hand side yields:

$$\int_{\Omega} q (\nabla \times \mathbf{u}) \, d\Omega = \int_{\Omega} (\nabla \times q) \cdot \mathbf{u} \, d\Omega - \int_{\partial\Omega} (\mathbf{u} \times \mathbf{n}) \cdot q \, dS, \quad (3.29)$$

where  $\mathbf{n}$  is the outward unit normal on  $\partial\Omega$ . In two dimensions  $q$  is a scalar test function, and its curl  $\nabla \times q = (-\partial_y q, \partial_x q)^T$  is a 2D vector. The boundary term involves the tangential velocity along  $\partial\Omega$ .

Combining (3.28) and (3.29):

$$\int_{\Omega} q \omega \, d\Omega - \int_{\Omega} (\nabla \times q) \cdot \mathbf{u} \, d\Omega = - \int_{\partial\Omega} (\mathbf{u} \times \mathbf{n}) q \, dS. \quad (3.30)$$

This states that the vorticity–velocity relation  $\omega = \nabla \times \mathbf{u}$  holds in a variational sense, with additional boundary contributions emerging naturally from the integration by parts.

### 3.5.3. Mimetic Discretisation

The vorticity  $\omega$  and test function  $q$  are expanded in the 0-form basis  $\Psi^{(0)} = h_i(\xi) h_j(\eta)$  (2.28), and the velocity  $\mathbf{u}$  in the 1-form basis  $\Psi^{(1)}$  (2.29)–(2.30). Applying Corollary 1 of [31], which guarantees that derivatives of nodal basis functions are exactly represented in the edge function space through the incidence matrix, the three terms in (3.30) discretise to:

$$\int_{\Omega} q \omega \, d\Omega = \mathbf{q}^T \mathbf{M}^0 \boldsymbol{\omega}, \quad (3.31)$$

$$\int_{\Omega} (\nabla \times q) \cdot \mathbf{u} \, d\Omega = \mathbf{q}^T (\mathbf{E}^{1,0})^T \mathbf{M}^1 \mathbf{u}, \quad (3.32)$$

$$- \int_{\partial\Omega} (\mathbf{u} \times \mathbf{n}) q \, dS = \mathbf{q}^T \tilde{\mathbf{B}}^1(\mathbf{u}_b), \quad (3.33)$$

where  $M^0$  and  $M^1$  are the 0-form and 1-form mass matrices defined in Section 2.5.4., and  $\tilde{B}^1$  encodes the boundary contributions.


Since (3.30) must hold for arbitrary  $\mathbf{q}$ , the discrete vorticity–velocity relation reads:

$$M^0 \boldsymbol{\omega} = (E^{1,0})^T M^1 \mathbf{u} + \tilde{B}^1(\mathbf{u}_b). \quad (3.34)$$

The mass matrices are mapped to the physical domain using the Jacobian transformations for affine maps (2.47)–(2.48).

### 3.5.4. Boundary Treatment

The vorticity degrees of freedom are defined on the primal mesh nodes, in the same manner as the streamfunction. Figure 3.6 illustrates this association for  $N = M = 3$ : each primal node carries a vorticity unknown  $\omega_i$ , interpreted as the counterclockwise circulation around the surrounding dual cell.



**FIGURE 3.6** Association of vorticity DoFs with the primal mesh. The primal grid is shown in black, dual control volumes in blue dashed lines. Each red node corresponds to a vorticity unknown  $\omega_i$ , interpreted as counterclockwise circulation around the surrounding dual cell (green arrows). The dual mesh does not extend beyond the element.

The boundary term  $\tilde{B}^1$  requires the tangential velocity along  $\partial\Omega$ . Two cases arise depending on whether the exact boundary velocity is available.

#### 3.5.4.1. Case 1: Exact Boundary Velocities Known

When the exact velocity is known on  $\partial\Omega$  (e.g., in manufactured solution tests, or when experimental data covers the boundary), the tangential

boundary velocities are computed directly. For each boundary edge, a one-dimensional integral is evaluated using Gauss–Lobatto–Legendre quadrature along that edge. The contributions on the bottom, top, left, and right boundaries (for  $i = 0, 1, \dots, N - 1$ ) are:

$$\mathbf{u}_{i,\text{bottom}} = - \sum_p \mathbf{u}(x_p, y_{\text{bottom}}) h_i(x_p) w_p J_{00}, \quad (3.35)$$

$$\mathbf{u}_{i,\text{top}} = - \sum_p \mathbf{u}(x_p, y_{\text{top}}) h_i(x_p) w_p J_{00}, \quad (3.36)$$

$$v_{i,\text{left}} = - \sum_p v(x_{\text{left}}, y_p) h_i(y_p) w_p J_{11}, \quad (3.37)$$

$$v_{i,\text{right}} = - \sum_p v(x_{\text{right}}, y_p) h_i(y_p) w_p J_{11}, \quad (3.38)$$

where  $h_i$  is the one-dimensional Lagrange basis function,  $w_p$  the quadrature weight, and  $J_{00}, J_{11}$  the diagonal Jacobian entries accounting for the mapping from reference to physical coordinates. Since the vorticity DoFs are defined counterclockwise, contributions on the left and top boundaries require an additional minus sign to enforce orientation consistency.


The resulting boundary vector  $\tilde{\mathcal{B}}^1$  has length  $NM$  (one entry per dual control volume). Figure 3.7 illustrates the assembly for  $N = M = 3$ . For instance, the bottom-left dual cell receives contributions from both the bottom and left boundaries:

$$\tilde{\mathcal{B}}_0^1 = \mathbf{u}_{0,\text{bottom}} - v_{0,\text{left}}. \quad (3.39)$$

The full vector for  $N = M = 3$  is:

$$\tilde{\mathcal{B}}^1 = \begin{pmatrix} \mathbf{u}_{0,\text{bottom}} - v_{0,\text{left}} \\ \mathbf{u}_{1,\text{bottom}} \\ \mathbf{u}_{2,\text{bottom}} + v_{0,\text{right}} \\ -v_{1,\text{left}} \\ 0 \\ v_{1,\text{right}} \\ -\mathbf{u}_{0,\text{top}} - v_{2,\text{left}} \\ -\mathbf{u}_{1,\text{top}} \\ -\mathbf{u}_{2,\text{top}} + v_{2,\text{right}} \end{pmatrix}. \quad (3.40)$$

Interior nodes (e.g.,  $\omega_4$ ) receive no boundary contribution. The resulting boundary vector is exact when the boundary velocities are known analytically.



**FIGURE 3.7** Assembly of boundary contributions. Orange arrows show tangential velocity components along the dual boundary, oriented upwards or rightwards. Green arrows denote counterclockwise vorticity DoFs. The sign pattern in  $\tilde{\mathcal{B}}^1$  arises from reconciling these two conventions.

3

#### 3.5.4.2. Case 2: Boundary Velocities Not Exactly Known

When the exact velocity is not available on  $\partial\Omega$ , the reconstructed velocity field  $\mathbf{u}_h$  is used to approximate the tangential boundary velocities. This situation is typical for experimental PTV data, where the measurement domain may not extend to the boundary or the boundary values are noisy. The boundary integrals (3.35)–(3.38) are evaluated using the 1-form basis functions  $\Psi^{(1)}$  restricted to the boundary nodes, with the velocity coefficients  $\mathbf{u}_k$  from the least-squares solve (Section 3.3.6).

Since  $\mathbf{u}_h$  is already computed from the constrained least-squares problem, the boundary values are self-consistent with the interior reconstruction. However, this introduces an additional approximation error beyond the interior discretisation error: the boundary contribution is no longer exact, and its accuracy depends on how well the polynomial reconstruction captures the true velocity near the boundary. For well-resolved interior fields, the boundary approximation error is expected to be comparable to the interior discretisation error established in Section 2.6..

#### 3.5.5. Multi-Element Assembly

In the multi-element setting, the local vorticity system (3.34) is continuous across all elements. Stacking all element contributions gives the block-

diagonal system:

$$\underbrace{\text{blkdiag}(M_0^0, \dots, M_{K-1}^0)}_{\mathcal{M}} \boldsymbol{\omega} = \underbrace{\text{blkdiag}((E_0^{1,0})^T M_0^1, \dots, (E_{K-1}^{1,0})^T M_{K-1}^1)}_{\mathcal{C}} \mathbf{u} + \tilde{\mathcal{B}}^1, \quad (3.41)$$

where  $\mathcal{M}$  is block-diagonal with the element 0-form mass matrices,  $\mathcal{C}$  is block-diagonal with the curl-mass blocks, and  $\tilde{\mathcal{B}}^1$  stacks the element boundary contributions.


### 3

#### 3.5.5.1. Interface Continuity via Constraints

Across each internal interface, the vorticity must be continuous (since  $\omega$  is a 0-form with nodal values). Let  $N_\omega$  be the sparse constraint matrix whose rows encode differences of matching interface DoFs; each row has a +1 at one element's node index and -1 at the neighbouring element's matched node index:

$$N_\omega \boldsymbol{\omega} = \mathbf{0}. \quad (3.42)$$

As illustrated in Figure 3.8, not all interface equalities are independent: closed cycles of interfaces induce one redundant constraint. For this reason, one constraint per closed interface loop is omitted (e.g., the connection between  $\omega_{20}$  and  $\omega_{27}$  in Figure 3.8), preventing rank deficiency in the constraint set.



**FIGURE 3.8** Four-element ( $K_x = K_y = 2$ ) layout with continuing vorticity indices  $\omega_i$  for each element ( $N = M = 3$ ). Double-headed red arrows  $\gamma_j$  indicate interface equalities to be enforced between matching nodal vorticities. One interface in each closed loop is intentionally left unconstrained to avoid redundancy and ensure a nonsingular constraint system.

### 3.5.5.2. KKT System with Lagrange Multipliers

Enforcing (3.42) via Lagrange multipliers  $\gamma$  yields the symmetric saddle-point system:

$$\begin{pmatrix} \mathcal{M} & N_\omega^T \\ N_\omega & 0 \end{pmatrix} \begin{pmatrix} \boldsymbol{\omega} \\ \boldsymbol{\gamma} \end{pmatrix} = \begin{pmatrix} \mathcal{C}\mathbf{u} + \tilde{\mathcal{B}}^1 \\ \mathbf{0} \end{pmatrix}. \quad (3.43)$$

The matrix  $\mathcal{M}$  is block-diagonal SPD, while the coupling through  $N_\omega$  enforces inter-element continuity. The formulation retains element-wise assembly: each element contributes its own mass matrix and right-hand side, and the interface constraints couple only shared boundary nodes.

### 3.5.6. Vorticity Reconstruction on Fine Grid

The continuous vorticity field is obtained by interpolating the discrete nodal values  $\omega_k$  within each element using the 0-form basis functions (2.28):

$$\omega(x, y) = \Psi^{(0)}(\xi, \eta) \boldsymbol{\omega}_k, \quad (3.44)$$

where  $\omega_k$  contains the nodal vorticity DoFs of element  $k$ . As for the streamfunction (Section 3.4.3.), the 0-form basis transforms trivially under an affine map (2.43): no Jacobian scaling is required. The reconstruction is performed independently on each element over the  $N_{\text{fine}} \times N_{\text{fine}}$  Gauss–Lobatto–Legendre grid and assembled into a global array of resolution  $K_x N_{\text{fine}} \times K_y N_{\text{fine}}$ .

## 3

### 3.6. Total and Static Pressure Reconstruction

This section derives the pressure field from the reconstructed velocity and vorticity using the incompressible Navier–Stokes equations. The total pressure is obtained by solving a Poisson-type equation cast as a saddle-point system in the mimetic framework. The static pressure is then recovered by subtracting the dynamic pressure.

#### 3.6.1. Governing Equation

We begin from the incompressible steady Navier–Stokes equations in non-dimensional form:

$$(\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \frac{1}{\text{Re}} \nabla^2 \mathbf{u}, \quad (3.45)$$

where  $\text{Re} = U_\infty L / \nu$  is the Reynolds number, with  $U_\infty$  a characteristic velocity,  $L$  a characteristic length, and  $\nu$  the kinematic viscosity. All quantities are non-dimensionalised accordingly, and body forces are assumed absent.

The advective term is rewritten using the Lamb vector identity:

$$(\mathbf{u} \cdot \nabla)\mathbf{u} = -\mathbf{u} \times \boldsymbol{\omega} + \frac{1}{2} \nabla |\mathbf{u}|^2, \quad (3.46)$$

where  $\boldsymbol{\omega} = \nabla \times \mathbf{u}$  is the vorticity. Substituting into (3.45) and defining the total pressure

$$\tilde{P} = p + \frac{1}{2} |\mathbf{u}|^2, \quad (3.47)$$

gives

$$\mathbf{u} \times \boldsymbol{\omega} = \nabla \tilde{P} - \frac{1}{\text{Re}} \nabla^2 \mathbf{u}. \quad (3.48)$$

Taking the divergence of both sides yields the pressure Poisson equation. For divergence-free flow,  $\nabla^2 \mathbf{u} = -\nabla \times \boldsymbol{\omega}$ , and since the divergence of any curl is identically zero, the viscous term vanishes:  $\nabla \cdot \nabla^2 \mathbf{u} = -\nabla \cdot (\nabla \times \boldsymbol{\omega}) = 0$ . This gives:

$$\Delta \tilde{P} = \nabla \cdot (\mathbf{u} \times \boldsymbol{\omega}). \quad (3.49)$$

This is a scalar Poisson equation for the total pressure, driven by the divergence of the vorticity force term  $\mathbf{C} = \mathbf{u} \times \boldsymbol{\omega}$ . In two dimensions,  $\mathbf{C}$  reduces to:

$$\mathbf{C} = \mathbf{u} \times \boldsymbol{\omega} = \begin{pmatrix} v\omega \\ -u\omega \end{pmatrix}, \quad (3.50)$$

where  $\omega$  is the scalar vorticity (Section 3.5.).

### 3.6.2. Weak Formulation

To cast the pressure Poisson equation (3.49) as a first-order system suitable for the mimetic framework, we introduce the auxiliary variable  $\mathbf{q} = \nabla \tilde{P}$  (the pressure gradient). Like the velocity,  $\mathbf{q}$  is a primal 1-form discretised in  $H(\text{div}; \Omega)$ . Taking the divergence gives  $\nabla \cdot \mathbf{q} = \Delta \tilde{P}$ , so the system becomes:

$$\mathbf{q} - \nabla \tilde{P} = \mathbf{0}, \quad (3.51)$$

$$\nabla \cdot \mathbf{q} = \nabla \cdot \mathbf{C}. \quad (3.52)$$

To derive the weak formulation, we multiply (3.51) by a 1-form test function  $\mathbf{v}$  and (3.52) by a 2-form test function  $g$ , then integrate over  $\Omega$ . Applying integration by parts to the  $\nabla \tilde{P}$  term:

$$\int_{\Omega} \mathbf{v} \cdot \nabla \tilde{P} \, d\Omega = \int_{\partial\Omega} (\mathbf{v} \cdot \mathbf{n}) \tilde{P} \, dS - \int_{\Omega} (\nabla \cdot \mathbf{v}) \tilde{P} \, d\Omega. \quad (3.53)$$

The boundary term in (3.53) vanishes when  $\tilde{P} = 0$  on  $\partial\Omega$ . This homogeneous Dirichlet condition on the total pressure eliminates the surface integral and closes the weak formulation. Since the total pressure is determined only up to a constant, prescribing  $\tilde{P} = 0$  on the boundary also fixes this constant. In practice, for the manufactured solution tests (Chapter 5), this condition is consistent with the analytical fields. For experimental PTV data, where no pressure measurement is available at the boundary, a reference pressure is instead prescribed at a single cell (Section 3.6.4.).

With the boundary term eliminated, the weak formulation reads:

$$\int_{\Omega} \mathbf{v} \cdot \mathbf{q} \, d\Omega + \int_{\Omega} (\nabla \cdot \mathbf{v}) \tilde{P} \, d\Omega = 0, \quad (3.54)$$

$$\int_{\Omega} g (\nabla \cdot \mathbf{q}) \, d\Omega = \int_{\Omega} g (\nabla \cdot \mathbf{C}) \, d\Omega. \quad (3.55)$$

Note that integration by parts is not required for the second equation, since the divergence of both  $\mathbf{q}$  and  $\mathbf{C}$  yields scalar quantities (2-forms) that are directly tested against  $g$ .

### 3.6.3. Mimetic Discretisation

The pressure gradient  $\mathbf{q}$  and the test function  $\mathbf{v}$  are discretised in the 1-form basis (2.29)–(2.30), while the total pressure  $\tilde{\mathbf{P}}$  and the test function  $g$  are discretised in the 2-form basis (2.31). Following the same procedure as for the vorticity (Section 3.5.), using Corollary 1 of [31] to replace derivatives of basis functions with incidence matrices, the weak equations (3.54)–(3.55) discretise to:

$$M^1 \mathbf{q} + (E^{2,1})^T \tilde{\mathbf{P}} = \mathbf{0}, \quad (3.56)$$

$$E^{2,1} \mathbf{q} = E^{2,1} \mathbf{C}, \quad (3.57)$$

where  $\tilde{\mathbf{P}} = M^2 \mathbf{P}$  is the  $L^2$  inner-product representation of the total pressure coefficients (Section 2.5.4.). This is written in matrix form as the saddle-point system:

$$\begin{pmatrix} M^1 & (E^{2,1})^T \\ E^{2,1} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{q} \\ \tilde{\mathbf{P}} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ E^{2,1} \mathbf{C} \end{pmatrix}. \quad (3.58)$$

The mass matrices  $M^1$  and  $M^2$  are defined in Section 2.5.4., with the physical-domain transformations (2.48) and (2.49) for affine maps.

