DOUGLAS SERSON

Numerical study of wings with wavy leading and trailing edges

São Paulo 2017

DOUGLAS SERSON

Numerical study of wings with wavy leading and trailing edges

Tese submetida à Escola Politécnica da Universidade de São Paulo para obtenção do título de Doutor em Ciências

São Paulo 2017

DOUGLAS SERSON

Numerical study of wings with wavy leading and trailing edges

Tese submetida à Escola Politécnica da Universidade de São Paulo para obtenção do título de Doutor em Ciências

Área de concentração: Engenharia Mecânica

Orientadores: Prof. Spencer J. Sherwin Prof. Julio R. Meneghini

São Paulo 2017

Este exemplar foi revisado e corrigi responsabilidade única do autor e c	do em relação à versão original, sob om a anuência de seu orientador.
São Paulo, de	de
Assinatura do autor:	
Assinatura do orientador:	

Catalogação-na-publicação

Serson, Douglas Numerical study of wings with wavy leading and trailing edges / D. Serson -- versão corr. -- São Paulo, 2017. 159 p. Tese (Doutorado) - Escola Politécnica da Universidade de São Paulo.

Departamento de Engenharia Mecânica.

1.Aerodinâmica 2.Mecânica dos fluidos computacional 3.Métodos espectrais I.Universidade de São Paulo. Escola Politécnica. Departamento de Engenharia Mecânica II.t.

Acknowledgements

First, I would like to thank my supervisors Professor Julio Meneghini and Professor Spencer Sherwin for giving me this opportunity. This work would not be possible without their guidance and encouragement.

I also thank my family, friends and colleagues for all the support during these past three and a half years.

Finally, I gratefully acknowledge the funding I received for this research. During the first one year and a half of this PhD, when I was in Brazil, I was funded by FAPESP under grant 2012/23493-0. In the two years I spent in the UK I was funded by CNPq under grant 231787/2013-8. Most of the simulations presented here were performed in the archer national supercomputer using allocation provided by the UK Turbulence Consortium (EPSRC grant EP/L000261/1).

Abstract

Inspired by the pectoral flippers of the humpback whale, the use of spanwise waviness in wings has been considered in the literature as a possible way of delaying the stall, and possibly also reducing the drag coefficient, allowing for improved aerodynamic characteristics. In order to provide a better understanding of this flow control mechanism, the present work investigates numerically the effect of the waviness on the flow around infinite wings with a NACA0012 profile. The study consists of direct numerical simulations employing the spectral/hp method, which is available through the nektar++ library. Considering the high computational cost of the simulations performed, several improvements were introduced to the method, making it more efficient and allowing higher Reynolds numbers to be analysed. These improvements to the method include a coordinate transformation technique to treat the waviness, changes to the parallelism strategy, and an adaptive polynomial order refinement procedure. Initially, simulations were performed for a very low value of the Reynolds number Re = 1,000, allowing the three-dimensional flow structures to be observed in detail. In this case, the results show that the waviness leads to a decrease in the lift-to-drag ratio, accompanied by a strong reduction in the fluctuations of the lift force. The reduction in the lift-to-drag ratio is the combined effect of lower drag and lift forces, and is associated with a regime where the flow remains attached behind the peaks of the leading edge while there are distinct regions of flow separation behind the troughs. Then, simulations with Re = 10,000 were considered. For high angles of attack, the results for this case are similar to the lower Re, with the waviness leading to separation behind the troughs and reducing both the lift and the drag. However, for a lower angle of attack the waviness leads to a large increase in the lift coefficient. This was observed to be related to the fact that flow around the straight wing is laminar in this case, with the waviness inducing transition to a turbulent state. Finally, the case Re = 50,000 was considered, with the results showing a good agreement with experiments presented in the literature.

Keywords: Aerodynamics. Flow control. Computational fluid dynamics. Spectral/hp methods.

Resumo

Inspirado na nadadeira peitoral da baleia jubarte, o uso de ondulações ao longo da envergadura de asas tem sido considerado na literatura como uma possível maneira de atrasar o estol, e possivelmente também reduzir o arrasto, levando a melhores características aerodinâmicas. Com o objetivo de obter um melhor entendimento desse mecanismo de controle do escoamento, o presente trabalho investiga numericamente o efeito de ondulações no escoamento ao redor de asas infinitas com o perfil NACA0012. O estudo consiste de simulações diretas do escoamento usando o método espectral/hp, que está disponível através da biblioteca nektar++. Considerando o alto custo computacional das simulações realizadas, diversas melhorias foram introduzidas no método, tornando-o mais eficiente e permitindo que números de Reynolds mais elevados fossem analisados. Essas melhorias ao método incluem uma técnica de mudança de coordenadas para tratar a ondulação, mudanças na estratégia de paralelismo e um procedimento de refinamento usando ordem polinomial variável. Inicialmente, simulações foram realizadas para um número de Reynolds muito baixo Re = 1,000, o que permitiu observar as estruturas tridimensionais do escoamento em detalhe. Nesse caso, os resultados mostram que a ondulação leva a uma diminuição da razão sustentação-arrasto, combinada com uma forte redução das flutuações da força de sustentação. A redução da razão sustentação-arrasto é consequência de uma combinação de arrasto e sustentação mais baixos e está associada a um regime no qual o escoamento permanece colado atrás dos picos do bordo de ataque, enquanto que regiões distintas de escoamento separado estão presentes atrás dos vales. Em seguida, simulações com Re = 10,000 foram consideradas. Para ângulos de ataque elevados, os resultados neste caso são similares àqueles com Re mais baixo, com a ondulação levando a separação atrás dos vales e provocando reduções na sustentação e no arrasto. No entanto, para um ângulo de ataque mais baixo a ondulação leva a um grande aumento na força de sustentação. Foi observado que isso está relacionado ao fato de que o escoamento ao redor da asa lisa é laminar neste caso, com a ondulação induzindo a transição para um estado turbulento. Finalmente, o caso Re = 50,000 foi considerado, com os resultados apresentando uma boa concordância com experimentos apresentados na literatura. Palavras-chave: Aerodinâmica. Controle (Escoamento). Mecânica dos fluídos computacional. Métodos espectrais.

List of Figures

1.1 1.2	Classification of flow control strategies. Extracted from Gad-el Hak (1996) Changes in the flow and their effect in engineering goals. Extracted from	4
	Gad-el Hak (2000).	4
1.3	Passive flow control mechanisms.	5
1.4	Reynolds number range for flight vehicles. Extracted from Mueller and De-	
	Laurier (2003)	8
2.1	On the left, picture of a humpback whale, extracted from Custodio (2007).	
	On the right, planform of a pectoral flipper, extracted from Fish and Battle	
	(1995)	12
2.2	On the left, models studied by Miklosovic et al. (2007). On the right, lift	
	coefficient results. The top figures correspond to the full-span model (two-	
	dimensional analysis), while the bottom figures correspond to the half-span	10
0 0	On the left, regults obtained by the model of yer Nieron et al. (2007).	13
2.3	the right experimental results from Johari et al. (2007). The symbols rep-	
	resent cases with 4 wavelengths along the span, while the lines represent	
	cases with 8 wavelengths along the span. The parameter η represents the	
	amplitude of the bumps. Extracted from van Nierop et al. (2008).	14
2.4	Comparison between instantaneous vorticity magnitude contours for a smooth	
	and a wavy wing. Extracted from Pedro and Kobayashi (2008).	14
2.5	Recirculation zones for $Re=800$ and $\alpha=20^{\circ}$. The first case in the left	
	corresponds to a straight wing, while the other cases correspond to different	
	combinations of amplitude and wavelength. Adapted from Favier et al. (2012).	15
2.6	Time-averaged streamlines illustrating physical structures induced by the	
	waviness, with colours representing pressure coefficient. Extracted from	
	Skillen et al. (2015)	16
2.7	Blunt trailing edge body studied by Bearman and Tombazis (1993)	17
4.1	Geometry of a wing and part of the domain in the different coordinate sys-	
	tems considered.	36