The right-hand side of (3.58) requires the primal 1-form representation of  $\mathbf{C} = \mathbf{u} \times \boldsymbol{\omega}$ . This is constructed in two steps.

First, the dual-space representation  $\tilde{\mathbf{C}}$  is assembled by evaluating the pointwise products  $\mathbf{v}\boldsymbol{\omega}$  and  $\mathbf{u}\boldsymbol{\omega}$  at the Gauss–Lobatto–Legendre quadrature points and integrating against the 1-form basis functions. For the  $x$ -component:

$$\tilde{C}_{x,i} = \sum_q \Psi_{x,i}^{(1)}(q) \cdot (\mathbf{v}\boldsymbol{\omega})_q \cdot w_q \cdot \det J, \quad (3.59)$$

where  $(\mathbf{v}\boldsymbol{\omega})_q = (\Psi_v^{(1)} \mathbf{v}_k) * (\Psi_\omega^{(0)} \boldsymbol{\omega}_k)$  is the pointwise product of the reconstructed velocity and vorticity at quadrature point  $q$ ,  $w_q$  is the quadrature weight, and  $*$  denotes the Hadamard (elementwise) product. Similarly, the  $y$ -component:

$$\tilde{C}_{y,i} = - \sum_q \Psi_{y,i}^{(1)}(q) \cdot (\mathbf{u}\boldsymbol{\omega})_q \cdot w_q \cdot \det J. \quad (3.60)$$

The assembled vector  $\tilde{\mathbf{C}} = (\tilde{C}_x, \tilde{C}_y)^T$  lives in the dual space.

Second, the primal representation is recovered by applying the inverse mass matrix:

$$\mathbf{C} = (M^1)^{-1} \tilde{\mathbf{C}}. \quad (3.61)$$

This ensures that  $\mathbf{C}$  is a primal 1-form, compatible with the divergence operator  $E^{2,1}$  in (3.58). The transformation from dual to primal space through  $(M^1)^{-1}$  is the discrete analogue of the inverse Hodge star.

### 3.6.4. Multi-Element Assembly

In the multi-element setting, the local saddle-point system (3.58) is replicated across all elements. The global mass matrices and incidence matrices are assembled using the Kronecker product:

$$M_t^1 = I_K \otimes M^1, \quad M_t^2 = I_K \otimes M^2, \quad E_t^{2,1} = I_K \otimes E^{2,1}, \quad (3.62)$$

where  $I_K$  is the  $K \times K$  identity matrix with  $K = K_x \cdot K_y$ .

The pressure gradient  $\mathbf{q}$  is a 1-form and therefore requires normal flux continuity across element interfaces, enforced through a continuity constraint matrix  $N_L$  with the same structure as for the velocity (Section 3.3.4.). Each row of  $N_L$  contains +1 and -1 entries equating matching pressure gradient flux DoFs on shared edges.

The complete multi-element saddle-point system, including the interface continuity constraints via Lagrange multipliers  $\lambda$ , is:

$$\begin{pmatrix} M_t^1 & N_L^T & (E_t^{2,1})^T \\ N_L & 0 & 0 \\ E_t^{2,1} & 0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{q} \\ \lambda \\ \tilde{\mathbf{P}} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ E_t^{2,1} \mathbf{C} \end{pmatrix}. \quad (3.63)$$

The Lagrange multiplier  $\lambda$  enforces inter-element continuity of the pressure gradient, and  $\tilde{\mathbf{P}}$  enforces the divergence constraint cell by cell.

#### 3.6.4.1. Reference Pressure

When the boundary condition  $\tilde{\mathbf{P}} = 0$  is not imposed on the entire boundary, the total pressure is determined only up to a constant. To remove this null-space, a reference pressure is prescribed by constraining  $\tilde{\mathbf{P}}$  in one cell to a known value (e.g.,  $\tilde{\mathbf{P}}_0 = 0$ ). In practice, this is implemented by adding a single row to the constraint system or by fixing the corresponding degree of freedom directly.

#### 3.6.4.2. Recovery of Primal Pressure

The solution of (3.63) yields the dual pressure  $\tilde{\mathbf{P}} = M^2 \mathbf{P}$ . The primal (volume-weighted) pressure is recovered by:

$$\mathbf{P} = (M_t^2)^{-1} \tilde{\mathbf{P}}. \quad (3.64)$$

### 3.6.5. Static Pressure Recovery

The static pressure  $p$  is obtained from the total pressure  $\tilde{\mathbf{P}}$  and the reconstructed velocity by subtracting the dynamic pressure:

$$p = \tilde{\mathbf{P}} - \frac{1}{2} |\mathbf{u}|^2. \quad (3.65)$$

The dynamic pressure  $\frac{1}{2}|\mathbf{u}|^2$  is evaluated pointwise on the fine reconstruction grid from the velocity field computed in Section 3.4.1.. Both the total pressure field (reconstructed via the 2-form basis  $\Psi^{(2)}$  (2.31)) and the velocity field (reconstructed via the 1-form basis  $\Psi^{(1)}$  (2.29)–(2.30)) are evaluated on the same  $N_{\text{fine}} \times N_{\text{fine}}$  Gauss–Lobatto–Legendre grid per element, so the subtraction is performed pointwise without interpolation.

### 3

#### 3.7. Algorithm Summary

This section collects the complete reconstruction procedure into four pseudocode algorithms. Each algorithm references the equations derived in Sections 3.3.–3.6.; no new material is introduced.

In practice, many quantities are computed once and reused across algorithms. The one-dimensional Lagrange basis functions  $h_i(\xi)$  and edge basis functions  $e_i(\xi)$ , together with the tensor-product shape functions  $\Psi^{(0)}$ ,  $\Psi^{(1)}$ , and  $\Psi^{(2)}$ , are evaluated once on the Gauss–Lobatto–Legendre nodes and stored. Similarly, the mass matrices  $M^0$ ,  $M^1$ , and  $M^2$  are assembled once per element during the vorticity computation (Algorithm 3) and reused directly in the pressure computation (Algorithm 4). The incidence matrices  $E^{1,0}$  and  $E^{2,1}$  are purely topological and constructed once for a given polynomial order. The Jacobian  $J$  and its determinant are likewise computed once per element. The algorithms below are therefore presented as sequential stages, but the shared quantities are not recomputed between them.

**Algorithm 1: Constrained Multi-Element Velocity Reconstruction**

**Input:** Experimental velocity data  $\mathbf{u}^d$ , mesh bounds  
 $[x_L, x_R] \times [y_L, y_R]$ , mesh sizes  $K_x, K_y$ , polynomial orders  
 $P_\xi, P_\eta$

**Output:** Edge-based velocity DoFs  $\mathbf{u}$ , divergence field,  
streamfunction  $\psi$

**Domain decomposition and local evaluation;**

Partition domain into  $K_x \times K_y$  elements (Section 3.2.);

**for each element k do**

Map data points  $(x, y) \rightarrow (\xi, \eta)$  via the affine mapping (2.3);  
Evaluate 1-form basis  $\Psi_k^{(1)}(\xi, \eta)$  at mapped locations  
(Section 3.3.1.);  
Apply Piola transformation:  $\Psi_k^{(1)}(x, y) = \frac{1}{\det J} J \Psi_k^{(1)}(\xi, \eta)$  (2.44);

**Global assembly and constrained solve;**

Assemble block-diagonal matrix  $A$  from local  $\Psi_k^{(1)}$  (Section 3.3.3.);  
Build continuity constraint matrix  $N_L$  (Section 3.3.4.);  
Build divergence constraint  $E_t^{2,1} = I_K \otimes E^{2,1}$  (Section 3.3.5.);  
Solve the KKT system (3.23) for  $(\mathbf{u}, \lambda, \zeta)$ ;

**Postprocessing;**

**for each element k do**

Evaluate  $\Psi_k^{(1)}$  on  $N_{\text{fine}} \times N_{\text{fine}}$  Gauss–Lobatto–Legendre grid;  
Reconstruct velocity:  $\mathbf{u}_k^{\text{fine}} = \Psi_k^{(1)}(x, y) \mathbf{u}_k$  (Section 3.4.1.);  
Verify divergence:  $(\nabla \cdot \mathbf{u})_k^{\text{fine}} = \Psi_k^{(2)}(x, y) E^{2,1} \mathbf{u}_k$  (Section 3.4.2.);  
Reconstruct streamfunction via Algorithm 2 (Section 3.4.3.);  
Evaluate:  $\psi_k^{\text{fine}} = \Psi_k^{(0)}(\xi, \eta) \psi_k$ ;

**Algorithm 2:** Streamfunction Marching**Input:** Incidence matrix  $E^{1,0} \in \mathbb{R}^{m \times n}$ , flux vector  $\mathbf{u} \in \mathbb{R}^m$ **Output:** Streamfunction DoFs  $\boldsymbol{\psi} \in \mathbb{R}^n$ **Initialisation;**Set  $\boldsymbol{\psi} \leftarrow \mathbf{0}$ ,  $\mathbf{k} \leftarrow \text{false}^{1 \times n}$ ;Set starting value:  $\psi_0 \leftarrow s$ ,  $k_0 \leftarrow \text{true}$ ;**March across edges;**changed  $\leftarrow \text{true}$ ;**while** *changed* **do**    changed  $\leftarrow \text{false}$ ;    **for** edge  $e = 0$  to  $m - 1$  **do**        Extract endpoints  $(i, j)$  and signs  $(\sigma_i, \sigma_j)$  from row  $e$  of  $E^{1,0}$ ;        **if** exactly one of  $k_i, k_j$  is true **then**            **if**  $k_i$  **then**                 $\psi_j \leftarrow (\mathbf{u}_e - \sigma_i \psi_i) / \sigma_j$ ,  $k_j \leftarrow \text{true}$ ;            **else**                 $\psi_i \leftarrow (\mathbf{u}_e - \sigma_j \psi_j) / \sigma_i$ ,  $k_i \leftarrow \text{true}$ ;            changed  $\leftarrow \text{true}$ ;**Multi-element extension;**

Process elements left-to-right, then row above (Section 3.4.3.);

Seed:  $\psi_0^{\text{right}} = \psi_{N-1}^{\text{left}}$ ; $\psi_0^{\text{top}} = \psi_{(M-1)N}^{\text{bottom}}$ ;

**Algorithm 3: Constrained Multi-Element Vorticity Computation**

**Input:** Velocity DoFs  $\mathbf{u}$ , mesh bounds  $[x_L, x_R] \times [y_L, y_R]$ , mesh sizes  $K_x, K_y$ , polynomial orders  $p_\xi, p_\eta$

**Output:** Global vorticity DoFs  $\boldsymbol{\omega} \in \mathbb{R}^{K^{n_0}}$  (and interface multipliers  $\boldsymbol{\gamma}$ )

**Build mass matrices and right-hand side;**

Compute  $M^0$  and  $M^1$  on the physical domain (2.47)–(2.48);

Assemble boundary contribution  $\tilde{B}^1$  (Section 3.5.4.);

Compute right-hand side:  $\mathbf{g} \leftarrow (E^{1,0})^T M^1 \mathbf{u} + \tilde{B}^1$  (3.34);

**Interface continuity constraints;**

Build vorticity continuity matrix  $N_\omega$  (Section 3.5.5.);

Remove one constraint per closed interface loop;

**Assemble and solve KKT system;**

$$\text{Solve (3.43): } \begin{pmatrix} \mathcal{M} & N_\omega^T \\ N_\omega & 0 \end{pmatrix} \begin{pmatrix} \boldsymbol{\omega} \\ \boldsymbol{\gamma} \end{pmatrix} = \begin{pmatrix} \mathbf{g} \\ \mathbf{0} \end{pmatrix};$$

**Postprocessing;**

**for each element  $k$  do**

└ Reconstruct:  $\omega_k^{\text{fine}} = \Psi_k^{(0)}(\xi, \eta) \omega_k$  (Section 3.5.6.);

**Algorithm 4:** Constrained Multi-Element Total Pressure Computation**Input:** Velocity DoFs  $\mathbf{u}$ , vorticity DoFs  $\boldsymbol{\omega}$ , mesh bounds $[\mathbf{x}_L, \mathbf{x}_R] \times [\mathbf{y}_L, \mathbf{y}_R]$ , mesh sizes  $K_x, K_y$ , polynomial orders $p_\xi, p_\eta$ **Output:** Pressure gradient DoFs  $\mathbf{q}$ , total pressure DoFs  $\tilde{\mathbf{P}}$  (and interface multipliers  $\boldsymbol{\lambda}$ )**Build mass matrices;**Compute  $M^1$  and  $M^2$  on the physical domain (2.48)–(2.49);**Construct vorticity force term;****for each element  $k$  do**Evaluate  $\Psi^{(0)}, \Psi_x^{(1)}, \Psi_y^{(1)}$  at Gauss–Lobatto–Legendre quadrature points;Compute pointwise products:  $(v\omega)_q, (u\omega)_q$ ;Integrate to dual space:  $\tilde{C}_{x,i}$  (3.59),  $\tilde{C}_{y,i}$  (3.60);Assemble global  $\tilde{\mathbf{C}}$  and transform to primal space:

$$\mathbf{C} = (M^1)^{-1} \tilde{\mathbf{C}} \quad (3.61);$$

**Interface continuity and global assembly;**Build pressure gradient continuity matrix  $N_L$  (same structure as Section 3.3.4.);Form global matrices:  $M_t^1 = I_K \otimes M^1, M_t^2 = I_K \otimes M^2,$ 

$$E_t^{2,1} = I_K \otimes E^{2,1} \quad (3.62);$$

**Assemble and solve KKT system;**Constrain  $\tilde{\mathbf{P}}$  in one cell to fix the reference pressure (Section 3.6.4.);

$$\text{Solve (3.63):} \quad \begin{pmatrix} M_t^1 & N_L^T & (E_t^{2,1})^T \\ N_L & 0 & 0 \\ E_t^{2,1} & 0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{q} \\ \boldsymbol{\lambda} \\ \tilde{\mathbf{P}} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ E_t^{2,1} \mathbf{C} \end{pmatrix};$$

**Postprocessing;**Recover primal pressure:  $\mathbf{P} = (M_t^2)^{-1} \tilde{\mathbf{P}} \quad (3.64);$ **for each element  $k$  do**Reconstruct pressure gradient:  $\mathbf{q}_k^{\text{fine}} = \Psi_k^{(1)}(x, y) \mathbf{q}_k$ ;Reconstruct total pressure:  $P_k^{\text{fine}} = \Psi_k^{(2)}(x, y) P_k$ ;Compute static pressure:  $p_k^{\text{fine}} = P_k^{\text{fine}} - \frac{1}{2} |\mathbf{u}_k^{\text{fine}}|^2 \quad (3.65);$

# Chapter 4

## Verification and Validation Setup

Before the mimetic spectral element reconstruction method can be applied to experimental flow measurements, it must first be verified and validated. Verification confirms that the numerical implementation correctly solves the discrete equations formulated in Chapter 3 and achieves the convergence rates established in Section 2.6. Validation assesses whether the method produces physically meaningful results when applied to real experimental data. This chapter describes the setup for both studies; the results are presented in Chapter 5.

## 4

The verification approach uses the Method of Manufactured Solutions (Section 4.1.). An analytical velocity field is prescribed from which all associated quantities (vorticity, streamfunction, and total pressure) can be computed exactly. The analytical fields are sampled at random points to emulate experimental measurements, and the reconstruction output is compared against the known solution. This enables systematic testing of spectral and algebraic convergence and divergence-free constraint satisfaction.

The validation data originate from a three-dimensional Lagrangian Particle Tracking campaign performed in [11] at Delft University of Technology. Flow seeded with helium-filled soap bubbles was measured around a cube of side length  $H = 120$  mm mounted on a flat plate in a low-speed open-jet wind tunnel at  $U_\infty = 10$  m s<sup>-1</sup> ( $Re_H \approx 80\,000$ ), using a seven-camera imaging system with object-aware monolithic particle tracking (OA<sub>7</sub>-STB). A schematic of the experimental setup is shown in Figure 4.2. The raw three-dimensional particle tracks within a spanwise slab  $-\delta \leq y \leq \delta$  centred at the mid-plane  $y = 0$  are projected onto the  $x$ - $z$  plane to produce scattered two-dimensional velocity data, which serve as input to the reconstruction method. The reconstructed fields are compared against an ensemble-averaged reference field obtained by second-order polynomial regression. In the spanwise mid-plane, the cube presents a square cross-section to the flow, whose sharp edges enforce fixed separation points and produce strong velocity gradients, recirculation, and reattachment. Three validation regions of increasing flow complexity are defined in Section 4.2..

### 4.1. Method Verification via Manufactured Solutions

The Mimetic Spectral Element Method developed in Chapter 3 relies on high-order polynomial basis functions to discretise the velocity field and preserve the algebraic structure of the de Rham complex at the discrete level. Before applying this method to experimental data, the numerical

implementation must be verified. Verification testing uses manufactured solutions to isolate the accuracy of the implementation. By constructing analytical flow fields where all quantities can be computed exactly, the numerical errors can be measured and compared against the theoretical predictions from Section 2.6..

#### 4.1.1. Purpose and Strategy

To verify the implementation, the Method of Manufactured Solutions is employed. An analytical velocity field is chosen from which the corresponding vorticity, streamfunction, and total pressure fields can be computed exactly. The analytical solution is sampled at random discrete points, simulating experimental measurements. The reconstruction algorithm processes these synthetic data, and the resulting numerical fields are compared directly against the known analytical solution.

The synthetic data are processed through the reconstruction algorithm, and the output is compared against the known solution. Different polynomial orders are tested to observe spectral convergence behaviour, and different mesh resolutions demonstrate the expected algebraic convergence rates. The multi-element configuration tests whether interface continuity constraints are properly imposed. Throughout all cases, the divergence-free constraint should be satisfied to near machine precision, and the reconstructed fields should maintain the mathematical relationships derived in Chapter 3.

#### 4.1.2. Analytical Test Case: Trigonometric Velocity Field

The manufactured solution is constructed from trigonometric functions that provide a smooth, infinitely differentiable velocity field. This smoothness is essential for observing the spectral convergence rates characteristic of high-order polynomial approximations. The velocity field is defined as

$$u(x, y) = A \sin(\alpha x) \cos(\beta y), \quad (4.1)$$

$$v(x, y) = B \cos(\gamma x) \sin(\delta y), \quad (4.2)$$

where  $A$ ,  $B$ ,  $\alpha$ ,  $\beta$ ,  $\gamma$ , and  $\delta$  are parameters that control the spatial frequency and amplitude. For the velocity field to satisfy the incompressibility condition, the divergence must vanish:

$$\nabla \cdot \mathbf{u} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0. \quad (4.3)$$

Substituting (4.1)–(4.2) yields

$$A\alpha \cos(\alpha x) \cos(\beta y) + B\delta \cos(\gamma x) \cos(\delta y) = 0,$$

which is satisfied identically when  $A\alpha = -B\delta$  and  $\alpha = \gamma$ ,  $\beta = \delta$ . The parameter values

$$A = 2\pi, \quad B = -2\pi, \quad \alpha = \beta = \gamma = \delta = 2\pi \quad (4.4)$$

satisfy this condition, and the velocity components become

$$u(x, y) = 2\pi \sin(2\pi x) \cos(2\pi y), \quad (4.5)$$

$$v(x, y) = -2\pi \cos(2\pi x) \sin(2\pi y). \quad (4.6)$$

The vorticity is computed from the curl of the velocity field:

$$\omega = \nabla \times \mathbf{u} = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}. \quad (4.7)$$

Evaluating the derivatives,

$$\begin{aligned} \frac{\partial v}{\partial x} &= 4\pi^2 \sin(2\pi x) \sin(2\pi y), \\ \frac{\partial u}{\partial y} &= -4\pi^2 \sin(2\pi x) \sin(2\pi y), \end{aligned}$$

gives

$$\omega(x, y) = 8\pi^2 \sin(2\pi x) \sin(2\pi y). \quad (4.8)$$

The streamfunction  $\psi$  is related to the velocity components through

$$\mathbf{u} = -\frac{\partial \psi}{\partial y}, \quad v = \frac{\partial \psi}{\partial x}. \quad (4.9)$$

Integrating the first relation and setting the integration constant to zero yields

$$\psi(x, y) = \sin(2\pi x) \sin(2\pi y). \quad (4.10)$$

This can be verified through the compatibility relation  $\nabla^2 \psi = -\omega$ :

$$\begin{aligned} \nabla^2 \psi &= -(2\pi)^2 \sin(2\pi x) \sin(2\pi y) - (2\pi)^2 \sin(2\pi x) \sin(2\pi y) \\ &= -8\pi^2 \sin(2\pi x) \sin(2\pi y) = -\omega. \end{aligned}$$

For the total pressure field, the Poisson equation from Section 3.6. is used:

$$\Delta \tilde{P} = \nabla \cdot (\mathbf{u} \times \boldsymbol{\omega}). \quad (4.11)$$

The total pressure satisfying this equation is

$$\tilde{P}(x, y) = 2\pi^2 \cos(4\pi x) \sin^2(2\pi y) + \pi^2 \cos(4\pi y). \quad (4.12)$$

The pressure gradient components follow from differentiation:

$$q_x = \frac{\partial \tilde{P}}{\partial x} = -8\pi^3 \sin(4\pi x) \sin^2(2\pi y), \quad (4.13)$$

$$q_y = \frac{\partial \tilde{P}}{\partial y} = -8\pi^3 \sin^2(2\pi x) \sin(4\pi y). \quad (4.14)$$

With all fields defined analytically, exact values can be computed at any point in the domain, providing the reference solution against which the numerical reconstruction is compared.

### 4.1.3. Test Configurations

The verification tests are designed to systematically assess the accuracy and convergence behaviour of the reconstruction method across different computational setups.

4

#### 4.1.3.1. Domain Configurations

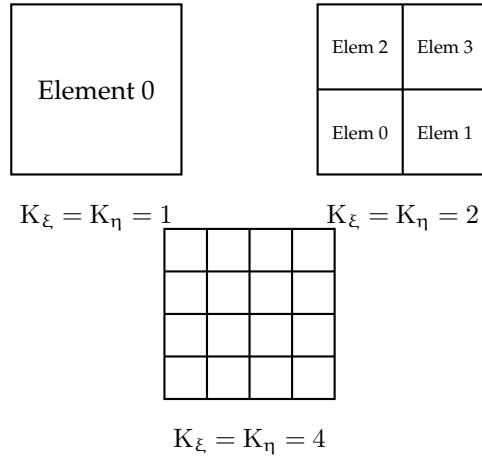
Two domains are considered. The reference element uses the standard computational domain  $[-1, 1] \times [-1, 1]$ , where basis functions are naturally defined. This configuration tests the method without any geometric transformation effects. The physical domain uses  $[0, 1] \times [0, 1]$ , requiring the Piola transformation (Section 2.5.5.) to map between reference and physical coordinates. This tests the correct implementation of the Jacobian scaling.

#### 4.1.3.2. Mesh Configurations

Both single-element and multi-element configurations are tested. The single-element case ( $K_\xi = K_\eta = 1$ ) covers the entire domain with a single spectral element, isolating the approximation properties of the high-order polynomial basis functions without interface coupling. The multi-element cases use structured Cartesian meshes with  $K_\xi = K_\eta = K$  for  $K \in \{2, 4, 7\}$ . These configurations test the continuity constraints enforced through the Lagrange multiplier matrix  $N_L$  (Section 3.3.4.) and verify that flux matching across element interfaces is correctly handled. Figure 4.1 illustrates examples of single-element and multi-element configurations.

#### 4.1.3.3. Discretisation Parameters

The polynomial order  $p$  determines the number of nodes in each direction ( $N = M = p + 1$ ) and the degree of the basis functions  $h_i(\xi)$  and  $e_i(\xi)$  described in Section 2.3.. For the  $p$ -refinement study,  $p$  ranges from 3 to 14. For  $h$ -refinement,  $p$  is fixed and the number of elements  $K$  is varied.



4

**FIGURE 4.1** Mesh configurations for verification tests. Top left: single element covering entire domain. Top right:  $2 \times 2$  multi-element mesh. Bottom:  $4 \times 4$  multi-element mesh.

Three data densities are used:  $n_r \in \{40\,000, 160\,000, 360\,000\}$  measurement points distributed uniformly at random over the domain. These densities range from relatively sparse to dense sampling, allowing evaluation of how data availability affects the reconstruction accuracy.

#### 4.1.4. Synthetic Data Generation

The analytical velocity field is sampled at randomly distributed points within the computational domain to simulate experimental measurements. For each element, points are generated with uniformly random coordinates within that element's spatial extent. This random distribution tests the method's ability to handle irregular data and verifies that the least-squares formulation remains stable when measurement locations do not align with the spectral element nodal structure. At each point  $(x_i, y_i)$ , the velocity components  $u(x_i, y_i)$  and  $v(x_i, y_i)$  are evaluated from the analytical solution (4.5)–(4.6) and stored as synthetic measurement data.

#### 4.1.5. Convergence Studies

The verification process examines how the reconstruction error decreases as the discretisation is refined. Two refinement strategies are employed.

##### 4.1.5.1. p-Refinement

The polynomial order  $p$  is systematically increased while the mesh configuration is held fixed. For smooth solutions such as the manufactured trigonometric velocity field, spectral element methods exhibit exponential

convergence: the error decreases as  $O(e^{-c p})$  for some constant  $c > 0$  that depends on the analyticity of the solution (Section 2.6.). Three mesh configurations ( $K = 1$ ,  $K = 4$ ,  $K = 7$ ) are tested at each data density. The convergence is visualised on a semi-log plot of the error against polynomial order.

#### 4.1.5.2. h-Refinement

The polynomial order is held fixed while the number of elements is increased, thereby decreasing the element size  $h$ . The convergence rate is assessed by plotting the error versus element size on a log-log scale, where algebraic convergence appears as a straight line whose slope indicates the order of accuracy. The theoretical rates from Table 2.3 predict a hierarchy of convergence orders: the streamfunction converges at  $O(h^{p+1})$ , the velocity  $H(\text{div}; \Omega)$ -error at  $O(h^p)$ , the  $L^2$ -errors of vorticity and total pressure at  $O(h^{p-1})$  to  $O(h^p)$ , and the  $H(\text{curl}; \Omega)$ -error of the vorticity at  $O(h^{p-2})$ .

#### 4.1.6. Error Norms

The convergence studies track five error quantities that together cover all reconstructed fields; these correspond directly to the theoretical rates in Table 2.3. The  $L^2$ -norm measures the pointwise error over the domain:

$$\|e\|_{L^2} = \sqrt{\int_{\Omega} |f_h - f_{\text{exact}}|^2 d\Omega}, \quad (4.15)$$

where  $f_h$  is the reconstructed field and  $f_{\text{exact}}$  is the analytical solution, with  $f$  representing any of the reconstructed quantities. The  $H(\text{div}; \Omega)$ -norm additionally includes the divergence error:

$$\|e\|_{H(\text{div}; \Omega)} = \sqrt{\|\mathbf{u}_h - \mathbf{u}_{\text{exact}}\|_{L^2}^2 + \|\nabla \cdot \mathbf{u}_h - \nabla \cdot \mathbf{u}_{\text{exact}}\|_{L^2}^2}. \quad (4.16)$$

Since the manufactured solution is divergence-free and the reconstruction enforces  $\nabla \cdot \mathbf{u}_h = 0$  exactly through the incidence matrix  $E^{2,1}$  (Section 2.5.2.), the divergence term should be near machine precision. The  $H(\text{curl}; \Omega)$ -norm combines the  $L^2$ -vorticity error with the  $L^2$ -error of the curl of vorticity:

$$\|e\|_{H(\text{curl}; \Omega)} = \sqrt{\|\omega_h - \omega_{\text{exact}}\|_{L^2}^2 + \|\nabla \times \omega_h - \nabla \times \omega_{\text{exact}}\|_{L^2}^2}. \quad (4.17)$$

The five error quantities evaluated are: (1) the  $H(\text{div}; \Omega)$ -norm of the velocity, (2) the  $L^2$ -norm of the streamfunction, (3) the  $L^2$ -norm of the vorticity, (4) the  $L^2$ -norm of the total pressure, and (5) the  $H(\text{curl}; \Omega)$ -norm of the vorticity. In practice, the integrals are approximated using Gauss–Lobatto–Legendre quadrature weights on the reconstruction grid.

## 4.2. Validation Setup: Surface-Mounted Cube Experiment

While the manufactured solution verification confirms that the numerical implementation is correct, validation against experimental data assesses whether the method produces physically meaningful results when applied to real flow measurements. The surface-mounted cube experiment provides a test case with complex flow physics including flow separation and recirculation. The sharp geometry change induces flow separation upstream of the cube leading edge, creating a recirculation zone with reversed flow. The separated shear layer exhibits strong vorticity gradients. These features test the method's ability to reconstruct velocity fields with strong spatial variations and regions of complex flow topology.

This section describes the experimental configuration, the procedure for extracting a two-dimensional velocity field from the three-dimensional particle tracking data, and the validation cases that will be used to assess the method's performance.

### 4.2.1. Experimental Configuration


The experimental data originate from a volumetric particle tracking campaign performed in [11] at Delft University of Technology. The experiments were conducted in a low-speed open-jet wind tunnel with a  $60 \times 60 \text{ cm}^2$  cross-section. A cube of side length  $H = 120 \text{ mm}$  was mounted on a flat plate with an elliptical leading edge, positioned downstream of the tunnel exit. The freestream velocity was  $U_\infty = 10 \text{ m s}^{-1}$ , corresponding to a Reynolds number based on the cube side length of

$$\text{Re}_H = \frac{U_\infty H}{\nu} \approx 80\,000, \quad (4.18)$$

where  $\nu$  is the kinematic viscosity of air. The turbulence intensity of the freestream was approximately 1%. A schematic of the experimental setup is shown in Figure 4.2.

The flow was seeded with helium-filled soap bubbles (HFSB) produced by a seeding rake containing 200 bubble generators upstream of the wind tunnel contraction. The mean bubble diameter was  $350 \text{ }\mu\text{m}$  and the production rate was  $6 \times 10^6 \text{ bubbles s}^{-1}$ , yielding a seeding concentration of approximately  $1.2 \text{ bubbles cm}^{-3}$  at the freestream velocity. The seeding density in the images was approximately 0.02 particles per pixel.

The imaging system comprised seven CMOS cameras ( $1024 \times 1024 \text{ px}$ ) positioned around the measurement volume, which had approximate dimensions  $40 \times 40 \times 30 \text{ cm}^3$ . Two LaVision LED illumination units provided



**FIGURE 4.2** Schematic of the experimental setup for the surface-mounted cube validation, adapted from [11]. An open-jet wind tunnel with a  $60 \times 60 \text{ cm}^2$  exit delivers a freestream velocity  $U_\infty = 10 \text{ m s}^{-1}$  over a cube of side length  $H = 120 \text{ mm}$  mounted on a flat plate. Seven CMOS cameras (C1–C7) and two LED illumination units record the motion of helium-filled soap bubbles seeded upstream. The dashed rectangle indicates the measurement volume. For the present validation, a two-dimensional velocity field is extracted from a spanwise slab centred at  $y = 0$ , yielding a two-dimensional view of the cube’s square cross-section in the  $x$ – $z$  plane.

pulsed illumination from two directions. Images were recorded at a rate of 3 kHz, and each measurement consisted of 5000 images for a duration of 1.67 s. The experimental parameters are summarised in Table 4.1.

**TABLE 4.1** Experimental parameters for the surface-mounted cube measurement [11].

Parameter	Symbol	Value
Cube side length	$H$	120 mm
Freestream velocity	$U_\infty$	$10 \text{ m s}^{-1}$
Reynolds number	$Re_H$	$\approx 80\,000$
Measurement volume		$40 \times 40 \times 30 \text{ cm}^3$
Number of cameras		7
Camera resolution		$1024 \times 1024 \text{ px}$
Acquisition frequency		3000 Hz
Images per run		5000
Seeding concentration		$1.2 \text{ bubbles cm}^{-3}$

The particle images were processed using the shake-the-box (STB) algorithm [2] available in the LaVision DaVis 11 software. Object-aware mono-

lithic particle reconstruction (OA<sub>7</sub>-STB) was employed, which uses the registered position of the cube to identify lines of sight obstructed by the object for each camera. This allows particle triangulation to proceed with the maximum number of unobstructed cameras at each point in the measurement volume, avoiding the measurement volume erosion that occurs when a monolithic imaging system requires all cameras to observe every particle [37].

#### 4.2.2. From 3D Particle Tracks to 2D Velocity Field

The reconstruction method developed in Chapter 3 operates on two-dimensional velocity data. Two distinct data sets are derived from the three-dimensional particle tracks produced by the STB algorithm: the raw scattered measurement data that serves as input to the reconstruction, and a reference field obtained by ensemble averaging that is used for comparison.

4

The input data for the reconstruction method consists of the raw particle track positions and velocities, without any regression or interpolation applied. A spanwise slab of thickness  $2\delta$  centred at the mid-plane  $y = 0$  is extracted from the full three-dimensional track data. All particle positions and velocities within this slab are projected onto the  $x$ - $z$  plane, producing a set of scattered two-dimensional velocity measurements  $\mathbf{u}^d = \{(x_i, z_i, u_i, w_i)\}_{i=1}^{N^d}$  with  $x$ -velocity component  $u$  and  $z$ -velocity component  $w$ . The slab half-thickness  $\delta$  is chosen per validation case to balance two competing requirements: a thin slab minimises out-of-plane contamination, making the flow more nearly two-dimensional, while a thick slab includes more particle tracks, improving the conditioning of the least-squares system. The specific values are reported alongside the results in Chapter 5. Within the slab, the flow is approximately uniform in the spanwise direction and edge effects from the tunnel sidewalls are negligible. Because the measurement locations originate from individual particle tracks, the data points  $(x_i, z_i)$  are irregularly distributed, similar to the random distributions used in the verification tests.