4.2	Convergence on the infinity norm of the u velocity error for the Kovasznay flow with different mappings	11
13	Contours of the w velocity component of the Kovasznav flow in the xz plane	
ч.0	for the simulations with no mapping and with the mapping of equation (4.32)	15
11	Contours of spanwise vorticity for square cylinder with and without spanwise	40
т.т	waviness with $B_c = 100$. The waviness has a wavelength $\frac{\lambda}{2} = 2.8$ and	
	waviness with $Re = 100$. The waviness has a wavelength $\frac{1}{H} = 2.8$ and applitude $h = 0.2$	16
4 5	Comparison of results for a square sylinder with waviness of wavelength λ –	40
4.5	Comparison of results for a square cylinder with wavelength $\frac{1}{H}$ =	
	2.8 at $Re = 100$, with the results norm bareful and Sherwin (2001b). Each	47
46	Lastenteneous contours of verticity for flow around two circular cylinder in	47
4.0	instantaneous contours of volticity of now around two circular cylinder in tandem with $P_{0} = 100$. The downetroom cylinder is fixed, while the upstream	
	tandem with $Re = 100$. The downstream cylinder is fixed, while the upstream over the construction of Re	10
47	Cylinder oscilates with frequency 0.3 and amplitude 0.75	48
4.7	Detail of computational mesh used in the simulation of the neordinate transfer	
	mation at two different time instante	40
		49
5.1	Illustration of different parallelization strategies using four processes. Ex-	
	tracted from Bolis (2013)	57
5.2	Scaling of the parallel simulations for the straight wing with $Re=10,000.$	59
5.3	Scaling of the parallel simulations for wavy wing with $Re = 10,000.$	61
5.4	Efficiency of the parallelization, compared to the reference $P_Z = 8$, $P_{XY} = 6$.	
	Simulations with $Re = 10,000.$	62
5.5	Isocontours of the efficiency of the parallelization (colours), for simulations	
	with $Re = 10,000$. The solid lines represent a constant number of processes	
	P_{TOT}	63
5.6	Polynomial order distribution obtained with adaptive order procedure. Simu-	
	lation with $Re = 50,000$ and $\alpha = 15^{\circ}$.	66
6.1	Geometry of a wavy wing with $h/c = 0.1$ and $\lambda/c = 0.5$.	72
6.2	Results obtained for a straight wing using two-dimensional and three-dimension	al
	simulations.	74
6.3	Comparison between L/D results obtained for wavy wings with different	
	wavelengths with the results from the reference straight wing.	75
6.4	Results for wavy wings with wavelength $\lambda/c = 0.5$.	76
6.5	Instantaneous skin-friction lines on the wall for the baseline, L05h05 and	
	L05h10 cases. The flow is from left to right and the colours represent the	
	orientation of the skin-friction in the chord direction, with blue corresponding	
	to reversed flow.	78
6.6	Recirculation zones for different cases with $\alpha = 12^{\circ}$.	79

6.	7 Pressure coefficient and streamwise tangential pressure gradient on the wing	
	surface for the case L05h10, with $lpha = 12^{\circ}$. The two-dimensional baseline	
	result at the same angle of attack is also presented as a reference.	81
6.	.8 Contours of spanwise pressure gradient and velocity at $z = 0.125$. The z	
	direction points out of the page. Case L05h10, with $\alpha = 12^{\circ}$	82
6.	.9 Contours of spanwise velocity (colours) at $x = -0.03$ and pressure on the	
	wing surface (grayscale). Case L05h10, with $\alpha = 12^{\circ}$.	83
6.	.10 Contours of spanwise velocity and streamwise vorticity(colours) at $x = 0.2$	
	and pressure on the wing surface (grayscale). Case L05h10, with $\alpha = 12^{\circ}$.	84
6.	.11 Comparison of results for L/D obtained using waviness only on leading edge	
	or trailing edge with the previous results.	85
6.	.12 Recirculation zones for cases with waviness only on leading edge or on trail-	
	ing edge, with $\alpha = 12^{\circ}$.	86
6.	.13 Comparison between lift to drag ratio of case L05h05_LE keeping a constant	
	thickness with the same case keeping the profile constant.	86
7.	.1 Time-averaged force coefficients for simulations with $Re = 10,000.$	88
7.	.2 Comparison between time series and spectra of the lift coefficient of the	
	L05h10 geometry with the baseline wing, for simulations with $Re = 10,000$.	89
7.	.3 Comparison between spectra of the lift coefficient of different geometries, for	
	$\alpha = 18^{\circ} \text{ and } Re = 10,000.$	90
7.	4 Time-averaged skin-friction lines on the wall for simulations with $Re = 10,000$.	
	The flow is from left to right and the colours represent the orientation of the	
	skin-friction in the chord direction, with blue corresponding to reversed flow.	92
7.	.5 Instantaneous recirculation regions for simulations with $Re = 10,000.$	93
7.	6 Time-averaged recirculation regions for simulations with $Re = 10,000.$	94
7.	.7 Time-averaged recirculation regions for simulations with $\alpha = 18^{\circ}$ and $Re =$	
	10,000	95
7.	.8 Time-averaged surface pressure coefficient distribution for simulations with	
	$\alpha = 6^{\circ} \text{ and } Re = 10,000. \dots$	96
7.	.9 Time-averaged surface pressure coefficient distribution for simulations with	
	$\alpha = 12^{\circ} \text{ and } Re = 10,000.$	97
7.	.10 Time-averaged surface pressure coefficient distribution for case L05h10 with	
	$\alpha = 18^{\circ} \text{ and } Re = 10,000.$	97
7.	.11 Comparison between u velocity profiles along the y direction for baseline and	
	L025h05 geometries with $\alpha = 6^{\circ}$ and $Re = 10,000$	99
7.	.12 Comparison between u velocity profiles along the y direction for baseline and	
	L025h05 geometries with $\alpha = 12^{\circ}$ and $Re = 10,000.$	100

7.13	Contours of time-averaged spanwise velocity at $z = 0.125$. Case L05h10, with $\alpha = 12^{\circ}$ and $Re = 10,000$.	101
7.14	Time-averaged streamlines projected in the xy planes at different sections for the L025b05 and L05b10 geometrics with $x = 12^{\circ}$ and $P_{0} = 10,000$. The	
	colours represent the orientation of the chordwise velocity	102
8.1	Results of force coefficients for simulations with $Re = 50,000$. For the time-	
	averaged lift coefficient, the solid lines are the experimental results from	
	Paula (2016) and the dashed line is the 2π slope.	105
8.2	Spectra of lift coefficient for simulations with $Re = 50,000.$	106
8.3	Moving averages of the lift coefficient for simulations with $Re = 50,000$ and	
	$\alpha = 15^{\circ}$.	106
8.4	Time-averaged skin-friction lines on the wall for simulations with $Re = 50,000$.	
	The flow is from top to bottom.	108
8.5	Time-averaged recirculation regions for simulations with $Re = 50,000.$	109
8.6	Oil visualization from experiments of Paula (2016) with $Re = 50,000$. The	
	flow is from top to bottom.	110
8.7	Surface pressure coefficient distribution for simulations with $Re = 50,000.$	111
8.8	Location of the probes on the xy plane	112
8.9	Spectra of <i>u</i> -velocity fluctuations for simulations of the baseline wing with	
	$Re = 50,000$. The dashed line corresponds to the $-\frac{5}{3}$ slope	113
8.10	Spectra of $u-velocity fluctuations for simulations of the wavy wing with \alpha=6^\circ$	
	and $Re = 50,000$. The dashed line corresponds to the $-\frac{5}{3}$ slope	114
8.11	Time-averaged contours of spanwise velocity at $z = 0.125$ for L05h10 wing	
	with $\alpha = 6^{\circ}$ and $Re = 50,000$.	115
8.12	Slices with contours of instantaneous spanwise vorticity (on the left) and of	
	turbulence kinetic energy (on the right) for simulations with $\alpha = 6^{\circ}$	116
8.13	Iso-surface of $TKE=0.015$ for wavy wing with $\alpha=6^\circ$ and $Re=50,000.$.	117
8.14	Recirculation region filtered using an exponential moving average with differ-	
	ent final times. Simulation with $\alpha = 15^{\circ}$ and $Re = 50,000.$	118
8.15	Recirculation regions for flow obtained using different time-averaging inter-	
	vals. Wavy wing with $\alpha = 15^{\circ}$ and $Re = 50,000$.	119
B.1	Initial mesh used in the convergence tests for $Re = 1,000$. On the left the	
	entire domain is shown, while on the right only the region close to the airfoil	
	is presented.	146
B.2	Final mesh obtained from the convergence tests for $Re = 1,000.$	149
B.3	Detail of the mesh used for simulations with $Re = 10,000.$	151