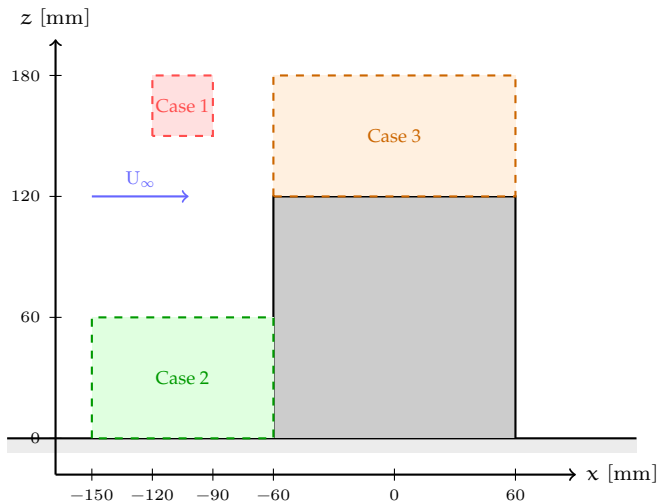
The reference field for comparison is obtained by ensemble-averaging the particle tracks onto a Cartesian grid following the procedure described in [38]. A bin size of  $10 \times 10 \times 10 \text{ mm}^3$  is used with an overlap of 75%, resulting in a grid spacing of 2.5 mm. Within each bin, the velocity is represented by a second-order polynomial regression over all tracks passing through that bin. This produces a time-averaged three-dimensional velocity field on a structured grid, from which the same spanwise slab is extracted to obtain a two-dimensional reference field in the  $x$ - $z$  plane.

The reconstructed fields from the Mimetic Spectral Element Method are compared against this reference to assess the quality of the reconstruction.

The coordinate system is defined with  $x$  as the (streamwise) horizontal direction and  $z$  as the (wall-normal) vertical direction. The origin is placed at the centre of the cube base, so that the cube occupies the region  $-60 \leq x \leq 60$  mm,  $0 \leq z \leq 120$  mm in the  $x$ - $z$  plane.

### 4.2.3. Validation Cases

The reconstruction method is validated against three distinct regions within the surface-mounted cube, each presenting different flow characteristics and challenges. These cases are selected to test the method's performance across a range of flow conditions. The locations of the three validation regions within the measurement domain are shown in Figure 4.3.



**FIGURE 4.3** Validation case regions in the  $x$ - $z$  plane. The cube cross-section (grey,  $120 \times 120$  mm) occupies  $-60 \leq x \leq 60$  mm,  $0 \leq z \leq 120$  mm. The coordinate origin is at the centre of the cube base. Dashed rectangles indicate the three reconstruction domains: Case 1 ( $-120 \leq x \leq -90$ ,  $150 \leq z \leq 180$  mm), Case 2 ( $-150 \leq x \leq -60$ ,  $0 \leq z \leq 60$  mm), and Case 3 ( $-60 \leq x \leq 60$ ,  $120 \leq z \leq 180$  mm).

#### 4.2.3.1. Case 1: Freestream Region

The first validation case focuses on a region upstream of and above the surface-mounted cube, where the flow is undisturbed by the obstacle. The reconstruction domain is  $-120 \leq x \leq -90$  mm,  $150 \leq z \leq 180$  mm. In this region, the  $x$ -velocity  $u$  approaches the freestream velocity  $U_\infty$  and the  $z$ -velocity  $w$  is small.

This case serves as a baseline to assess the method's ability to reconstruct simple flow fields accurately. The expected low vorticity and nearly divergence-free character of the freestream provide an opportunity to verify that the method does not introduce spurious oscillations or artifacts in regions where the flow is smooth. Additionally, this case tests whether the divergence-free constraint is properly satisfied in a region where the experimental measurements themselves should naturally exhibit low divergence.

#### 4.2.3.2. Case 2: Upstream of the Cube Base

The second validation case examines the flow upstream of the cube at the base level. The reconstruction domain is  $-150 \leq x \leq -60$  mm,  $0 \leq z \leq 60$  mm. This region is characterised by the development of the boundary layer along the flat plate and the onset of flow deceleration as the approaching flow encounters the upstream face of the cube. The  $x$ -velocity decelerates and the  $z$ -velocity shows an upward deflection as the flow is redirected over the obstacle. An adverse pressure gradient develops near the cube face.

This case tests the method's ability to handle moderate velocity gradients and the transition from undisturbed boundary layer flow to the stagnation region at the cube face. The pressure field reconstruction is relevant in this region, as the adverse pressure gradient approaching the cube drives the deceleration process.

#### 4.2.3.3. Case 3: Above the Surface-Mounted Cube

The third validation case focuses on the region directly above the surface-mounted cube. The reconstruction domain is  $-60 \leq x \leq 60$  mm,  $120 \leq z \leq 180$  mm. This region presents the most challenging conditions for reconstruction due to the presence of strong velocity gradients, flow separation at the upstream edge, flow reversal within the recirculation zone, and high vorticity magnitudes in the separated shear layer. The  $x$ -velocity exhibits both positive (downstream) and negative (upstream) values, and the  $z$ -velocity shows significant variations.

The reconstruction must handle steep gradients without introducing oscillations or excessive smoothing that would obscure the physical flow features. The separated flow region challenges the divergence-free constraint enforcement, as local errors in the experimental measurements may be larger in regions of high turbulence and strong shear. The pressure field reconstruction is particularly relevant in this region, as the adverse pressure gradient drives the separation process.

# Chapter 5

## Results and Discussion

This chapter presents the results of the verification and validation studies described in Chapter 4. Section 5.1. reports the verification results obtained with the manufactured solution, confirming that the implementation achieves the theoretical convergence rates from Section 2.6. and that the divergence-free constraint is satisfied to machine precision. Section 5.2. presents the validation results from applying the reconstruction method to the surface-mounted cube experiment.


## 5.1. Verification Results: Manufactured Solutions

The manufactured trigonometric velocity field (4.5)–(4.6) is used to verify the implementation on single-element and multi-element configurations, followed by systematic convergence studies under p-refinement and h-refinement. The five error quantities defined in Section 4.1.6. are evaluated throughout: the  $H(\text{div}; \Omega)$ -norm of the velocity, the  $L^2$ -norms of the streamfunction, vorticity, and total pressure, and the  $H(\text{curl}; \Omega)$ -norm of the vorticity.

5

### 5.1.1. Single-Element Reconstruction on the Reference Element


The reference element tests use the computational domain  $[-1, 1] \times [-1, 1]$  with a single spectral element at polynomial order  $p = 16$ . Three data densities  $n_r \in \{400, 1600, 3600\}$  are tested. The measurement points are distributed with uniform random probability throughout the domain. Figure 5.1 shows the spatial distribution of the measurement points for the three densities, coloured by the velocity magnitude  $|\mathbf{u}|$  evaluated from the analytical solution. The characteristic pattern of the trigonometric velocity field is visible even in the sparsest case.



**FIGURE 5.1** Spatial distribution of measurement points on the reference element  $[-1, 1] \times [-1, 1]$ , coloured by velocity magnitude  $|\mathbf{u}|$ . (a)  $n_r = 400$ , (b)  $n_r = 1600$ , (c)  $n_r = 3600$ . Single element,  $p = 16$ .


Figure 5.2 compares the analytical and reconstructed x-velocity component  $u$  for the three data densities. The rows correspond to  $n_r = 400$  (top),  $n_r = 1600$  (middle), and  $n_r = 3600$  (bottom). The reconstructed

field captures the sinusoidal structure of the manufactured solution in all three cases, with no discernible difference between the analytical and reconstructed contours. The y-velocity component  $v$  is shown in Figure 5.3 and exhibits the same level of agreement.




**FIGURE 5.2** Analytical (left) and reconstructed (right)  $x$ -velocity component  $u$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

Figure 5.4 shows the divergence of the reconstructed velocity field. The analytical divergence is identically zero (left column), while the recon-



**FIGURE 5.3** Analytical (left) and reconstructed (right) y-velocity component  $v$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

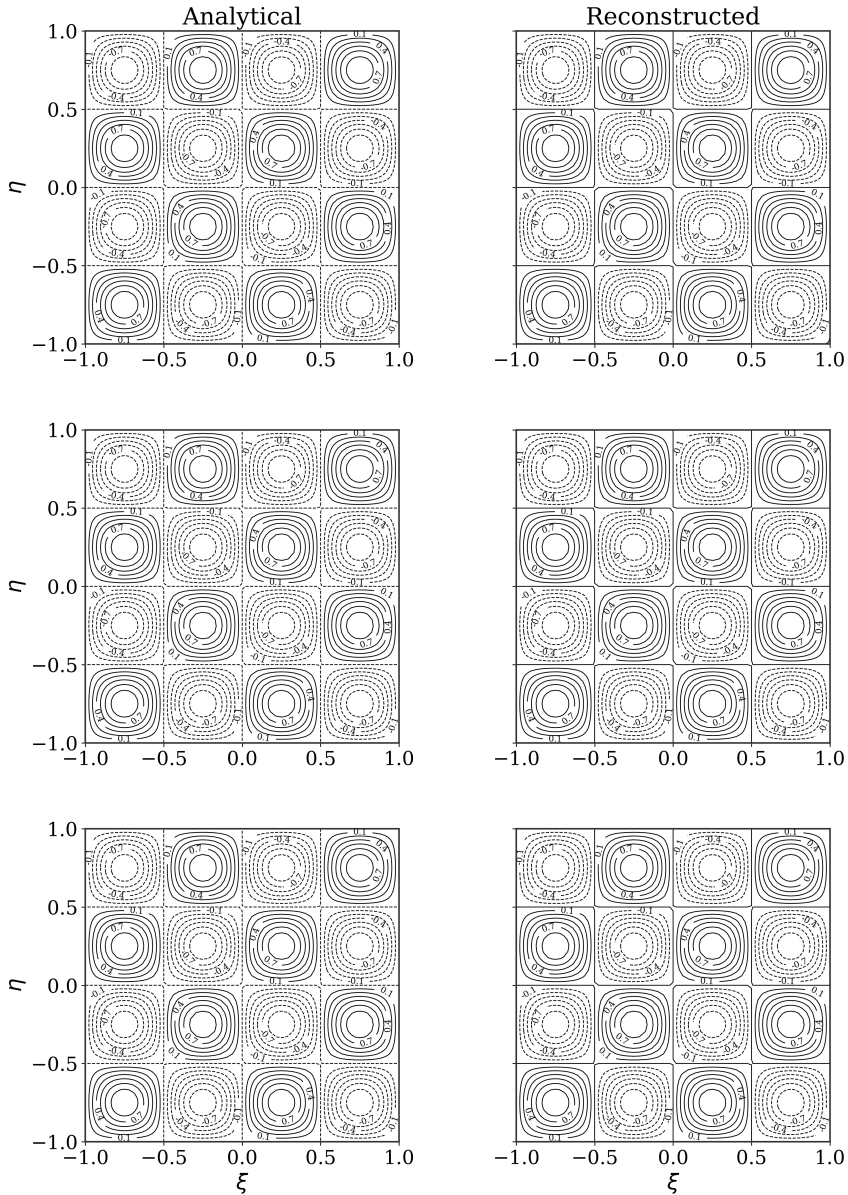
structed divergence (right column) is on the order of  $10^{-11}$  for all three data densities. This confirms that the saddle-point formulation enforces the discrete incompressibility constraint  $E^{2,1}\mathbf{u}_h = 0$  to near machine precision, regardless of the data density.



**FIGURE 5.4** Analytical (left) and reconstructed (right) divergence  $\nabla \cdot \mathbf{u}$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

The streamfunction, recovered from the velocity through the incidence matrix  $E^{1,0}$  (Section 3.4.), is shown in Figure 5.5. The contour lines of the


reconstructed field match the analytical solution closely at all data densities.



**FIGURE 5.5** Analytical (left) and reconstructed (right) streamfunction  $\psi$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

The vorticity, obtained from the weak curl formulation (Section 3.5.), is shown in Figure 5.6. The sinusoidal pattern is well captured at all data


densities, though the reconstruction at  $n_r = 400$  shows slightly more variation than the denser cases, reflecting the effect of the inverse estimate on limited data.



**FIGURE 5.6** Analytical (left) and reconstructed (right) vorticity  $\omega$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

The total pressure  $\tilde{P}$  and static pressure  $p$ , defined as  $p = \tilde{P} - \frac{1}{2}|\mathbf{u}|^2$ , are computed from the second saddle-point system (Section 3.6.) and shown


in Figures 5.7 and 5.8. Both fields are accurately reconstructed across all data densities.



**FIGURE 5.7** Analytical (left) and reconstructed (right) total pressure  $\bar{P}$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

### 5.1.2. Single-Element Reconstruction on the Physical Domain

The same single-element tests are repeated on the physical domain  $[0, 1] \times [0, 1]$  to verify that the Piola transformation (Section 2.5.5.) and the Jaco-



**FIGURE 5.8** Analytical (left) and reconstructed (right) static pressure  $p$  on  $[-1, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

bian scaling in the mass matrices (2.47)–(2.49) are correctly implemented. The polynomial order, data densities, and reconstruction resolution are identical to the reference element tests.

Figures 5.9 and 5.10 show the reconstructed velocity components on the physical domain. The fields exhibit the same level of agreement with the analytical solution as on the reference element, confirming that the affine mapping introduces no additional approximation error.

The divergence of the reconstructed velocity on the physical domain (Figure 5.11) is again on the order of  $10^{-11}$ , confirming that the Piola transformation preserves the incompressibility condition exactly at the discrete level.

The streamfunction, vorticity, total pressure, and static pressure on the physical domain are shown in Figures 5.12–5.15. All fields match their analytical counterparts with accuracy comparable to the reference element results.

5


These results confirm that the Piola transformation is correctly implemented: the coordinate mapping from  $[-1, 1]^2$  to  $[0, 1]^2$  introduces no degradation in reconstruction accuracy, the divergence-free constraint remains satisfied to machine precision, and all derived fields are reconstructed with the same fidelity as on the reference element.

### 5.1.3. Multi-Element Reconstruction


The multi-element verification tests the implementation when the domain  $[-1, 1] \times [-1, 1]$  is decomposed into a  $4 \times 4$  mesh ( $K_\xi = K_\eta = 4$ ) with polynomial order  $p = 4$ . A total of  $n_r = 1600$  randomly distributed measurement points are used. This configuration introduces element interfaces where flux continuity must be enforced through the Lagrange multiplier constraints (Section 3.3.4.) and each element requires the Piola transformation to map between reference and physical coordinates.

Figures 5.16 and 5.17 compare the analytical and reconstructed velocity components. The reconstructed fields reproduce the sinusoidal pattern of the manufactured solution, with no visible discontinuities at element interfaces. The lower polynomial order ( $p = 4$  compared to  $p = 16$  in the single-element tests) means fewer lobes of the trigonometric field are sharply resolved, but the overall structure is well captured.


The divergence of the reconstructed velocity (Figure 5.18) is on the order of  $10^{-13}$ , confirming that the discrete incompressibility constraint is satisfied to near machine precision in the multi-element configuration. The



**FIGURE 5.9** Analytical (left) and reconstructed (right) x-velocity component  $u$  on  $[0, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .

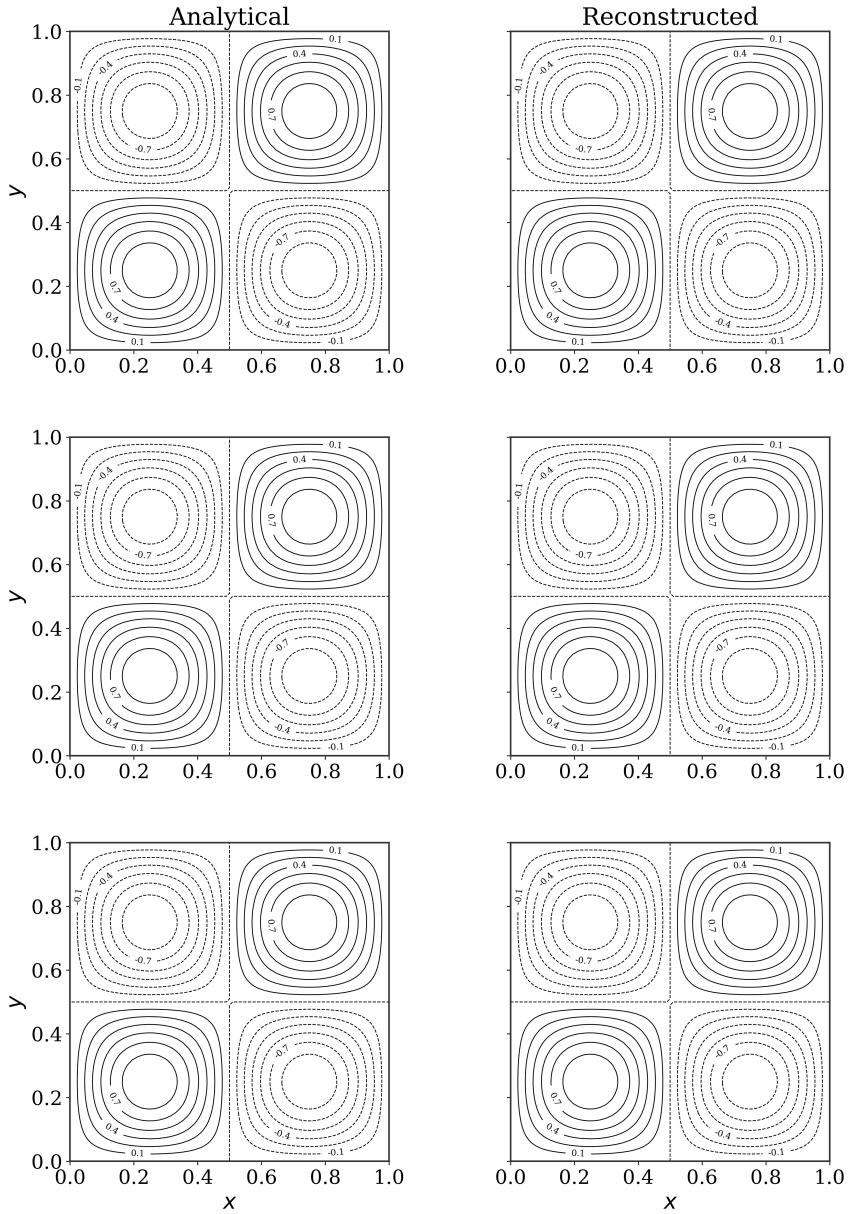


**FIGURE 5.10** Analytical (left) and reconstructed (right) y-velocity component  $v$  on  $[0, 1]^2$ . Top to bottom:  $n_T = 400, 1600, 3600$ . Single element,  $p = 16$ .




5

**FIGURE 5.11** Analytical (left) and reconstructed (right) divergence  $\nabla \cdot \mathbf{u}$  on  $[0, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .




5

**FIGURE 5.12** Analytical (left) and reconstructed (right) streamfunction  $\psi$  on  $[0, 1]^2$ . Top to bottom:  $n_\tau = 400, 1600, 3600$ . Single element,  $p = 16$ .




5

**FIGURE 5.13** Analytical (left) and reconstructed (right) vorticity  $\omega$  on  $[0, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .




5


**FIGURE 5.14** Analytical (left) and reconstructed (right) total pressure  $\tilde{P}$  on  $[0, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .



**FIGURE 5.15** Analytical (left) and reconstructed (right) static pressure  $p$  on  $[0, 1]^2$ . Top to bottom:  $n_r = 400, 1600, 3600$ . Single element,  $p = 16$ .



**FIGURE 5.16** Analytical (left) and reconstructed (right) x-velocity component  $u$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .




**FIGURE 5.17** Analytical (left) and reconstructed (right) y-velocity component  $v$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .


element boundaries are faintly visible in the divergence contour, reflecting the slightly different local truncation errors in each element, but the magnitude remains negligible.

Figure 5.19 shows the absolute flux mismatch at each interface degree of freedom for the velocity ( $|\mathbb{N}_L \mathbf{u}_h|$ ) and the pressure gradient ( $|\mathbb{N}_L \mathbf{q}_h|$ ). The velocity flux mismatch is on the order of  $10^{-17}$  to  $10^{-16}$ , at or below machine precision. The pressure gradient mismatch is slightly higher, reaching up to  $10^{-15}$ , but still well within machine precision. These values confirm that the Lagrange multiplier constraints correctly enforce flux continuity across all element interfaces.

The streamfunction (Figure 5.20) shows contour lines that are continuous across element boundaries, with the element grid visible as thin dashed




**FIGURE 5.18** Analytical (left) and reconstructed (right) divergence  $\nabla \cdot \mathbf{u}$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .



**FIGURE 5.19** Absolute flux mismatch at element interfaces for velocity and pressure gradient. Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .


lines. The vorticity, total pressure, and static pressure are shown in Figures 5.21–5.23. All derived fields are accurately reconstructed.

Figure 5.24 compares the error norms across three configurations: the single-element case ( $p = 16$ ,  $K = 1$ ), and two multi-element cases ( $p = 4$ ,  $K = 4$  and  $p = 8$ ,  $K = 4$ ), all at  $n_r = 1600$ . The multi-element case with  $p = 8$  achieves comparable or lower error levels than the single-element case for all quantities except the pressure, while the  $p = 4$  case shows higher errors as expected from the reduced polynomial order. The hierarchy of error levels across quantities is consistent in all three configurations: the streamfunction has the lowest error, followed by the velocity, then the vorticity and pressure.




**FIGURE 5.20** Analytical (left) and reconstructed (right) streamfunction  $\psi$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .


5



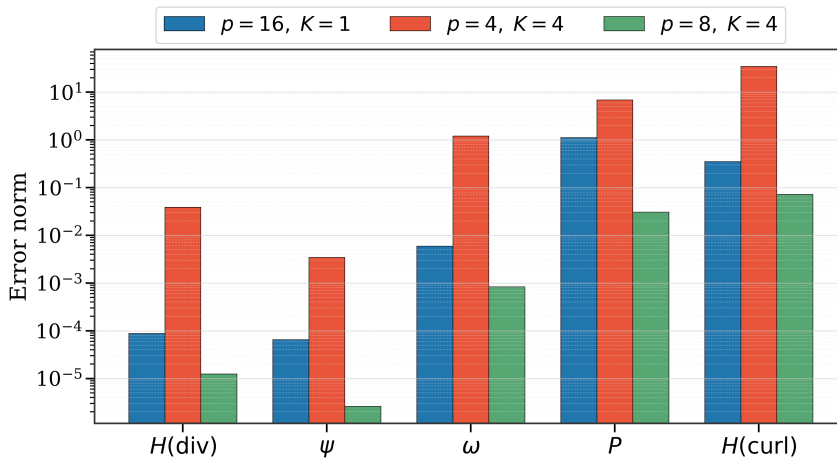
**FIGURE 5.21** Analytical (left) and reconstructed (right) vorticity  $\omega$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .



**FIGURE 5.22** Analytical (left) and reconstructed (right) total pressure  $\tilde{P}$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .



**FIGURE 5.23** Analytical (left) and reconstructed (right) static pressure  $p$ . Multi-element mesh  $K = 4$ ,  $p = 4$ ,  $n_r = 1600$ .



**FIGURE 5.24** Error norms for single-element ( $p = 16$ ,  $K = 1$ ) and multi-element ( $p = 4$  and  $p = 8$ ,  $K = 4$ ) configurations.  $n_r = 1600$ .

### 5.1.4. Convergence Studies

The convergence studies systematically investigate how the reconstruction error decreases with increasing polynomial order ( $p$ -refinement) and increasing number of elements ( $h$ -refinement). These studies verify that the implementation achieves the theoretical convergence rates derived in Section 2.6. and summarised in Table 2.3. Three data densities  $n_r \in \{40\,000, 160\,000, 360\,000\}$  are used throughout.

#### 5.1.4.1. $p$ -Refinement Results

The  $p$ -refinement study holds the mesh fixed and increases the polynomial order  $p$  from 3 to 14. Three mesh configurations are tested ( $K = 1$ ,

$K = 4$ , and  $K = 7$ ) on the reference domain  $[-1, 1] \times [-1, 1]$ . Because the manufactured solution is analytic, the polynomial approximation result (2.52) predicts that the error should decrease exponentially with  $p$ , appearing as a straight line on a semi-log plot. For derived quantities such as the vorticity and pressure, the inverse estimate (2.51) introduces algebraic prefactors of the form  $p^2/h$  that shift the convergence curves vertically but do not change the exponential rate (Section 2.6.2.). The theory therefore predicts that all five error quantities should appear as approximately parallel straight lines on the semi-log plot, ordered from bottom to top according to the number of differentiation steps separating each quantity from the streamfunction.

Figure 5.25 shows the  $H(\text{div}; \Omega)$ -error of the velocity. Since the divergence contribution to the  $H(\text{div}; \Omega)$ -norm vanishes to machine precision (the constraint  $E^{2,1}\mathbf{u}_h = 0$  is enforced exactly), the  $H(\text{div}; \Omega)$ -norm reduces to the  $L^2$ -velocity error and the convergence rate follows (2.52) directly.

5

For  $K = 1$ , the error decreases from approximately  $10^1$  at  $p = 3$  to approximately  $10^{-3}$  at  $p = 14$ . The convergence curve exhibits a visible even-odd staggering pattern, which is not predicted by the general theory but arises from the definite parity of the manufactured solution. The two velocity components have complementary symmetry,


$$\begin{aligned} u &= 2\pi \sin(2\pi x) \cos(2\pi y), \\ v &= -2\pi \cos(2\pi x) \sin(2\pi y), \end{aligned}$$

where  $u$  is odd in  $x$  and even in  $y$ , while  $v$  is even in  $x$  and odd in  $y$ . In the 1-form basis,  $u$  is represented by edge functions (degree  $p - 1$ ) in  $x$  tensored with Lagrange functions (degree  $p$ ) in  $y$ , and  $v$  uses the transposed combination. For even  $p$ , the highest-degree basis function has the wrong parity to contribute in one direction for each component, effectively reducing the useful polynomial degree by one for both  $u$  and  $v$ . Since the  $H(\text{div}; \Omega)$ -error adds both component errors in quadrature, the staggering remains visible in the combined norm.

For  $K = 4$ , the element boundaries fall at  $x = -1, -0.5, 0, 0.5, 1$ . These coincide with the zeros of  $\sin(2\pi x)$ , so each element retains definite parity and the staggering persists. The smaller element size does however produce substantially faster convergence, with errors reaching approximately  $10^{-13}$  at  $p = 14$ . For  $K = 7$ , the interior element boundaries fall at irrational multiples of the domain length ( $x = -5/7, -3/7, \dots$ ), which do not coincide with any zeros of the trigonometric functions. This breaks

the parity on every element, eliminating the staggering entirely. The error decreases smoothly, reaching approximately  $10^{-12}$  by  $p = 12$ .

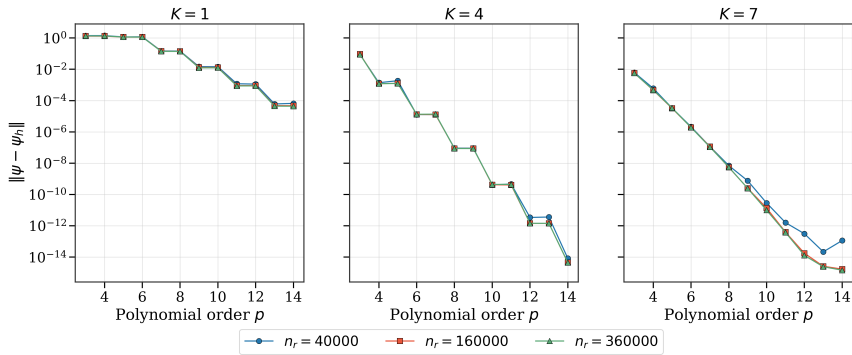
For all three mesh configurations, the three data densities produce nearly identical curves until the error approaches machine precision. For  $K = 7$  at high polynomial orders, the  $n_r = 40\,000$  curve separates slightly from the denser data sets. With 49 elements, each element receives approximately  $40\,000/49 \approx 816$  measurement points, which is still well above the number of velocity degrees of freedom per element but becomes less generous at high  $p$ .



**FIGURE 5.25**  $H(\text{div}; \Omega)$ -error versus polynomial order  $p$  for  $K = 1$ ,  $K = 4$ , and  $K = 7$ .

Figure 5.26 shows the  $L^2$ -error of the streamfunction. The streamfunction is recovered from the velocity by integration through the incidence matrix  $E^{1,0}$ . According to (2.60), this integration step gains one algebraic order relative to the velocity. Under  $p$ -refinement this translates to a vertical downward shift of the convergence curve relative to the velocity. The figure confirms this prediction. The streamfunction error is consistently the lowest of all five quantities across all mesh configurations. For  $K = 7$  the error reaches approximately  $10^{-14}$  by  $p = 12$ , approaching the limit of double-precision arithmetic. The data density sensitivity becomes visible at high  $p$  for  $K = 7$ , where the  $n_r = 40\,000$  curve levels off while the denser data sets continue to decrease. This is consistent with integration acting as a low-pass filter (Section 2.6.2.), so the streamfunction is less sensitive to limited data than the velocity itself.

Figure 5.27 shows the  $L^2$ -error of the vorticity. The vorticity is obtained from the velocity through the weak curl formulation (Section 3.5.). The error equation (2.55) shows that the vorticity error depends on how well the discrete curl operator approximates the continuous one. The inverse




**FIGURE 5.26** Streamfunction  $L^2$ -error versus polynomial order  $p$  for  $K = 1$ ,  $K = 4$ , and  $K = 7$ .

estimate (2.51) enters because the curl of a discrete test function  $\varepsilon^h$  satisfies

5

$$\|\nabla \times \varepsilon^h\| \leq C h^{-1} \|\varepsilon^h\|,$$


which introduces a factor  $h^{-1}$  in the upper bound (2.58). Under  $p$ -refinement at fixed  $h$  this becomes the prefactor  $p^2/h$ , shifting the vorticity convergence curve upward relative to the velocity. The figure confirms this. The vorticity error is consistently one to two orders of magnitude above the velocity error across all configurations. For  $K = 1$  the vorticity decreases from approximately  $10^2$  to  $10^{-2}$ , while for  $K = 7$  it reaches approximately  $10^{-11}$  by  $p = 13$ . The  $n_r = 40\,000$  curve separates from the denser data sets more strongly than for the velocity, which the theory also predicts since the factor  $h^{-1}$  amplifies the errors arising from limited data coverage in the velocity reconstruction.



**FIGURE 5.27** Vorticity  $L^2$ -error versus polynomial order  $p$  for  $K = 1$ ,  $K = 4$ , and  $K = 7$ .

Figure 5.28 shows the  $L^2$ -error of the total pressure. The total pressure is computed from the saddle-point system whose right-hand side involves the vorticity force term  $\mathbf{C}_h = \mathbf{u}_h \times \omega_h$ . According to (2.62), the error in this forcing is dominated by  $\|\mathbf{u}\|_{L^\infty} \|\omega - \omega_h\|_{L^2}$ , and the saddle-point stability theory bounds the pressure error between  $\mathcal{O}(h^{p-1})$  and  $\mathcal{O}(h^p)$  (2.63). The pressure is therefore sensitive to how well both the velocity and the vorticity are resolved, because the forcing term involves their product.


For  $K = 1$ , the pressure error remains between  $10^0$  and  $10^2$  across all polynomial orders, showing essentially no net convergence. The parity mismatch simultaneously degrades both the velocity and the vorticity at even  $p$ , and their product in the forcing term compounds the error. Since the single-element domain  $[-1, 1]^2$  preserves the definite parity at all polynomial orders, the forcing error (2.62) does not decrease with  $p$ . For  $K = 4$  the pressure converges but the staggering is still visible. For  $K = 7$  the parity alignment is completely broken and the pressure converges smoothly and exponentially, reaching approximately  $10^{-10}$  at  $p = 14$ . This confirms that the lack of convergence at  $K = 1$  is a property of the test case symmetry and not a limitation of the method.



**FIGURE 5.28** Total pressure  $L^2$ -error versus polynomial order  $p$  for  $K = 1$ ,  $K = 4$ , and  $K = 7$ .

Figure 5.29 shows the  $H(\text{curl}; \Omega)$ -error of the vorticity. The  $H(\text{curl}; \Omega)$ -norm combines the  $L^2$ -vorticity error with the  $L^2$ -error of the curl of vorticity (4.17). The curl of the vorticity is effectively a second derivative of the velocity, requiring two applications of the inverse estimate (2.51). Under  $p$ -refinement this introduces a prefactor  $(p^2/h)^2$ , which is the largest algebraic prefactor of all five quantities. The theory (2.59) therefore predicts the  $H(\text{curl}; \Omega)$ -error to be the highest and the most sensitive to data density.

The figure confirms both predictions. For  $K = 1$  the error decreases only from approximately  $10^3$  to  $10^1$ , the slowest convergence of all five quantities. For  $K = 7$  the error reaches approximately  $10^{-8}$  at  $p = 13$ , which is several orders of magnitude above the velocity and streamfunction at the same  $p$ . The  $n_r = 40\,000$  curve separates from the denser data sets more strongly than for any other quantity. Each application of the inverse estimate amplifies data sparsity effects by a factor  $p^2/h$ , so two applications produce the strongest sensitivity to limited data.