B.4	Moving average of the lift coefficient for different time-averaging intervals,	
	described by the width of the interval T_{ma} and by the final time. The dotted	
	lines represent 2% deviations from the value 0.76	152
B.5	Comparison of the lift coefficient obtained using the chosen averaging inter-	
	val to the result with a short interval, which is close to the instantaneous lift	
	coefficient	153
B.6	Modal energy for simulation of baseline wing with $Re = 10,000$ and $\alpha = 18^{\circ}$.	154
B.7	Detail of the mesh used for simulations with $Re = 50,000.$	155
B.8	Modal energy for simulations of baseline wing with $Re = 50,000.$	155
C.1	Enstrophy for the Taylor-Green Vortex simulations with different configurations	.158
C.2	Iso-surfaces of the second invariant of the velocity gradient tensor $Q = 10$	
	coloured by u velocity for simulations of the Taylor-Green Vortex at $t = 10$.	
	Only one eighth of the domain is displayed.	159

List of Tables

3.1	Coefficients for stiffly-scheme time integration.	31
5.1 5.2	Parameters for the spectral-vanishing viscosity employed in the present work. Comparison of the computational cost of adaptive order case of figure 5.6 with constant uniform polynomial order and with variable order without adap- tive procedure. The computational costs are normalized with respect to the adaptive order case.	55 67
6.1	Parameters of the waviness for the cases analyzed.	73
8.1 8.2	Parameters used in simulations with $Re = 50,000$. L_z is the spanwise length and ϵ is the SVV diffusion coefficient	103 112
B.1	Convergence tests for the simulation time with $\alpha = 5^{\circ}$.	147
B.2	Convergence tests for the simultion time with $\alpha = 18^{\circ}$	147
B.3	Convergence tests for basis functions polynomial order with $lpha=5^\circ.$	148
B.4	Convergence tests for basis functions polynomial order with $\alpha = 18^{\circ}$	148
B.5	Convergence tests for domain size with $\alpha = 5^{\circ}$.	148
B.6	Convergence tests for domain size with $\alpha = 18^{\circ}$.	149
B.7	Convergence tests for time-step for simulations with $\alpha = 5^{\circ}$	150
B.8	Convergence tests for the time-step for simulations with $\alpha = 18^{\circ}$	150
B.9	Convergence tests for number of Fourier modes in the span direction. Simu-	
	lations with $\alpha = 5^{\circ}$.	150
B.10	Convergence tests for number of Fourier modes in the span direction. Simu-	
	lations with $\alpha = 18^{\circ}$	151
B.11	Convergence tests for number of Fourier modes in the span direction. Simu-	
	lations with $Re = 10,000$	153
B.12	Convergence tests for basis functions polynomial order. Simulations with	
	Re = 10,000.	153

Nomenclature

- $(\bar{x}, \bar{y}, \bar{z})$ Transformed (computational) coordinate system
- (ξ_1,ξ_2) Coordinates in standard quadrilateral element
- (u, v, w) Velocity components
- (x, y, z) Cartesian coordinate system

α Angle of attack

- α_q Coefficients for backward differentiation formula
- $\bar{\mathbf{G}}(p)$ Generalized pressure gradient
- $\bar{\mathbf{L}}(\mathbf{u})$ Generalized viscous terms
- $\bar{\mathbf{N}}(\mathbf{u})$ Generalized convective terms
- $ar{u}^i$ Components of vector ${f u}$ in the transformed coordinate system
- \bar{V}^{j} Components of velocity of coordinate system
- β Wave-number for Fourier expansion
- β_q Coefficients for polynomial extrapolation
- Δt Time step
- ϵ Spectral-vanishing viscosity diffusion coefficient
- ϵ^{ijk} Permutation symbol
- ϵ_l Tolerance for reducing the order in adaptive procedure
- ϵ_u Tolerance for increasing the order in adaptive procedure
- η Amplitude of waviness from van Nierop et al. (2008)
- η_p Parallelism efficiency

- $\Gamma_{\mathfrak{D}}$ Boundary with Dirichlet conditions
- $\Gamma_{\mathfrak{N}}$ Boundary with Neumann conditions
- Γ Boundary of the domain
- γ_0 Coefficient for backward differentiation formula
- Û Matrix representing the spectral-vanishing viscosity filter in modal space
- \hat{u}^i Coefficients of the solution
- \hat{u}_q Global coefficients
- \hat{u}_l Local coefficients in all elements
- κ Constant for transformation restricted to leading or trailing edge
- λ Waviness wavelength
- $\bar{\mathbf{u}},\,\hat{\mathbf{u}},\,\hat{\bar{\mathbf{u}}}\,$ Intermediate velocity fields
- $\mathbf{M}_{\mathbf{e}}$ Concatenation of elemental mass matrices
- A(u, p) Forcing term for explicit treatment of mappings
- B Matrix with values of basis functions at quadrature points
- \mathbf{D}_i Matrix to calculate derivative in *i* direction in physical space
- F General vector field
- J Diagonal matrix containing Jacobian of the transformation in quadrature points
- $\mathbf{L}(\mathbf{u})$ Linear viscous terms
- $\mathbf{M}_{\mathbf{g}}$ Global mass matrix
- \mathbf{N}^* Polynomial extrapolation of operator \mathbf{N}
- $\mathbf{N}(\mathbf{u})~$ Non-linear convective terms
- Q Generalized curl-curl operator
- \mathbf{S}_{VV}^{e} Spectral-vanishing viscosity operator in a single element
- \mathbf{u}^+ Explicit part of backward differentiation formula
- u Velocity vector

- \mathbf{u}^0 Initial condition for velocity vector
- $\mathbf{u}_{\mathfrak{D}}$ Dirichlet boundary conditions
- $\mathbf{u}_{\mathfrak{N}}$ Neumann boundary conditions
- W Diagonal matrix containing quadrature weights
- \mathbf{W}_{std} Diagonal matrix containing quadrature weights in standard element
- \mathcal{A} Connectivity map for global assembly
- μ Coefficient for Helmholtz equation
- Ω Domain
- $\overline{C_d}$ Mean drag coefficient
- $\overline{C_l}$ Mean lift coefficient
- Φ^i Basis functions
- ϕ_p Legendre polynomial basis
- Φ_{pqr} Three-dimensional expansion basis
- $\Phi_{\it pq} ~~ {\rm Two-dimensional\ expansion\ basis}$
- ψ_p Modified polynomial basis
- σ^{ij} Stress tensor
- $\tilde{\mathbf{B}}$ Matrix with values of orthogonal basis functions at quadrature points
- $\tilde{\mathbf{S}}_i$ Matrix for obtaining derivative in *i* direction in the orthogonal basis
- ε^{ijk} Generalization of the permutation symbol
- ξ Coordinate in standard element
- ξ_i Nodes for Gaussian quadrature
- ξ_q Nodes for Lagrange polynomials
- *c* Mean chord length
- C_p Pressure coefficient
- C'_l Lift coefficient fluctuation (*rms*)