**FIGURE 5.29**  $H(\text{curl}; \Omega)$ -error versus polynomial order  $p$  for  $K = 1$ ,  $K = 4$ , and  $K = 7$ .

Figure 5.30 combines all five error quantities for  $K = 7$  at the highest data density ( $n_r = 360\,000$ ). The theory predicts that all quantities should converge exponentially and that the curves should be approximately parallel, separated by vertical offsets that reflect the number of differentiation or integration steps separating each quantity from the velocity. The figure confirms this. The streamfunction converges fastest, reaching approximately  $10^{-14}$  around  $p = 12$ . The  $H(\text{div}; \Omega)$ -velocity error follows at approximately  $10^{-12}$  around  $p = 12$ . The vorticity and total pressure are close together at approximately  $10^{-10}$  to  $10^{-11}$  around  $p = 13$ . The  $H(\text{curl}; \Omega)$ -error converges slowest at approximately  $10^{-8}$  around  $p = 13$ . The vertical separation between successive quantities is approximately two orders of magnitude, consistent with the inverse estimate (2.51) at  $p \approx 7$  and  $h = 2/7 \approx 0.29$ .

#### 5.1.4.2. $h$ -Refinement Results

The  $h$ -refinement study holds the polynomial order  $p$  fixed and increases the number of elements  $K_\xi = K_\eta = K$ , decreasing the element size  $h = 2/K$ . Polynomial orders  $p \in \{3, 4, 5, 6, 7, 8\}$  and element counts  $K \in \{4, 8, 12, 16, 20, 24, 28\}$  are tested on the reference domain  $[-1, 1] \times [-1, 1]$ . Under  $h$ -refinement,



**FIGURE 5.30** All five error quantities versus polynomial order  $p$  at  $K = 7$  and  $n_r = 360\,000$ .


the polynomial approximation result (2.50) predicts algebraic convergence: the error decreases as  $\mathcal{O}(h^r)$  for some rate  $r$ , appearing as a straight line on a log-log plot with slope  $r$ .

The theoretical rates for each reconstructed quantity follow from the error analysis in Section 2.6.3. The velocity, as the primary unknown, converges at  $\mathcal{O}(h^p)$  in the  $H(\text{div}; \Omega)$ -norm (2.54). The divergence contribution vanishes to machine precision because  $E^{2,1}\mathbf{u}_h = 0$  is enforced exactly. The streamfunction, recovered by integration through  $E^{1,0}$ , gains one order and converges at  $\mathcal{O}(h^{p+1})$  (2.60). The vorticity is bounded between  $\mathcal{O}(h^p)$  and  $\mathcal{O}(h^{p-1})$  (2.58), where the upper bound arises from the inverse estimate applied to the error equation (2.55). The total pressure is bounded between  $\mathcal{O}(h^p)$  and  $\mathcal{O}(h^{p-1})$  (2.63), with the upper bound set by the vorticity force error (2.62). The  $H(\text{curl}; \Omega)$ -norm of the vorticity incurs a second inverse estimate penalty, giving  $\mathcal{O}(h^{p-2})$  (2.59).

The individual convergence plots at  $p = 7$  are presented in Figures 5.31–5.35.


Figure 5.31 shows the  $H(\text{div}; \Omega)$ -error versus element size for  $p = 7$ . The theory (2.54) predicts a slope of  $p = 7$  on the log-log plot. The mea-

sured slope for  $n_r = 360\,000$  is approximately 7.08, confirming the prediction. For the coarsest data ( $n_r = 40\,000$ ), the errors deviate upward from the asymptotic line at small  $h$ . At  $K = 28$ , each element contains only  $40\,000/784 \approx 51$  measurement points, which is close to the number of velocity degrees of freedom per element ( $2p(p+1) = 112$  for  $p = 7$ ). The least-squares problem becomes poorly overdetermined and data sparsity begins to dominate the polynomial approximation error. The two denser data sets produce nearly identical results, confirming that above a threshold data density the convergence rate is controlled entirely by the polynomial approximation.



**FIGURE 5.31**  $H(\text{div}; \Omega)$ -error versus element size  $h$  for  $p = 7$  and three data densities.

The streamfunction  $L^2$ -error (Figure 5.32) should converge one order faster than the velocity according to (2.60), predicting a slope of  $p + 1 = 8$  for  $p = 7$ . The measured slope is approximately 8.02, in good agreement with the prediction. The streamfunction achieves the lowest absolute error levels of all quantities, reaching approximately  $10^{-8}$  at the finest mesh. The sensitivity to data density is weaker than for the velocity: the three  $n_r$  curves remain close together even at the smallest element sizes. This is predicted by the theory, since the integration through  $E^{1,0}$  acts as a discrete antidifferentiation that smooths out local sampling errors in the velocity.



**FIGURE 5.32** Streamfunction  $L^2$ -error versus element size  $h$  for  $p = 7$  and three data densities.

The vorticity  $L^2$ -error (Figure 5.33) is bounded between  $\mathcal{O}(h^p)$  and  $\mathcal{O}(h^{p-1})$  (2.58). The upper bound arises from applying the inverse estimate (2.51) to the error equation (2.55): the term  $C h^{-1} \|\mathbf{u} - \mathbf{u}_h\|$  produces  $h^{-1} \cdot h^p = h^{p-1}$ , predicting a slope between 6 and 7 for  $p = 7$ . The measured slope is approximately 6.37, falling within these bounds. The vorticity error also exhibits greater sensitivity to data density than the velocity: the three  $n_r$  curves are clearly separated at small  $h$ . The factor  $h^{-1}$  in the upper bound of (2.58) amplifies the errors arising from limited data coverage in the velocity reconstruction. As  $h$  decreases, this amplification factor grows, so the gap between data-limited and data-sufficient runs widens.

The total pressure  $L^2$ -error (Figure 5.34) is bounded between  $\mathcal{O}(h^p)$  and  $\mathcal{O}(h^{p-1})$  (2.63). The upper bound is set by the vorticity force error (2.62):  $\|\mathbf{C} - \mathbf{C}_h\|_{L^2} = \mathcal{O}(h^{p-1})$ , which enters the saddle-point system as a right-hand side perturbation. For  $p = 7$ , the theory predicts a slope between 6 and 7. The measured slope is approximately 6.27, falling within these bounds. At this polynomial order the total pressure convergence rate closely tracks the vorticity, consistent with the prediction from (2.63) that the pressure error is dominated by the vorticity force error (2.62).

5




FIGURE 5.33 Vorticity  $L^2$ -error versus element size  $h$  for  $p = 7$  and three data densities.




FIGURE 5.34 Total pressure  $L^2$ -error versus element size  $h$  for  $p = 7$  and three data densities.

The  $H(\text{curl}; \Omega)$ -error of the vorticity (Figure 5.35) requires two applications of the inverse estimate (2.51), giving the theoretical rate  $\mathcal{O}(h^{p-2})$  (2.59) and a predicted slope of  $p-2 = 5$  for  $p = 7$ . The measured slope of approximately 5.24 confirms this prediction. This quantity exhibits the strongest sensitivity to data density of all five error measures: the three  $n_r$  curves remain widely separated across most of the element size range. The two successive inverse estimates each introduce a factor  $h^{-1}$ , so the total amplification of data sparsity effects scales as  $h^{-2}$ . At small  $h$  this amplification is large, and the  $n_r = 40\,000$  curve departs from the asymptotic regime much earlier than for the velocity or streamfunction.





FIGURE 5.35  $H(\text{curl}; \Omega)$ -error versus element size  $h$  for  $p = 7$  and three data densities.

Figure 5.36 combines all five error quantities on a single log-log plot for  $p = 7$  at the highest data density ( $n_r = 360\,000$ ). The theory predicts the hierarchy of slopes  $\psi$  at  $p+1 = 8$ ,  $H(\text{div}; \Omega)$  at  $p = 7$ ,  $\omega(L^2)$  and  $\bar{P}$  between  $p - 1 = 6$  and  $p = 7$ , and  $H(\text{curl}; \Omega)$  at  $p - 2 = 5$ . The measured slopes are 8.02 for the streamfunction, 7.08 for the  $H(\text{div}; \Omega)$ -velocity, 6.37 for the vorticity, 6.27 for the total pressure, and 5.24 for the  $H(\text{curl}; \Omega)$ -norm, all in good agreement with the predictions. The gap of approximately one algebraic order between successive quantities in the hierarchy is clearly visible and consistent with the  $p$ -refinement results in Figure 5.30, where the same ordering was observed as vertical offsets on the semi-log plot.



**FIGURE 5.36** All five error quantities versus element size  $h$  at  $p = 7$  and  $n_r = 360\,000$ .

5


Table 5.1 compares the theoretical and measured convergence rates at  $p = 7$  for the densest data set ( $n_r = 360\,000$ ).

**TABLE 5.1** Theoretical and measured  $h$ -refinement convergence rates at  $p = 7$  ( $n_r = 360\,000$ ).

Quantity	Norm	Theoretical rate	Measured ( $p = 7$ )
Streamfunction $\psi$	$L^2$	$\mathcal{O}(h^{p+1})$	$\approx 8.02$
Velocity $\mathbf{u}$	$H(\text{div}; \Omega)$	$\mathcal{O}(h^p)$	$\approx 7.08$
Total pressure $\tilde{P}$	$L^2$	$\mathcal{O}(h^{p-1})$ to $\mathcal{O}(h^p)$	$\approx 6.27$
Vorticity $\omega$	$L^2$	$\mathcal{O}(h^{p-1})$ to $\mathcal{O}(h^p)$	$\approx 6.37$
Vorticity $\omega$	$H(\text{curl}; \Omega)$	$\mathcal{O}(h^{p-2})$	$\approx 5.24$

Figure 5.37 presents the measured convergence slopes for all five error quantities as a function of polynomial order, computed from the densest data set. The theory predicts that the slopes should increase linearly with  $p$ , with constant offsets between quantities:  $\psi$  at  $p + 1$ ,  $H(\text{div}; \Omega)$  at  $p$ ,  $\omega$  ( $L^2$ ) between  $p - 1$  and  $p$ ,  $\tilde{P}$  between  $p - 1$  and  $p$ , and  $H(\text{curl}; \Omega)$  at  $p - 2$ . The figure confirms this hierarchy. The streamfunction consistently achieves the highest slope, increasing from approximately 4.3 at  $p = 3$  to 7.6 at  $p = 8$ . The  $H(\text{div}; \Omega)$ -error and vorticity track each other closely, with the vorticity slope approximately one unit below the velocity. The  $H(\text{curl}; \Omega)$ -error has the lowest slope throughout, approximately two units below the velocity.

The total pressure exhibits a notable departure at low polynomial orders: the measured slope at  $p = 3$  is approximately 1.4, well below the theoretical lower bound of  $p - 1 = 2$ . This is related to the even-odd staggering discussed in the  $p$ -refinement study: at low  $p$ , the coupling between velocity and vorticity errors in the forcing term  $C_h$  (2.62) is more sensitive to parity effects. At higher  $p$ , the pressure slope recovers and reaches approximately 7.8 at  $p = 8$ , exceeding  $p$  and approaching  $p + 1$ . A systematic even-odd staggering is visible across all quantities: even polynomial orders ( $p = 4, 6, 8$ ) produce slightly lower slopes than their odd neighbours ( $p = 3, 5, 7$ ). This is the same parity effect observed in the  $p$ -refinement study, now manifesting as reduced algebraic convergence rates rather than zigzag patterns on a semi-log plot.



**FIGURE 5.37** Measured  $h$ -refinement convergence slopes versus polynomial order  $p$  for all reconstructed quantities at  $n_r = 360\,000$ .

A clear pattern emerges in the sensitivity to data density across the convergence studies. The streamfunction, being an integral of the velocity through  $E^{1,0}$ , is the most robust to limited data: the integration step acts as a low-pass filter that smooths out local sampling errors. The velocity itself shows moderate sensitivity. The vorticity, obtained through the weak curl involving the inverse estimate (2.51), amplifies data sparsity effects by

a factor  $h^{-1}$ . The  $H(\text{curl}; \Omega)$ -error, involving two successive applications of the inverse estimate, is the most sensitive. This ordering matches the hierarchy predicted by the error analysis and is consistently observed across both  $p$ -refinement and  $h$ -refinement studies.

### 5.1.5. Summary of Verification Findings

The verification tests confirm that the mimetic spectral element reconstruction method is correctly implemented. The divergence-free constraint is satisfied to machine precision in all configurations, confirming that the algebraic constraint  $E^{2,1}\mathbf{u}_h = 0$  is exactly enforced. The Piola transformation introduces no additional approximation error. The multi-element configuration correctly handles interface coupling through the Lagrange multiplier constraints, with flux mismatch at machine precision.

Under  $p$ -refinement, exponential convergence is observed for all reconstructed quantities, with the hierarchy of convergence rates predicted by the error analysis preserved. Under  $h$ -refinement, algebraic convergence is observed at rates consistent with the theoretical predictions from Table 2.3. The streamfunction converges at  $\mathcal{O}(h^{p+1})$  (2.60), the velocity at  $\mathcal{O}(h^p)$  (2.54), the vorticity and total pressure within the bounds of  $\mathcal{O}(h^{p-1})$  to  $\mathcal{O}(h^p)$  (2.58) (2.63), and the  $H(\text{curl}; \Omega)$ -error at  $\mathcal{O}(h^{p-2})$  (2.59).

5

## 5.2. Validation Results: Surface-Mounted Cube

The reconstruction method is now applied to the experimental surface-mounted cube data described in Section 4.2.. For each of the three validation cases defined in Section 4.2.3., the raw particle tracks within a spanwise slab of half-thickness  $\delta$  are projected onto the  $x$ - $z$  plane and used as input to the reconstruction. Two configurations are tested per case: a single-element reconstruction with  $p = 8$  and  $K = 1$ , and a multi-element reconstruction with  $p = 4$  and  $K = 4$ . The reconstructed velocity magnitude, vorticity, divergence, streamfunction, total pressure, and static pressure are compared against the ensemble-averaged reference field from [11] where applicable. The static pressure is obtained by subtracting the dynamic head from the reconstructed total pressure,  $p_{\text{static}} = \bar{P} - \frac{1}{2}|\mathbf{u}_h|^2$ , and inherits the same accuracy hierarchy as the total pressure since both share the velocity solve as their upstream input. The Hendriksen reference does not include a pressure field, so the static pressure is reported as a reconstruction-only diagnostic. The Hendriksen vorticity is plotted as  $-\omega_y$  to match the sign convention of the two-dimensional curl  $\omega = \partial w/\partial x - \partial u/\partial z$  used in the reconstruction. The Hendriksen divergence

is computed as  $\partial u/\partial x + \partial w/\partial z$  from the individual gradient components available in the bin-averaged data.

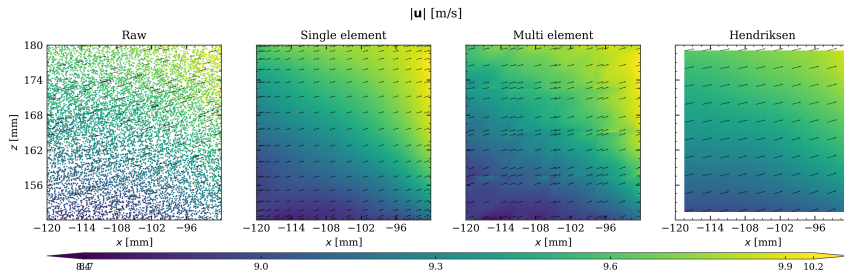
Unlike the verification study, no analytical solution exists for the experimental data. The comparison is therefore qualitative: the velocity and vorticity fields are compared against Hendriksen's ensemble average, while the streamfunction, total pressure, and static pressure, which are not available from the reference data, are assessed for physical consistency. The divergence comparison is particularly informative because the reconstruction enforces  $\nabla \cdot \mathbf{u} = 0$  exactly through the constraint  $E^{2,1} \mathbf{u}_h = 0$ , whereas the Hendriksen reference field has no such constraint and retains nonzero divergence from the bin averaging process and noise in the data.

### 5.2.1. Case 1: Freestream Region

Case 1 covers a  $30 \times 30$  mm domain in the freestream upstream of the cube ( $-120 \leq x \leq -90$  mm,  $150 \leq z \leq 180$  mm), well above and upstream of the cube. The slab half-thickness is  $\delta = 3$  mm. No boundary conditions are applied ( $\mathbf{u} \cdot \mathbf{n} = 0$  is not enforced on any edge). This region serves as a baseline where the flow is expected to be nearly uniform with velocity close to  $U_\infty = 10 \text{ m s}^{-1}$ , low vorticity, and nearly constant total pressure.


Figure 5.38 shows the velocity magnitude from the raw particle data, the single-element and multi-element reconstructions, and the Hendriksen reference. The raw scatter data shows velocities in the range of 7 to 11 m/s, reflecting the instantaneous turbulent fluctuations in the individual particle tracks. Both reconstructions produce a smooth field with velocity magnitudes between approximately 8.8 and 9.8 m/s, consistent with the Hendriksen reference which shows values between 9.0 and 9.7 m/s. The slight spatial gradient visible in the Hendriksen field (higher velocity toward the upper-left corner) is also captured by both reconstructions. The velocity vectors are predominantly horizontal and directed to the right, as expected for freestream flow. The single-element and multi-element results are qualitatively similar, indicating that the polynomial approximation is adequate for this smooth flow region.

Figure 5.39 compares the vorticity from the two reconstructions and the Hendriksen reference. The Hendriksen field shows vorticity values of approximately  $\pm 15 \text{ s}^{-1}$ , which is close to zero and consistent with freestream flow where no strong shear layers are present. The single-element reconstruction produces vorticity values of approximately  $\pm 100 \text{ s}^{-1}$ , while the multi-element reconstruction shows values up to  $\pm 250 \text{ s}^{-1}$  with a visible grid pattern at element interfaces. The higher vorticity levels in the recon-



**FIGURE 5.38** Velocity magnitude  $|\mathbf{u}|$  for Case 1. Raw particle data, single-element ( $p = 8$ ,  $K = 1$ ), multi-element ( $p = 4$ ,  $K = 4$ ), and Hendriksen reference.  $\delta = 3$  mm.


struction compared to the Hendriksen reference are expected. The Hendriksen ensemble average uses 10 mm bins with second-order polynomial regression, which acts as a spatial low-pass filter that suppresses velocity fluctuations at scales below the bin size. The reconstruction, operating at higher effective resolution, preserves these small-scale variations. When differentiated to obtain the vorticity, these variations are amplified. The grid pattern in the multi-element result reflects the fact that the vorticity is obtained through the weak curl formulation, and its tangential continuity across element interfaces (guaranteed by the  $H(\text{curl}; \Omega)$ -regularity) does not prevent jumps in the normal derivative.



**FIGURE 5.39** Vorticity  $\omega$  for Case 1. Single-element, multi-element, and Hendriksen reference.  $\delta = 3$  mm.

Figure 5.40 shows the divergence of the velocity field on a base-10 logarithmic scale. The single-element reconstruction averages  $\log_{10} |\nabla \cdot \mathbf{u}| \approx -12$  and the multi-element reconstruction averages  $-11$ , both at machine precision and confirming that the algebraic constraint  $E^{2,1} \mathbf{u}_h = 0$  is exactly


enforced regardless of the quality or density of the input data. The Hendriksen reference, by contrast, sits at  $\log_{10} |\nabla \cdot \mathbf{u}| \approx +1$ . The mimetic reconstructions are therefore roughly 12–13 orders of magnitude tighter on the divergence-free condition than the bin averaging. This residual divergence is inherent to the bin-averaging process and noise in the data, neither of which enforces the incompressibility constraint. In the freestream the residual is modest for the Hendriksen plot, but it can become much larger in regions with strong velocity gradients.



**FIGURE 5.40** Divergence  $\log_{10} |\nabla \cdot \mathbf{u}|$  for Case 1. Single-element, multi-element, and Hendriksen reference.  $\delta = 3$  mm.


Figure 5.41 shows the reconstructed streamfunction. Both the single-element and multi-element configurations produce nearly parallel, evenly spaced contour lines, consistent with a nearly uniform freestream. The slight upward tilt of the streamlines (increasing  $\psi$  from lower-right to upper-left) reflects the small  $z$ -velocity component present in the freestream at this location, which is upstream of the cube and therefore subject to a weak upwash. The streamfunction exists as a well-defined scalar field precisely because the velocity is divergence-free. For the Hendriksen field, where the divergence is nonzero, a consistent streamfunction cannot be defined.

Figure 5.42 shows the reconstructed total pressure. For uniform freestream flow, the total pressure  $\tilde{P} = p + \frac{1}{2}|\mathbf{u}|^2$  should be approximately constant. The single-element reconstruction produces a nearly uniform field with variations of approximately  $\pm 0.25 \text{ m}^2 \text{ s}^{-2}$ , which is small relative to the dynamic pressure  $\frac{1}{2}U_\infty^2 \approx 50 \text{ m}^2 \text{ s}^{-2}$ . The multi-element result shows somewhat larger variations with a visible grid pattern. The pressure is the most sensitive quantity in the reconstruction chain because it depends on the product  $\mathbf{u} \times \boldsymbol{\omega}$ , and any inaccuracies in the vorticity are amplified through this nonlinear coupling.




5

**FIGURE 5.41** Streamfunction  $\psi$  for Case 1. Single-element and multi-element reconstruction.  $\delta = 3$  mm.



**FIGURE 5.42** Total pressure  $\tilde{P}$  for Case 1. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

Figure 5.43 shows the reconstructed static pressure. The values lie in a narrow band between approximately  $-48$  and  $-38 \text{ m}^2 \text{ s}^{-2}$  with a weak gradient from lower-left (higher) to upper-right (lower). This near-uniform behaviour is consistent with the freestream interpretation of Case 1, where the absence of strong velocity gradients implies a small dynamic-head contribution and therefore little spatial variation in the static pressure. The single-element reconstruction is the smoothest; the multi-element field shows faint element seams but follows the same gradient.




**FIGURE 5.43** Static pressure  $\tilde{p}_{\text{static}}$  for Case 1. Single-element and multi-element reconstruction.  $\delta = 3 \text{ mm}$ .

### 5.2.2. Case 2: Upstream Cube Base

Case 2 covers a  $90 \times 60 \text{ mm}$  domain at the upstream base of the cube ( $-150 \leq x \leq -60 \text{ mm}$ ,  $0 \leq z \leq 60 \text{ mm}$ ). The slab half-thickness is  $\delta = 3 \text{ mm}$ . The right boundary ( $x = -60 \text{ mm}$ , the upstream face of the cube) and the bottom boundary ( $z = 0 \text{ mm}$ , the flat plate surface) are walls, where the normal flux condition  $\mathbf{u} \cdot \mathbf{n} = 0$  is enforced through the Lagrange multiplier matrix  $N_L$  (Section 3.3.4.).

Figure 5.44 shows the velocity magnitude. The freestream flow enters from the left at approximately  $8$  to  $10 \text{ m/s}$  and decelerates as it approaches the cube face at  $x = -60 \text{ mm}$ . A low-velocity region is visible near the bottom-right corner where the flat plate meets the cube, consistent with the adverse pressure gradient and boundary layer thickening expected

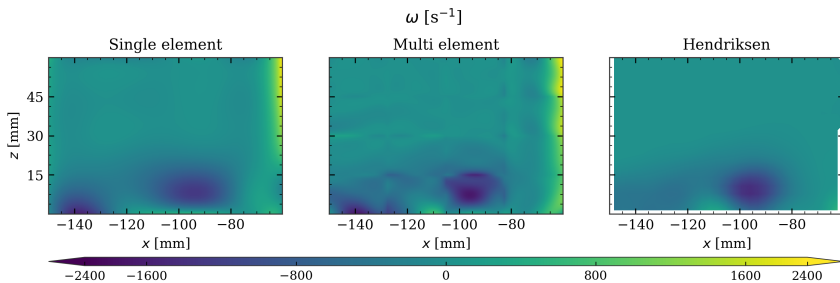
upstream of a bluff body. The raw particle data shows velocities ranging from near zero to above 10 m/s, with strong scatter near the wall where turbulent fluctuations are large relative to the mean velocity. Both reconstructions capture the deceleration pattern, and the spatial structure agrees well with the Hendriksen reference. The velocity magnitude near the cube face approaches zero at the wall, consistent with the enforced no-penetration condition.



**FIGURE 5.44** Velocity magnitude  $|\mathbf{u}|$  for Case 2. Raw particle data, single-element ( $p = 8$ ,  $K = 1$ ), multi-element ( $p = 4$ ,  $K = 4$ ), and Hendriksen reference.  $\delta = 3$  mm.

5

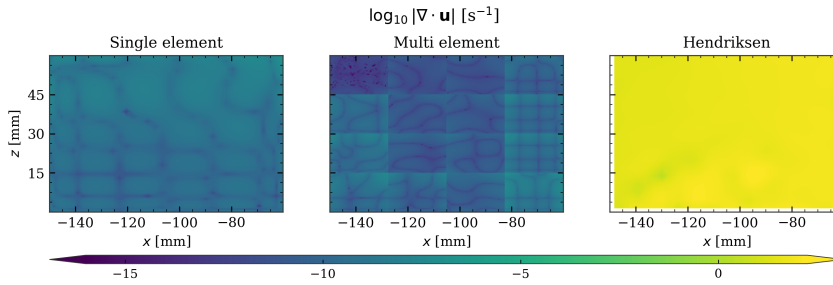
Figure 5.45 shows the vorticity. A concentrated region of negative vorticity is visible near the base of the cube face (around  $x \approx -80$  mm,  $z \approx 10$  mm), corresponding to the boundary layer separating from the flat plate as it encounters the adverse pressure gradient imposed by the cube. The single-element and multi-element reconstructions both capture this feature with peak vorticity values of approximately  $\pm 1500$   $\text{s}^{-1}$ , while the Hendriksen reference shows peak values of approximately  $\pm 800$   $\text{s}^{-1}$ . The spatial pattern is consistent across all three fields. The higher peak values in the reconstruction are attributed to the higher effective spatial resolution compared to the 10 mm bin averaging used by Hendriksen.



**FIGURE 5.45** Vorticity  $\omega$  for Case 2. Single-element, multi-element, and Hendriksen reference.  $\delta = 3$  mm.

Figure 5.46 shows the divergence on a base-10 logarithmic scale. The single-element reconstruction averages  $\log_{10} |\nabla \cdot \mathbf{u}| \approx -9$  and the multi-

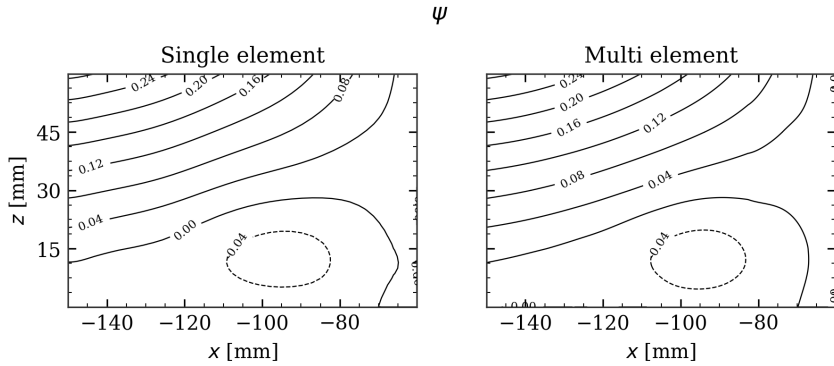
element averages  $-10$ , both effectively zero. The Hendriksen reference averages  $\log_{10} |\nabla \cdot \mathbf{u}| \approx +2$ , the largest of all three cases. The mimetic reconstructions are roughly 11–12 orders of magnitude tighter on the divergence-free condition than the bin averaging. The largest reference values appear near the cube face, where the velocity gradients are strongest, demonstrating that the bin-averaging procedure and noise in the data do not preserve mass conservation and that the violation is largest precisely in the regions of greatest physical interest. The mimetic reconstruction eliminates this problem entirely.



**FIGURE 5.46** Divergence  $\log_{10} |\nabla \cdot \mathbf{u}|$  for Case 2. Single-element, multi-element, and Hendriksen reference.  $\delta = 3$  mm.


Figure 5.47 shows the reconstructed streamfunction. The contour lines reveal the flow topology upstream of the cube. In the upper part of the domain the streamlines are approximately parallel, indicating freestream flow. Near the bottom-right corner a recirculation zone is visible, identified by the closed dashed contour line at  $\psi \approx -0.04$ . This recirculation corresponds to the separation bubble that forms upstream of the cube where the adverse pressure gradient causes the boundary layer to separate from the flat plate. Both the single-element and multi-element reconstructions capture this feature, with the recirculation centre located at approximately  $(x, z) \approx (-100, 15)$  mm. The existence of this well-defined streamfunction is a direct consequence of the velocity field that is divergence-free up to machine precision.

Figure 5.48 shows the reconstructed total pressure. The total pressure decreases from left to right as the flow approaches the cube, with the lowest values near the cube face and the flat plate junction. The spatial structure is physically plausible, with higher total pressure in the freestream and lower values in the decelerating and separated flow regions. The single-element result shows values ranging from approximately  $-100$  to  $0 \text{ m}^2 \text{ s}^{-2}$ , while the multi-element result shows a similar pattern with a larger range.




**FIGURE 5.47** Streamfunction  $\psi$  for Case 2. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

5



**FIGURE 5.48** Total pressure  $\tilde{P}$  for Case 2. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

Figure 5.49 shows the reconstructed static pressure. The field ranges from approximately  $-150$  to  $+70 \text{ m}^2 \text{ s}^{-2}$ . A pronounced low-pressure region is located near the upper-right of the domain, around the upstream upper edge of the cube ( $x \approx -60$  mm,  $z \approx 60$  mm), where the flow accelerates around the leading corner. Higher static pressure is observed upstream and along the lower-left part of the domain, consistent with the stagnation rise as the flow approaches the cube and decelerates toward the wall. Both configurations produce the same qualitative pattern; the multi-element field shows slight grid artefacts but agrees with the single-element distribution in magnitude and structure.




**FIGURE 5.49** Static pressure  $\tilde{p}_{\text{static}}$  for Case 2. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

### 5.2.3. Case 3: Above the Surface-Mounted Cube

Case 3 covers a  $120 \times 60$  mm domain above the top face of the cube ( $-60 \leq x \leq 60$  mm,  $120 \leq z \leq 180$  mm). The slab half-thickness is  $\delta = 3$  mm. The bottom boundary ( $z = 120$  mm, the top surface of the cube) is a wall where  $\mathbf{u} \cdot \mathbf{n} = 0$  is enforced. This is the most challenging validation case because it includes the separation at the leading edge of the cube top, the resulting shear layer, and the recirculation zone between the shear layer and the cube surface.

Figure 5.50 shows the velocity magnitude. The flow separates at the upstream edge of the cube top ( $x = -60$  mm,  $z = 120$  mm) and forms a shear layer that curves upward and to the right. Above the shear layer, the flow accelerates to approximately 12 to 14 m/s as it is compressed between the shear layer and the freestream. Below the shear layer, a low-velocity region extends along the cube surface, with velocities below 5 m/s. Both the single-element and multi-element reconstructions capture this structure, and the spatial pattern is consistent with the Hendriksen reference. The raw particle data shows the same structure but with large scatter, particularly in the low-velocity region where the turbulent fluctuations are large relative to the mean.


Figure 5.51 shows the vorticity. A band of strong negative vorticity traces the shear layer from the upstream separation point at ( $x \approx -60$  mm,  $z \approx 125$  mm) upward and to the right, reaching peak values of approximately  $-2500 \text{ s}^{-1}$  in the single-element reconstruction and  $-1500 \text{ s}^{-1}$  in the multi-element and Hendriksen fields. The vorticity in the freestream



**FIGURE 5.50** Velocity magnitude  $|u|$  for Case 3. Raw particle data, single-element ( $p = 8$ ,  $K = 1$ ), multi-element ( $p = 4$ ,  $K = 4$ ), and Hendriksen reference.  $\delta = 3$  mm.

above the shear layer is close to zero. The spatial pattern of the shear layer is consistent across all three fields. As in Case 1, the higher peak values in the single-element reconstruction reflect the higher spatial resolution compared to the bin-averaged reference.

5



**FIGURE 5.51** Vorticity  $\omega$  for Case 3. Single-element, multi-element, and Hendriksen reference.  $\delta = 3$  mm.

Figure 5.52 shows the divergence on a base-10 logarithmic scale. The single-element reconstruction averages  $\log_{10} |\nabla \cdot \mathbf{u}| \approx -12$  and the multi-element averages  $-11$ , both at machine precision. The Hendriksen reference averages  $\log_{10} |\nabla \cdot \mathbf{u}| \approx +1$ . The mimetic reconstructions are therefore roughly 12–13 orders of magnitude tighter on the divergence-free condition than the bin averaging. The nonzero divergence in the reference field is concentrated in the shear layer region where the velocity gradients are strongest. This is where accurate derived quantities (vorticity, pressure) are most needed, and it is precisely where the lack of a divergence-free constraint in the bin averaging introduces the largest errors.

Figure 5.53 shows the reconstructed streamfunction. The streamlines reveal the separation topology clearly. The main flow passes over the shear layer with compressed, closely spaced streamlines indicating the acceleration observed in the velocity magnitude plot. Near the bottom corners, closed streamline contours indicate recirculation zones. These recirculation regions correspond to the separated flow between the shear layer




FIGURE 5.52 Divergence  $\log_{10} |\nabla \cdot \mathbf{u}|$  for Case 3. Single-element, multi-element, and Hendriksen reference.  $\delta = 3$  mm.

and the cube surface. The downstream recirculation (near  $x \approx 50$  mm,  $z \approx 125$  mm) is larger than the upstream one, consistent with the asymmetry of the flow over a bluff body where the downstream separation region grows as the shear layer develops. Both the single-element and multi-element results show the same topology.