- *C*_{pb} Base pressure coefficient
- $D(\mathbf{u})$ Generalized divergence operator
- dS Area element
- *f* Forcing term for Helmholtz equation
- F_i Righ-hand side of discretized Helmholtz equation
- g^{ij} Inverse of metric tensor
- g_{ij} Metric tensor
- *H* Square cylinder height
- *h* Waviness peak-to-peak amplitude
- h_p Lagrange polynomial basis
- J Jacobian of a transformation
- *k* Constant for the Kovasznay flow
- *L* Reference length
- L/D Lift-to-drag ratio
- L_p Legendre polynomials
- L_z Spanwise length of the domain
- L_{ij} Weak Laplacian matrix
- M_{ij} Mass matrix
- $N_{dof,H}$ Number of degrees of freedom for the homogeneous solution
- N_{dof} Number of degrees of freedom
- N_{el} Number of elements in the mesh
- n_i Normal vector
- N_{modes} Number of modes of polynomial expansion
- *n_{runs}* Number of repetitions of adaptive procedure
- n_{steps} Number of time steps for adaptive polynomial order

- P Polynomial order
- *p* Pressure field
- P_1, P_2 Polynomial orders in two dimensions
- *P_{cut}* Cut-off for spectral-vanishing viscosity
- Pmax Maximum polynomial order in adaptive procedure
- *P_{min}* Minimum polynomial order in adaptive procedure
- $P_p^{\alpha,\beta}$ Jacobi polynomials
- P_{TOT} Total number of parallel processes
- P_{XY} Partitions in xy plane
- P_Z Groups of processes for Fourier parallelism
- *Q* Number of quadrature points in one dimension
- Q_1, Q_2 Number of quadrature points in two dimensions
- *Q_i* Spectral-vanishing viscosity filter
- Q_{st} Standard quadrilateral element
- $R(u^{\delta})$ Residual from approximation
- *Re* Reynolds number
- rms Root mean square
- S_e Adaptive refinement sensor in an element
- *S_{VV}* Spectral-vanishing viscosity operator
- St Strouhal number
- *T_{ma}* Window size for moving average
- *U* Reference velocity
- *u* Unknown for Helmholtz equation
- u^{δ} Approximate solution
- u^D Part of the solution satisfying Dirichlet conditions

- u^H Homogeneous part of the solution
- u_P Velocity u represented with P^{th} order polynomials
- v^j Test functions
- w_i Weights for Gaussian quadrature
- PSD Power spectral density
- TKE Turbulence kinetic energy
Contents

I	Intr	roduction	1			
1	Introduction					
	1.1	Flow control	3			
	1.2	Flow regimes for wings at low Reynolds numbers	6			
	1.3		6			
	1.4	Publications	8			
	1.5	Thesis outline	9			
2	Lite	rature review	11			
	2.1	Use of leading edge waviness	11			
	2.2	Use of trailing edge waviness	16			
II	Nu	imerical methods	19			
3	The	spectral/hp element method	21			
	3.1	The Galerkin Method	21			
	3.2	Local formulation	23			
		3.2.1 Expansion basis	23			
		3.2.2 Operations in standard region	26			
		3.2.3 Local operations	28			
	3.3	Global assembly and global operations	29			
	3.4	Time integration of the Navier-Stokes equations	30			
4	Trea	atment of the wavy geometries	35			
	4.1	Explicit treatment of the mapping	37			
	4.2	Semi-implicit treatment of the mapping	40			
	4.3	Post-processing and forces calculation	42			
	4.4	Test cases	43			
		4.4.1 Kovasznay flow	43			
		4.4.2 Square cylinder with waviness	45			
		4.4.3 Flow around moving cylinders	47			

	4.5	Discussion	49		
5	5 Numerical considerations for high Reynolds number simulations				
	5.1	The spectral-vanishing viscosity	53		
	5.2	Parallelization strategy	56		
	5.3	Adaptive order refinement	64		
111	R	esults	69		
c	NI	norical simulations with Day 1,000	74		
O	NUN 6 1	Accodynamic forces	71		
	0.1 6.2	Separation and regiroulation zonos	75		
	0.2 6.3		70		
	0.0 6.4		80		
	0.4 6 5	Effect of waviness with different shapes	83		
	0.0		00		
7	Nun	nerical simulations with ${f Re}=10,000$	87		
	7.1	Aerodynamic forces	87		
	7.2	Separation and recirculation zones	91		
	7.3	Surface pressure distribution	93		
	7.4	Velocity profiles	98		
	7.5	Discussion	98		
8	Nun	nerical simulations with ${f Re}={f 50},000$	103		
	8.1	Aerodynamic forces	104		
	8.2	Separation and recirculation zones	107		
	8.3	Surface pressure distribution	109		
	8.4	Spectra of velocity fluctuations	112		
	8.5	Discussion	114		
IV	C	onclusions and future work	121		
9	Cor	iclusions	123		
	9.1	Numerical methods	123		
	9.2	Effect of spanwise waviness on the flow	125		
10 Suggestions for future work 127					
Bibliography 12					
A	Appendices 13				

Α	Tens	sor calculus	137
	A.1	Coordinate systems transformations	137
	A.2	Contravariant vectors	138
	A.3	Covariant vectors	138
	A.4	Higher order tensors	138
	A.5	The metric tensor	139
	A.6	The covariant derivative	140
	A.7	Divergence and curl of vectors	141
	A.8	The generalized Navier-Stokes equations	142
	A.9	Application to wings with waviness	142
в	Mes	h generation and convergence tests	145
В	Mes B.1	h generation and convergence tests Convergence tests for $Re = 1,000$	145 145
в	Mes B.1	h generation and convergence testsConvergence tests for $Re = 1,000$ B.1.1Simulation time	145 145 146
В	Mes B.1	h generation and convergence tests Convergence tests for Re = 1,000 B.1.1 Simulation time B.1.2 Basis functions polynomial order	145 145 146 147
В	Mes B.1	h generation and convergence tests Convergence tests for Re = 1,000 B.1.1 Simulation time B.1.2 Basis functions polynomial order B.1.3 Domain size	145 145 146 147 147
В	Mes B.1	h generation and convergence tests Convergence tests for Re = 1,000 B.1.1 Simulation time B.1.2 Basis functions polynomial order B.1.3 Domain size B.1.4 Time-step	145 145 146 147 147 149
В	Mes B.1	h generation and convergence tests Convergence tests for Re = 1,000 B.1.1 Simulation time B.1.2 Basis functions polynomial order B.1.3 Domain size B.1.4 Time-step B.1.5	145 145 146 147 147 149 150
В	Mes B.1 B.2	h generation and convergence testsConvergence tests for $Re = 1,000$ B.1.1 Simulation timeB.1.2 Basis functions polynomial orderB.1.3 Domain sizeB.1.4 Time-stepB.1.5 Number of Fourier modes in the span directionConvergence tests for $Re = 10,000$	145 145 146 147 147 149 150 150
В	Mes B.1 B.2 B.3	h generation and convergence testsConvergence tests for $Re = 1,000$ B.1.1 Simulation timeB.1.2 Basis functions polynomial orderB.1.3 Domain sizeB.1.4 Time-stepB.1.5 Number of Fourier modes in the span directionConvergence tests for $Re = 10,000$ Tests for $Re = 50,000$	 145 146 147 147 149 150 150 154

Part I

Introduction

1. Introduction

1.1 Flow control

The development of wings with improved aerodynamic performance is a widely researched subject. This is motivated mainly by the aeronautical industry, where even a small reduction in the drag or an increase in the lift forces can lead to significant fuel savings. However, these improvements are also important in many other applications, like wind and gas turbines, the development of rotor blades, and car racing competitions.

In this context, an important role has always been played by flow control, which can be defined as the use of devices causing a beneficial change to the flow (Gad-el Hak, 1996). For the flow around wings, the benefits we want usually involve an increase of the lift force and a decrease of the drag force.