5

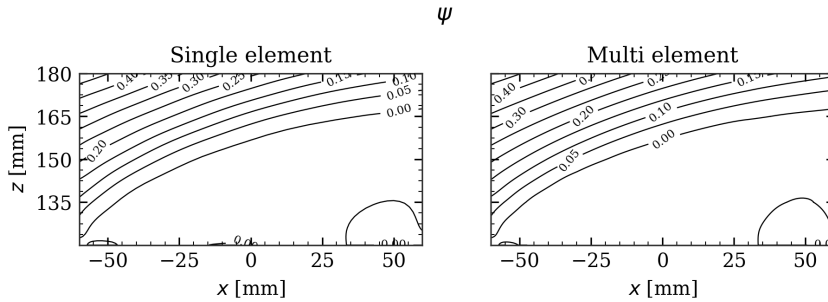



FIGURE 5.53 Streamfunction  $\psi$  for Case 3. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

Figure 5.54 shows the reconstructed total pressure. The lowest pressure values are found in the shear layer and recirculation zone near the bottom of the domain, where the velocity gradients are strongest and the vorticity reaches its peak values. Above the shear layer the total pressure is higher and more uniform. The single-element reconstruction produces values ranging from approximately  $-60$  to  $0 \text{ m}^2 \text{ s}^{-2}$ , while the multi-element result shows a similar spatial pattern with slightly larger range. Some grid artefacts are visible in the multi-element pressure field near the bottom-left corner, where the shear layer passes close to an element boundary.


Figure 5.55 shows the reconstructed static pressure. The field ranges from approximately  $-100$  to  $+50 \text{ m}^2 \text{ s}^{-2}$ . A low-pressure region is located in



**FIGURE 5.54** Total pressure  $\tilde{P}$  for Case 3. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

5

the lower-left of the domain, near  $(x \approx -40$  mm,  $z \approx 135$  mm), immediately above the upstream edge of the cube top where the flow accelerates around the leading corner and forms the separated shear layer. Higher static pressure is observed at the downstream-right edge of the domain, consistent with the gradual pressure recovery as the flow re-attaches above the cube. The single-element reconstruction produces a smoother field; the multi-element result captures the same low-pressure feature but shows visible element seams near the shear layer, where the gradients are strongest.



**FIGURE 5.55** Static pressure  $\tilde{p}_{\text{static}}$  for Case 3. Single-element and multi-element reconstruction.  $\delta = 3$  mm.

#### 5.2.4. Summary of Validation Findings

The validation results demonstrate that the mimetic spectral element reconstruction method produces physically consistent flow fields from experimental particle tracking data across three regions of increasing flow complexity.

The velocity fields agree well with the Hendriksen ensemble-averaged reference in all three cases. The spatial patterns of velocity magnitude, flow direction, deceleration zones, and shear layers are captured by both the single-element and multi-element configurations. The multi-element configuration shows some sensitivity to the ratio of measurement points per element, producing visible grid patterns in the vorticity and pressure when the number of particles per element becomes small relative to the number of degrees of freedom.

The divergence is the quantity where the mimetic approach shows its most distinctive advantage. The reconstructed velocity is divergence-free up to machine precision in all three cases, regardless of the flow complexity or the quality of the input data, with  $\log_{10} |\nabla \cdot \mathbf{u}|$  averaging  $-9$  to  $-12$  for the single-element and  $-10$  to  $-11$  for the multi-element configuration. The Hendriksen reference, which does not enforce incompressibility, averages  $\log_{10} |\nabla \cdot \mathbf{u}| \approx +1$  in the freestream (Case 1) and above the cube (Case 3), and  $\approx +2$  at the upstream cube base (Case 2) where the velocity gradients are strongest. The mimetic reconstructions are therefore roughly 11–13 orders of magnitude tighter on the divergence-free condition than the bin averaging, with the largest absolute gap precisely in the regions where accurate derived quantities are most needed.

The streamfunction, which exists only because the velocity field is divergence-free up to machine precision, reveals the flow topology in each case. In the freestream (Case 1) the streamlines are nearly parallel. At the cube base (Case 2) a recirculation zone is visible upstream of the cube. Above the cube (Case 3) the separation and recirculation structure is clearly resolved, with closed streamline contours identifying the separated flow regions. These topological features are a direct consequence of the structure-preserving property of the mimetic discretisation and could not be obtained from a velocity field with nonzero divergence.

The total and static pressures show physically plausible spatial distributions in all three cases, with the single-element reconstruction producing smoother fields than the multi-element configuration. The static pressure highlights stagnation rises upstream of the cube and low-pressure pockets at the upstream and upper leading edges where the flow accelerates around the corners. Of all reconstructed quantities, the pressure is the most affected by error amplification, since it depends on the nonlinear coupling  $\mathbf{u} \times \boldsymbol{\omega}$  and inherits errors from both the velocity and the vorticity. This sensitivity is consistent with the error analysis in Section 2.6., where

the pressure convergence rate was shown to be bounded by the vorticity force error (2.62).

# Chapter 6

## Conclusions and Recommendations

## 6.1. Conclusions

This thesis developed a constrained least-squares reconstruction method based on mimetic spectral elements for obtaining physically consistent velocity, vorticity, streamfunction, and pressure fields from scattered particle tracking data. The method was verified using a manufactured trigonometric solution and validated on three regions of separated flow over a surface-mounted cube. The main conclusions are summarised below, followed by the answers to the research questions posed in Section 1.2..

The most distinctive feature of the method is that the incompressibility constraint is satisfied exactly at the discrete level through the algebraic identity  $E^{2,1}\mathbf{u}_h = 0$ , enforced as a hard constraint in the saddle-point system. The reconstructed divergence is at machine precision in all configurations, both for the manufactured solution and for the experimental data. The Hendriksen ensemble-averaged reference, which does not enforce incompressibility, shows residual divergence up to  $\mathcal{O}(10^2) \text{ s}^{-1}$  in regions of strong velocity gradients. Because the velocity is divergence-free up to machine precision, a well-defined streamfunction exists and reveals the flow topology through closed contour lines that identify recirculation zones. This would not be possible from a velocity field with nonzero divergence.

6

All derived quantities are computed through the same discrete operators that preserve the de Rham complex, ensuring mutual consistency. The velocity lies in  $H(\text{div}; \Omega)$ , the vorticity is recovered through a weak curl formulation producing  $\omega \in H(\text{curl}; \Omega)$ , and the total pressure is obtained from a second saddle-point system involving the vorticity force  $\mathbf{u} \times \omega$ . The convergence rates under both p-refinement and h-refinement match the theoretical predictions from the error analysis, with a clear hierarchy that reflects the number of differentiations separating each quantity from the primary unknown. Under h-refinement at  $p = 7$ , the measured rates are  $\psi$  at 8.02,  $\mathbf{u}$  at 7.08,  $\omega$  at 6.37,  $\tilde{P}$  at 6.27, and  $H(\text{curl}; \Omega)$  at 5.24, confirming the predicted hierarchy from Table 2.3.

The validation on the surface-mounted cube demonstrates that the method captures flow deceleration, boundary layer development, separation, and recirculation. The single-element configuration ( $p = 8, K = 1$ ) produces smoother fields, while the multi-element configuration ( $p = 4, K = 4$ ) is more sensitive to the ratio of measurement points per element. The pressure is the most sensitive quantity in the reconstruction chain because it depends on the nonlinear coupling  $\mathbf{u} \times \omega$ , consistent with the error analysis.

These conclusions answer the three research questions as follows.

**SQ1: Formulation.** The velocity is placed in  $H(\text{div}; \Omega)$  with 1-form basis functions, enforcing normal flux continuity through the Piola transformation. The divergence-free constraint is imposed exactly through  $E^{2,1}\mathbf{u}_h = 0$ , and inter-element continuity is enforced through Lagrange multipliers. The vorticity and pressure are recovered through weak formulations that use the same discrete operators, ensuring consistency with the velocity.

**SQ2: Convergence.** Under p-refinement all quantities converge exponentially. Under h-refinement the rates follow the predicted hierarchy: each differentiation step reduces the rate by approximately one algebraic order, and each integration step gains one order. The sensitivity to data sparsity follows the same hierarchy, with the streamfunction being the most robust and the  $H(\text{curl}; \Omega)$ -norm the most sensitive.

**SQ3: Validation.** The method captures the essential flow features in all three validation cases and produces divergence-free velocity fields to machine precision from experimental data. The divergence-free velocity (up to machine precision) enables a well-defined streamfunction that reveals the flow topology, including recirculation zones at the cube base and above the cube top.

6

## 6.2. Recommendations

The method currently uses uniform Cartesian meshes. Non-uniform or locally refined meshes would allow the element size to be adapted to the local data density and flow complexity, placing more degrees of freedom in regions with strong gradients such as shear layers and boundary layers.

The reconstruction is restricted to two dimensions. The mimetic framework extends naturally to three dimensions through the three-dimensional de Rham complex, eliminating the need for the spanwise slab projection and allowing the method to be applied directly to the full volumetric particle tracking data.

The current formulation treats the measurement data as exact. Incorporating a statistical model of the measurement uncertainty into the least-squares formulation could improve the reconstruction in regions where the data density is low or the turbulent fluctuations are large relative to the mean velocity.

The pressure reconstruction depends on the vorticity through the nonlinear coupling  $\mathbf{u} \times \boldsymbol{\omega}$ , which makes it the quantity most affected by error amplification. Investigating alternative pressure formulations that reduce

this sensitivity, or incorporating pressure boundary conditions from the normal momentum equation at solid walls, could improve the pressure field for experimental data.

The method should be compared quantitatively against other reconstruction techniques such as VIC+ and FlowFit on the same experimental data to assess the relative advantages of exact constraint satisfaction versus approximate enforcement through penalty terms.

A systematic study of the slab thickness  $\delta$  and its effect on the two-dimensional assumption would provide practical guidelines for selecting  $\delta$  in future applications.

The validation was performed on time-averaged data. Extending the method to time-resolved reconstruction, where each instantaneous snapshot is reconstructed independently, would enable the study of unsteady flow phenomena such as vortex shedding and turbulent structures.

# References

- [1] Markus Raffel et al. *Particle Image Velocimetry: A Practical Guide*. 3rd. Springer, 2018.
- [2] D. Schanz, S. Gesemann, and A. Schröder. “Shake-The-Box: Lagrangian particle tracking at high particle image densities.” In: *Experiments in Fluids* 57.5 (2016), p. 70. DOI: 10.1007/s00348-016-2157-1.
- [3] F. F. J. Schrijer and F. Scarano. “Effect of predictor–corrector filtering on the stability and spatial resolution of iterative PIV interrogation.” In: *Experiments in Fluids* 45.5 (2008), pp. 927–941. DOI: 10.1007/s00348-008-0511-7.
- [4] Jan F. G. Schneiders and Fulvio Scarano. “Dense velocity reconstruction from tomographic PTV with material derivatives.” In: *Experiments in Fluids* 57.9 (2016), p. 139. DOI: 10.1007/s00348-016-2225-6.
- [5] Sebastian Gesemann et al. “From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties.” In: *18th International Symposium on Applications of Laser and Imaging Techniques to Fluid Mechanics*. Lisbon, Portugal, 2016.
- [6] John J. Charonko et al. “Assessment of pressure field calculations from particle image velocimetry measurements.” In: *Measurement Science and Technology* 21.10 (2010), p. 105401. DOI: 10.1088/0957-0233/21/10/105401.
- [7] B. W. van Oudheusden. “PIV-based pressure measurement.” In: *Measurement Science and Technology* 24 (2013), p. 032001.
- [8] Marc Gerritsma. “Edge functions for spectral element methods.” In: *Spectral and High Order Methods for Partial Differential Equations* (2011), pp. 199–207.
- [9] Jasper Kreeft, Artur Palha, and Marc Gerritsma. “Mimetic framework on curvilinear quadrilaterals of arbitrary order.” In: *arXiv:1111.4304* (2011).

- [10] Artur Palha et al. “Physics-compatible discretization techniques on single and dual grids, with application to the Poisson equation of volume forms.” In: *Journal of Computational Physics* 257 (2014), pp. 1394–1422.
- [11] L. A. Hendriksen, A. Sciacchitano, and F. Scarano. “Object registration techniques for 3D particle tracking.” In: *Measurement Science and Technology* 35.12 (2024), p. 125202. DOI: 10.1088/1361-6501/ad715c.
- [12] R. J. Adrian. “Twenty years of particle image velocimetry.” In: *Experiments in Fluids* 39 (2005), pp. 159–169.
- [13] J. Westerweel, G. E. Elsinga, and R. J. Adrian. “Particle image velocimetry for complex and turbulent flows.” In: *Annual Review of Fluid Mechanics* 45 (2013), pp. 409–436.
- [14] F. Scarano. “Iterative image deformation methods in PIV.” In: *Measurement Science and Technology* 13 (2002), R1–R19.
- [15] A. Melling. “Tracer particles and seeding for particle image velocimetry.” In: *Measurement Science and Technology* 8 (1997), pp. 1406–1416.
- [16] J. Westerweel and F. Scarano. “Universal outlier detection for PIV data.” In: *Experiments in Fluids* 39 (2005), pp. 1096–1100.
- [17] C. Rapp and C. Manhart. “Flow over periodic hills—Numerical and experimental study.” In: *Journal of Fluid Mechanics* 663 (2010), pp. 478–504.
- [18] M. Breuer et al. “Flow over periodic hills – Numerical and experimental results at large Reynolds numbers.” In: *Computers & Fluids* 38.2 (2009), pp. 433–457.
- [19] L. Li and Z. Pan. “Three-dimensional Lagrangian flow field reconstruction using constrained least squares and radial basis functions.” In: *arXiv preprint arXiv:2308.09227* (2023).
- [20] N. K. Chaurasia and S. Chakraborty. “Physics-informed neural networks for turbulent flow reconstruction from sparse measurements.” In: *Physics of Fluids* 36.8 (2024), p. 085187.
- [21] Claudio Canuto et al. *Spectral Methods: Fundamentals in Single Domains*. Springer, 2006.

- [22] Anthony T. Patera. "A spectral element method for fluid dynamics: Laminar flow in a channel expansion." In: *Journal of Computational Physics* 54.3 (1984), pp. 468–488.
- [23] George Em Karniadakis and Spencer J. Sherwin. *Spectral/hp Element Methods for Computational Fluid Dynamics*. 2nd. Oxford University Press, 2005.
- [24] Michel O. Deville, Paul F. Fischer, and Ernest H. Mund. *High-Order Methods for Incompressible Fluid Flow*. Cambridge University Press, 2002.
- [25] Ivo Babuška and Manil Suri. "The p and h-p versions of the finite element method, basic principles and properties." In: *SIAM Review* 36.4 (1994), pp. 578–632.
- [26] Christoph Schwab. *p- and hp-Finite Element Methods: Theory and Applications in Solid and Fluid Mechanics*. Oxford University Press, 1998.
- [27] David A. Kopriva. "A Polynomial Spectral Calculus for Analysis of DG Spectral Element Methods." In: *arXiv:1704.00709* (2017).
- [28] M. Duruflé, P. Grob, and P. Joly. "Influence of Gauss and Gauss–Lobatto Quadrature Rules on the Accuracy of a Quadrilateral Finite Element Method in the Time Domain." In: *Int. J. Numer. Methods Eng.* (2009).
- [29] Hao Li, Daniel Appelö, and Xiangxiong Zhang. "Accuracy of Spectral Element Method for Wave, Parabolic and Schrödinger Equations." In: *arXiv:2103.00400* (2021).
- [30] Francis X. Giraldo. "The Lagrange–Galerkin Spectral Element Method on Unstructured Quadrilateral Grids." In: *Journal of Computational Physics* 147 (1998), pp. 114–146.
- [31] V. Jain et al. "Construction and application of algebraic dual polynomial representations for finite element methods on quadrilateral and hexahedral meshes." In: *Computers and Mathematics with Applications* (2020). Available online. DOI: 10.1016/j.camwa.2020.09.022.
- [32] Daniele Boffi, Franco Brezzi, and Michel Fortin. *Mixed Finite Element Methods and Applications*. Vol. 44. Springer Series in Computational Mathematics. Springer, 2013.

- [33] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. "Finite element exterior calculus, homological techniques, and applications." In: *Acta Numerica* 15 (2006), pp. 1–155.
- [34] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. "Finite element exterior calculus: from Hodge theory to numerical stability." In: *Bulletin of the American Mathematical Society* 47.2 (2010), pp. 281–354.
- [35] Pavel B. Bochev and James M. Hyman. "Principles of mimetic discretizations of differential operators." In: *Compatible Spatial Discretizations, IMA Volumes in Mathematics and its Applications* 142 (2006), pp. 89–119.
- [36] Susanne C. Brenner and L. Ridgway Scott. *The Mathematical Theory of Finite Element Methods*. 3rd. Vol. 15. Texts in Applied Mathematics. Springer, 2008.
- [37] B. Wieneke and T. Rockstroh. "Lagrangian particle tracking in the presence of obstructing objects." In: *Measurement Science and Technology* 35.5 (2024), p. 055303. DOI: 10.1088/1361-6501/ad289d.
- [38] N. Agüera et al. "Ensemble 3D PTV for high resolution turbulent statistics." In: *Measurement Science and Technology* 27.12 (2016), p. 124011. DOI: 10.1088/0957-0233/27/12/124011.