In a broad sense, flow control mechanisms can be classified in active and passive devices, as illustrated in figure 1.1. The difference between these categories is that active devices require energy expenditure, while passive devices do not. When compared to active control, the main advantage of the passive mechanisms is their simplicity, which leads to a higher reliability, lower weight, and lower costs.

Flow control mechanisms can also be classified considering the effect they have on the flow. In this sense, we can have, for example, mechanisms affecting boundary layer separation, suppressing or inducing transition to turbulence, inducing the growth or decay of structures, or affecting the large scales of the flow. Figure 1.2 illustrates the impact a few of these changes can have on the performance of a wing. Another way of classifying the passive flow control mechanisms is according to the length scale they are designed to influence. A particular mechanism may act directly inside the boundary layer; it may act on scales slightly larger than the boundary layer with the purpose of influencing the boundary layer; or it may act on the large scales of the flow, possibly by modifying the potential flow over the wing.

A few examples of passive flow control mechanisms are the fences (Hieser, 1953), the vortex generators (Hieser, 1953; Storms and Jang, 1994), and the Gurney flap (Neuhart and Pendergraft, 1988; Storms and Jang, 1994). The goal of using these three devices is to obtain a delay in the stall, leading to a higher maximum lift coefficient. These mechanisms are illustrated in figure 1.3.



Figure 1.1: Classification of flow control strategies. Extracted from Gad-el Hak (1996).



Figure 1.2: Changes in the flow and their effect in engineering goals. Extracted from Gad-el Hak (2000).



(b) Gurney flap (Storms and Jang, 1994)

(a) Vortex generators (Shun and Ahmed, 2012)

Entrained Airflow

Vortex Generator

Vortex



Wing

Deployed Flap

(c) Fence (Williams et al., 2010)

Figure 1.3: Passive flow control mechanisms.

The present work is concerned with the use of waviness along the span of wings, which is proposed in the literature as another form of passive flow control capable of improving the aerodynamic performance. For example, the use of different trailing edge shapes was proposed by Petrusma and Gai (1989) as a way to reduce the drag in bodies with a blunt trailing edge. The references available in the literature on the effect of changes in shape of the trailing edge concern only these type of bodies, so the effect on bodies with sharp trailing edges is not yet known. More recently, several works have considered the effects of using waviness in the leading edge. Most of these studied focus on low Reynolds number (Re) on the order of 10^5 . This can be explained by the recent interest in wind turbines and micro-air-vehicles (MAV), which often operate at these low Re.

As will be discussed in chapter 2, the inspiration for using wavy leading edges comes from observing the shape of the humpback whale flippers, which is covered by tubercles in the leading edge region. Therefore, this can also be seen as part of the recent trend towards biomimetics, which is the use of engineering solutions based on highly-efficient systems encountered in nature. Another example of this is the development of poroelastic trailing edges inspired in the serrated wings of owls, in an attempt to reduce noise levels (Jaworski and Peake, 2013).

1.2 Flow regimes for wings at low Reynolds numbers

The present study considers a wide range of Reynolds numbers, extending over more than one order of magnitude of what are considered low Reynolds numbers. In order to provide some insight into how the flow around an airfoil changes as Re is increased, we now consider a brief overview of the different flow regimes for Re < 200,000, based on the summary from Carmichael (1981). These regimes are:

- Fractional *Re* : the flow is dominated by viscosity, and is outside the range of interest for the present work.
- Re < 150: this regime is of interest for turbulence reducing screens, where a minimum disturbance to the flow is desired.
- 1,000 < Re < 10,000: this regime is observed in nature for many insects. The flow is strongly laminar, being difficult to artificially transition the boundary layer. However, we note that at high angles of attack, when the airfoil behaves like a bluff body, it is possible to observe transition in the wake.
- 10,000 < Re < 30,000: according to Carmichael (1981), the boundary layer is laminar and artificial tripping has not been successful in this regime. The lift coefficient is restricted, since attempting to increase the angle of attack leads to separation without reattachment.
- 30,000 < Re < 70,000: Carmichael (1981) considers this to be the regime of most interest for technically oriented model aircraft builders. Reattachment of the separated shear layer can occur in the upper part of this range, although in the lower part the distance to the trailing edge is insufficient for reattachment. Boundary layer tripping also becomes possible.
- 70,000 < Re < 200,000: extensive laminar flow is easy to obtain and performance improves markedly in this regime. The laminar separation bubble can still cost some performance.

1.3 Contributions

In order to investigate the effect of spanwise waviness on the flow around wings, numerical simulations were performed for infinite wings with sinusoidal waviness. The simulations employed the spectral/hp method, which is a high order finite element method available

through the Nektar++ library (Cantwell et al., 2015). Considering that the computational cost of some of these simulations is large, a secondary goal of the present study was to develop improvements to the method. These improvements were focused on obtaining the solution more efficiently, extending the envelope of viable simulations, and ultimately allowing the study to consider Reynolds numbers for which it was previously not feasible to obtain a solution in a reasonable time.

The main contribution of the present work in terms of the numerical method was the development and implementation of an approach to include general coordinate transformations into the spectral/hp method. This allowed the infinite wavy wings to be mapped into straight wings, which could then be studied using a quasi-3D approach, leading to significant savings in terms of computational cost. Then, an investigation into the parallelism of the code was performed, allowing the simulations to run efficiently on modern supercomputers. These are in most cases distributed-memory clusters which increasingly rely on parallelism, and therefore it is important to understand and improve the scaling properties of the numerical methods employed. In addition, an adaptive polynomial order procedure was developed, which allowed the simulations to distribute the degrees of freedom in a more efficient manner, leading to lower computational costs for the same accuracy. These contributions to the numerical method were implemented and incorporated into Nektar++ as part of the development of this open-source project.

The numerical simulations initially considered a very low Reynolds number Re = 1,000. Although this is too low for most practical applications, the computational cost in this case is relatively low, allowing for a larger parameter space to be considered when compared to higher Re, and also serving as a good test for the methodology being employed. Another advantage of these simulations is that at this low Re the flow is laminar, and therefore it is easier to understand the main flow characteristics, serving as a guide for the rest of the work. The study was then extended to Re = 10,000 and Re = 50,000. This last case is of significant importance for two reasons. First, it allowed comparisons with experimental data available in Paula (2016). Furthermore, as illustrated in figure 1.4, this value of Re is close to the typical values encountered for MAVs, bringing the present work closer to practical applications. Considering the discussion of section 1.2, we note that these simulations cover different flow regimes, starting at a fully laminar flow and moving towards a regime where a laminar-separation bubble is expected to occur.

From the results of the numerical simulations, three important physical mechanisms by which the waviness affects the flow were identified. The first mechanism is the weakening of the suction peak on the sections corresponding to the waviness peaks. This leads to weaker adverse pressure gradients, explaining why the flow tends to remain attached on these sections. As a second mechanism, the waviness has a significant effect on the stability of the separated shear layers, with transition occurring earlier for the wavy wing. The last mechanism corresponds to a sub-harmonic behaviour, with the flow being notably different



Figure 1.4: Reynolds number range for flight vehicles. Extracted from Mueller and DeLaurier (2003).

between subsequent wavelengths in some cases. This allows the wing to maintain higher lift coefficients in some portions of the span.

1.4 Publications

The following papers have been submitted to publication during the research:

- Serson, D.; Meneghini, J. R.; Sherwin, S. J. Velocity-correction schemes for the incompressible Navier-Stokes equations in general coordinate systems. Journal of Computational Physics, v. 316, p. 243–254, 2016c
- Serson, D.; Meneghini, J. R.; Sherwin, S. J. Direct numerical simulations of the flow around wings with spanwise waviness at a very low Reynolds number. Computers & Fluids (submitted to publication), 2016a
- Serson, D.; Meneghini, J. R.; Sherwin, S. J. Direct numerical simulations of the flow around wings with spanwise waviness. Journal of Fluid Mechanics (submitted to publication), 2016d
- Bolis, A.; Cantwell, C. D.; Moxey, D.; Serson, D.; Sherwin, S. J. An adaptable parallel algorithm for the direct numerical simulation of incompressible turbulent flows using a Fourier spectral/hp element method and MPI virtual topologies. Computer Physics Communications, v. 206, p. 17–25, 2016

In addition, the following conference papers were presented:

- Serson, D.; Meneghini, J. R. Numerical study of wings with wavy leading and trailing edges. In Proceedings of the IUTAM-ABCM 8th Symposium on Laminar Turbulent Transition, 2014
- Serson, D.; Meneghini, J. R.; Sherwin, S. J. Extension of the velocity-correction scheme to general coordinate systems. In Proceedings of the International Conference on Spectral and High Order Methods, 2016b

1.5 Thesis outline

This thesis is divided in four parts, each of them organized in chapters. Appendices are also presented to deepen the discussion of topics which are not essential to the understanding of the text.

Part I contains this introductory chapter and the literature review presented in chapter 2, where the main research works concerning the use of spanwise waviness on wings are briefly presented, discussing the extension and limitations of the knowledge currently available on this topic.

The numerical methods employed in this work are presented in part II. First, chapter 3 briefly introduces spectral/hp methods. This includes only the basic aspects of the method required to understand the topics presented here, with details left for the indicated references. Then, chapter 4 describes the approach employed to include the waviness on the numerical simulations. This consists of a new method proposed to introduce general co-ordinate transformations into the velocity-correction scheme, which is the time-integration method generally employed when solving the incompressible Navier-Stokes equations with spectral/hp discretisations. To finish this part on the numerical methods, chapter 5 discusses a few topics related to practical aspects that were considered when performing direct numerical simulations at Reynolds number leading to transition to a turbulent flow.

The results of the simulations are contained in part III. This consists of three chapters, each of them presenting results at a different Reynolds number. To be specific, chapters 6 to 8 correspond to Re = 1,000, Re = 10,000 and Re = 50,000, in this order.

Finally, part IV summarizes the main conclusions of this work (chapter 9) and suggests lines of research that can be followed from the findings of this thesis (chapter 10).

2. Literature review

This chapter consists of a brief description and discussion of the main references available in the literature which are relevant to the present study. Section 2.1 considers works related to the use of waviness on the wing leading edge, while section 2.2 presents the references related to the use of wavy trailing edges. Whenever possible, the references are presented in chronological order, in order to show the evolution of the knowledge about these topics.

2.1 Use of leading edge waviness

The use of leading edge waviness or protuberances as a possible means of improving the aerodynamic characteristics of wings started receiving attention recently following the work of Fish and Battle (1995), which analyses the morphology of the pectoral flippers of humpback whales (*Megaptera novaeangliae*). These flippers have tubercles in the leading edge region, as shown in figure 2.1. Fish and Battle (1995) suggest that these tubercles might have the purpose of delaying stall, allowing the flipper to maintain high values of lift coefficient even at high angles of attack, providing the whale with its well-known good manoeuvrability.

In order to verify the possibility of obtaining an increase in the lift coefficient with the use of wavy leading edges, Watts and Fish (2001) carried out numerical simulations for the flow around a finite rectangular wing (with aspect ratio close to 2) employing a panel method. Two simulations were performed, one for a wing with a straight leading edge and the other with a wavy leading edge, both at an angle of attack of 10 degrees. Their results show that the modifications increase the lift by 4.8% and reduce the induced drag by 10.9%. However, it should be noted that the panel method they employed might be unable to predict accurately viscous effects like the stall. Also, another limitation of this study is that they only considered one angle of attack, and therefore it is not possible to know if the increase in lift at this particular angle of attack will be reflected in an increase in the maximum lift coefficient.

Miklosovic et al. (2004) present results of wind tunnel experiments for two wing models with a NACA0020 profile and planform based on the humpback whale flipper. However, only one model had protuberances in the leading edge, with the other model acting as a reference. The experiments were performed for a Reynolds number (*Re*) between 5.05×10^5



Figure 2.1: On the left, picture of a humpback whale, extracted from Custodio (2007). On the right, planform of a pectoral flipper, extracted from Fish and Battle (1995).

and 5.20×10^5 . They present lift and drag coefficient curves as a function of the angle of attack, and the waviness clearly postpone the stall, leading to better aerodynamic characteristics at high angles of attack and to a higher value for the maximum lift coefficient. Miklosovic et al. (2007) complement the study of Miklosovic et al. (2004), including results of similar experiments for the previous half-span models (i.e., models with a free extremity), and adding results for wings with a rectangular planform, mounted in a full-span configuration (i.e., with both extremities close to the wind-tunnel walls) to approximate the behaviour of infinite wings. The experiments for the half-span wing were performed for a Reynolds number similar to the previous work (between 534,000 and 631,000), while for the rectangular wing the Reynolds number was lower, between 274,000 and 277,000. Figure 2.2 contains the main results of their work, showing a clear distinction between the two cases, with the use of waviness in the rectangular wing leading to a premature stall and improving the performance only in the post-stall regime.

The work described in Levshin et al. (2006), Johari et al. (2007) and Custodio (2007) presents experiments performed in a water channel with rectangular wing models with a NACA63₄ – 021 profile, for $Re = 1.83 \times 10^5$. They used sinusoidal waviness with amplitudes between 2.5% and 12% of the mean chord, and wavelengths between 25% and 50% of the mean chord. Their results showed a similar behaviour to the full-span model of Miklosovic et al. (2007), with the waviness decreasing the lift and increasing the drag for moderate angles of attack (below the baseline stall angle), and increasing the lift and having negligible effect on the drag in the post-stall regime. They also observed that for this range of the parameters, the waviness amplitude is the parameter with highest effect on the results, with the wavelength and the leading edge radius having secondary importance. Finally, they present visualizations of the flow on the wing surface showing that, for the cases



Figure 2.2: On the left, models studied by Miklosovic et al. (2007). On the right, lift coefficient results. The top figures correspond to the full-span model (two-dimensional analysis), while the bottom figures correspond to the half-span model (three-dimensional analysis). Adapted from Miklosovic et al. (2007).

with waviness, stall occurs by separation behind the valleys of the waviness, with the flow remaining more resistant to separation behind the protuberances.

An attempt to model the effect of wavy leading edges on the aerodynamic characteristics of a wing is presented in van Nierop et al. (2008). In this work, elliptic wings with waviness were studied applying a lifting line theory, assuming the hypothesis that spanwise flow is negligible. They compared their results with experiments from Johari et al. (2007), with the comparison reproduced in figure 2.3. The model was capable of capturing some characteristics of the experimental results, such as the smoothing in the stall characteristics and the fact that separation starts in the valleys. However, in general there is a large difference between the obtained behaviours. Also, it should be noted that the model was applied to an elliptic wing, while the experimental results presented are for a rectangular wing, making comparisons more difficult to be performed.

Some of the experimental results of Miklosovic et al. (2004) were reproduced numerically by Pedro and Kobayashi (2008), using a detached eddy simulation (DES) formulation. The use of numerical simulations allowed the flow field to be described in more detail than in the experiments. For example, the vorticity magnitude contours reproduced in figure 2.4 show the presence of streamwise vortices originated at the protuberances, which is one hypothesis presented to explain the delay in the stall that they may cause.

The references presented above indicate that for a finite wing the leading edge waviness



Figure 2.3: On the left, results obtained by the model of van Nierop et al. (2008). On the right, experimental results from Johari et al. (2007). The symbols represent cases with 4 wavelengths along the span, while the lines represent cases with 8 wavelengths along the span. The parameter η represents the amplitude of the bumps. Extracted from van Nierop et al. (2008).



Figure 2.4: Comparison between instantaneous vorticity magnitude contours for a smooth and a wavy wing. Extracted from Pedro and Kobayashi (2008).



Figure 2.5: Recirculation zones for Re = 800 and $\alpha = 20^{\circ}$. The first case in the left corresponds to a straight wing, while the other cases correspond to different combinations of amplitude and wavelength. Adapted from Favier et al. (2012).

delays the stall, while for an infinite wing (or a full-span model, in the case of experiments) they only lead to improvements in the post-stall regime. However, experiments for a half-span model presented in Stanway (2008) contradict this idea. In this work, experiments were performed at four different values of Reynolds number between 44, 648 and 119, 060. The results have a strong dependence on Re, with an increase in the maximum lift coefficient being observed only for the highest value of Re. The explanation they give for this is that the waviness effect depends on the type of stall that occurs in the wing, with improvements in the aerodynamic characteristics occurring only in the case of leading edge stall.

Another important result is presented in Hansen et al. (2010). In this study, experiments were performed with rectangular wing models (with and without leading edge waviness) with a NACA0021 profile, for $Re = 1.2 \times 10^5$. These models were mounted in the wind tunnel in two different configurations, one with the model close to the walls of the tunnel, and the other where the model has a free extremity. This way, the same model was studied in a full-span configuration and in a half-span configuration, allowing for the determination of how the three-dimensional effects imposed by the wing tip affect the effectiveness of the waviness. Both configurations presented similar results, and therefore it was concluded that the waviness has little effect on the formation of the wing tip vortices. Since three-dimensional effects from the wing tip have little importance, it was suggested that the main factor to explain the different behaviours observed in the literature is the Reynolds number. However, it is important to note that although the three-dimensional effects from wing tip do not play an important role in understanding the impact of the waviness, the effects from using different planforms might still be significant.

Favier et al. (2012) studied this problem numerically considering infinite wings with a NACA0020 profile and Re = 800. They employed an immersed boundary method, carrying out a parametric study for the amplitude and wavelength, for a fixed angle of attack 20 degrees. For this angle of attack, the waviness leads to reductions in lift and drag in all cases. Also, the lift coefficient fluctuations decreased, becoming negligible in some cases. Figure 2.5 reproduce the recirculation regions presented for some of the cases analysed, showing some structures which might be present in the flow around wings with a wavy leading edge.





Skillen et al. (2015) present results of large-eddy simulations (LES) reproducing the smooth wing and one of the wavy wings of Hansen et al. (2010), for an angle of attack $\alpha = 20^{\circ}$. Although the lift coefficient did not match the experiments, the flow is described in detail, illustrating some of the features which might be expected for this flow. For example, they support that an important effect of the waviness is to deflect the flow around the leading edge, inducing a secondary flow. This is illustrated by the streamlines of figure 2.6.

Other works available in the literature study the effect of using wavy leading edges, both numerically (Arai et al., 2010; Weber et al., 2011; Yoon et al., 2011; Malipeddi, 2011; Malipeddi et al., 2012) and experimentally (Arai et al., 2010; Guerreiro and Sousa, 2012; Paula, 2016; Rostamzadeh et al., 2012; Zhang et al., 2013), with results that are consistent with the ones discussed in this section. Despite the large quantity of studies about this topic, many of them treat the problem only in a superficial manner, containing only experimental force measurements, or focusing on a single angle of attack. Therefore, it is still not clear what are the physical mechanisms that lead to the observed behaviours, and under what exact circumstances the waviness improve the aerodynamic performance of the wing.

2.2 Use of trailing edge waviness

The study of trailing edge waviness is generally related to bodies with a blunt trailing edge. The use of wings with blunt trailing edges, discussed in Summers and Page (1950) and Dugan (1952), can lead to an increase in the slope of the lift coefficient curve and in the maximum lift coefficient. However, it also increases the drag, because of the vortex shedding at the trailing edge. To reduce this problem, Summers and Page (1950) proposed the use of a splitter plate to suppress vortex shedding, and this solution was tested experimentally by Bearman (1965) for a semi-elliptic body with Reynolds numbers between 1.4×10^5 and 2.45×10^5 . The results of Bearman (1965) showed that a splitter plate with



Figure 2.7: Blunt trailing edge body studied by Bearman and Tombazis (1993).

length l/h = 1 (where *h* is the trailing edge height) leads to a significant increase in the base pressure coefficient, with $-C_{pb}$ falling approximately by half, what was attributed to an increase in the vortex formation length. For longer splitter plates, with l/h > 2, vortex shedding ceased to occur, with $-C_{pb}$ falling to a slightly lower value than for l/h = 1.

An alternative presented later to reduce the drag caused by the blunt trailing edge consists in changing the three-dimensional geometry of the separation line. The first studies in this direction, presented in Gai and Sharma (1983) and Petrusma and Gai (1989), analyse the effect of segmented (discontinuous) trailing edges. This modification caused an increase in the base pressure, therefore reducing the drag.

Following the idea of changing the geometry of the separation line, Bearman and Tombazis (1993) performed experiments for a semi-elliptic body with sinusoidal waviness in the trailing edge, using the setup shown in figure 2.7. The waviness lead to reductions of up to 35% in the suction at the base of the body. This work was complemented by Tombazis and Bearman (1997), where the modifications in the wake that lead to the drag reductions are discussed. In another study, Bearman and Owen (1998) applied the same concept to flat plates perpendicular to the flow and to rectangular cylinders. For these cases, drag reductions around 30% where observed.

The geometry considered by Bearman and Tombazis (1993) was studied numerically by Cai et al. (2008) using direct numerical simulations for Re = 2500. They also observed a drag reduction, which depends on the wavelength of the trailing edge waviness.

Based on the references discussed above, it can be noted that wavy trailing edges lead to a significant drag reduction in bodies with blunt trailing edges. However, there are no studies in the literature about how the shape of sharp trailing edges alter the aerodynamic performance of a wing. Therefore, this problem will be briefly discussed in chapter 6.

Part II

Numerical methods

3. The spectral/hp element method

In order to study the flow around wings with waviness, the Navier-Stokes equations for an incompressible flow were solved numerically using the spectral/hp element method, which is available through the Nektar++ library. This method combines characteristics of the finite element method (FEM) with characteristics from the spectral methods. In the spectral/hp element method, the domain is first divided in non-overlapping elements, in the same way as in the FEM, so it offers geometric flexibility and allows for local refinement. Then, the solution in each element is approximated by high order polynomials, giving the method an exponential spatial convergence. This chapter presents a brief introduction on this method, focusing on the topics which are more relevant to the present study. Further details about the method can be obtained in Karniadakis and Sherwin (2005), which served as a basis for much of this chapter.

3.1 The Galerkin Method

This section begins the description on how to proceed in order to numerically solve a partial differential equation. This will be done by considering the particular example of the Helmholtz equation

$$\nabla^2 u - \mu u - f = 0, \tag{3.1}$$

where u is the unknown, μ is a constant and f is a known forcing term. This equation was chosen because it appears in several applications, including the solution of the incompressible Navier-Stokes equations, which is the object of the present work.

We want to obtain the solution u in the domain Ω , subject to appropriate conditions on the boundary Γ . However, representing the exact solution would require knowing the values of u on all points of the domain, what in general is not possible from a numerical point of view. Therefore, the first step to solve the equation is approximating u by a linear combination of a finite number of basis functions, leading to

$$u^{\delta} = \sum_{i=1}^{N_{dof}} \hat{u}^i \Phi^i, \tag{3.2}$$

where u^{δ} is the approximate solution, N_{dof} is the number of degrees of freedom in the

approximation, \hat{u}^i are the coefficients of the expansion and Φ^i are the basis functions. It is also convenient to decompose the solution in a part u^D which satisfies the Dirichlet boundaries condition and an homogeneous solution u^H vanishing on those boundaries, in such a way that

$$u^{\delta} = u^{D} + u^{H} = u^{D} + \sum_{i=1}^{N_{dof,H}} \hat{u}^{i} \Phi^{i},$$
(3.3)

where the basis functions are now zero on the boundaries where Dirichlet boundary conditions are imposed.

For a given set of coefficients, it is obvious that u^{δ} will not necessarily satisfy equation (3.1). This only happens if u^{δ} is equal to the exact solution, and thus the equation can only be satisfied by u^{δ} if the exact solution belongs to the function space spanned by the basis functions, what is not likely to happen in practical scenarios. Therefore, in order to obtain an approximate solution, we must choose a set of conditions to be imposed on u^{δ} . Denoting the residual from equation (3.1) by $R(u^{\delta}) = \nabla^2 u^{\delta} - \mu u^{\delta} - f$, a widely used family of conditions is given by

$$\int_{\Omega} R(u^{\delta}) v^{j} d\Omega = 0, \qquad (3.4)$$

where v^j are called the test functions. This is called the weighted residuals method, and particular choices of the test functions lead to different numerical methods, like the Finite Volume Method and the least-square method. In the context of the spectral/hp element method, the approach used is the Galerkin method, which consists in taking the test functions to be the same as the basis functions Φ^i . Therefore, the approximate form of equation (3.1) is

$$\int_{\Omega} \left(\nabla^2(u^{\delta}) - \mu u^{\delta} - f \right) \Phi^j d\Omega = 0, \ 1 \le j \le N_{dof,H}.$$
(3.5)

Since the index j must be considered in the whole range from 1 to $N_{dof,H}$, this represents a set of $N_{dof,H}$ equations for $N_{dof,H}$ unknowns, which leads to a determined system.

Equation (3.5) can be further simplified by applying the divergence theorem to the first term in the integral, leading to

$$\int_{\Omega} \nabla u^{\delta} \cdot \nabla \Phi^{j} + \mu u^{\delta} \Phi^{j} d\Omega = \int_{\Omega} -f \Phi^{j} d\Omega + \int_{\Gamma} \Phi^{j} \nabla u^{\delta} \cdot \mathbf{n} dS,$$
(3.6)

where the last term vanishes on the boundaries with Dirichlet conditions, and therefore can be obtained from the Neumann boundary conditions. Substituting the definition of u^{δ} from equation (3.3), this can be represented in matrix form as

$$(L+\mu M)\hat{u}=F,\tag{3.7}$$

where

$$L_{ij} = \int_{\Omega} \nabla \Phi^{i} \cdot \nabla \Phi^{j} d\Omega,$$

$$M_{ij} = \int_{\Omega} \Phi^{i} \Phi^{j} d\Omega,$$
(3.8)

are the weak Laplacian matrix and the mass matrix, respectively. The right-hand side of the equation is:

$$F_j = -\int_{\Omega} \left(f \Phi^j + \nabla u^D \cdot \nabla \Phi^j + \mu u^D \Phi^j \right) \, d\Omega + \int_{\Gamma} \Phi^j \nabla u^\delta \cdot \mathbf{n} \, dS. \tag{3.9}$$

The previous formulation contains a general framework for finding an approximate solution for a partial differential equation. However, it is still necessary to choose a particular set of basis functions, and to define how to perform the several operations required for obtaining the matrices presented above. These topics will be briefly discussed in the next sections.

3.2 Local formulation

In the spectral/hp element methods, the computational domain is divided in several nonoverlapping elements, with the solution being represented by a high-order expansion (usually polynomial) inside each element. Most operations are performed in the element level, with an assembly operation being used to obtain a global solution from these local expansions. This section first presents some of the expansion bases employed to represent the solution in a single element. Then, the several local operations required to solve a partial differential equation are described. The focus will be on quadrilateral elements, since the meshes used in the present work were based on them; however, it is also possible to consider triangles or three-dimensional elements, as discussed in Karniadakis and Sherwin (2005). The process of moving from a local to a global expansion will be the topic of the next section.

3.2.1 Expansion basis

A good starting point for introducing the expansion basis consists in analysing the onedimensional case, since it acts as a building block for higher-dimension expansions. First, we consider a standard segment represented by $\xi \in [-1, 1]$. After having defined the expansion and operations in this standard region, the treatment can be extended to a general region by a mapping procedure, as discussed in section 3.2.3. Now that the standard region was defined, we must choose a convenient representation of the solution inside this region. This is done by approximating the solution by a polynomial of degree P, which can be seen as a linear combination of P+1 basis functions. There are several possible choices for the basis of this polynomial space. Two well-known examples are

$$h_p(\xi) = \frac{\prod_{q=0, q \neq p}^{P}(\xi - \xi_q)}{\prod_{q=0, q \neq p}^{P}(\xi_p - \xi_q)}, \ p = 0, ..., P,$$

$$\phi_p(\xi) = P_p^{0,0}(\xi) = L_p(\xi), \ p = 0, ..., P,$$
(3.10)

where ξ_q are a set of nodes and $P_p^{\alpha,\beta}$ represent the Jacobi polynomials. The first expansion consists of the Lagrange polynomials through the points ξ_q . It is called a nodal expansion, since each of its functions assumes the value 1 on one of the nodes and is zero on all other nodes. The second expansion is formed by the Legendre polynomials, with its components being orthogonal. This is a modal (or hierarchical) basis, since the mode corresponding to a given p is independent of the polynomial order P, and therefore the expansion with order P+1 contains all modes from the expansion of order P.

Although the orthogonality of the Legendre polynomials would seem to lead to attractive properties for this representation, in the context of the spectral element methods employing a continuous Galerkin formulation it is usually not a reasonable choice. This is because when we consider a multi-element expansion, imposing continuity across elements using this basis would lead to a coupling involving all modes. To circumvent this problem, this basis can be modified in order to obtain an expansion with the property usually called boundary-interior decomposition, which happens when only a few modes are non-zero on the element boundary. These modes are called boundary modes, while the others are the interior modes. By using this construction, continuity between elements can be imposed considering only the boundary modes, reducing the coupling. One possibility for a modal expansion with boundary-interior decomposition is

$$\psi_{p}(\xi) = \begin{cases} \frac{1-\xi}{2}, & p=0\\ \frac{1+\xi}{2}, & p=1\\ \left(\frac{1-\xi}{2}\right) \left(\frac{1+\xi}{2}\right) P_{p-2}^{1,1}(\xi), & 2 \le p \le P. \end{cases}$$
(3.11)

This is a hierarchical basis where on each vertex there is only one non-zero mode. The basis formed by Lagrange polynomials including nodes on both ends of the interval also satisfies the boundary-interior decomposition. Because of that, when continuity of the solution is required, the most widely used expansion basis employed in this method are the nodal basis formed by the Lagrange polynomials and the modified modal basis of equation (3.11). The latter was the one employed in the simulations of the present work. The extension of the one-dimensional basis to quadrilateral elements is straightforward. First, we define the standard quadrilateral region as

$$Q_{st} = \{(\xi_1, \xi_2) | -1 \le \xi_1, \xi_2 \le 1\}.$$
(3.12)

Then, the polynomial basis in this region is obtained by a tensor product between onedimensional basis in each direction:

$$\Phi_{pq}(\xi_1,\xi_2) = \Phi_p(\xi_1)\Phi_q(\xi_2), \ 0 \le p,q \le P,$$
(3.13)

where $\Phi_p(\xi_1)$ and $\Phi_q(\xi_2)$ can be any of the polynomial basis described previously. Note that if the one-dimensional basis has the boundary-interior decomposition property, then by using this definition the two-dimensional basis will also have it.

When in a three-dimensional domain all xy planes are identical and the boundary conditions are periodic in z, we say that the third direction is homogeneous. In this situation, usually called 2.5D or quasi-3D case, it is convenient to represent the homogeneous direction by a Fourier expansion, following the procedure proposed by Karniadakis (1990). By doing this, the three-dimensional expansion basis becomes

$$\Phi_{pqr}(\xi_1,\xi_2,\xi_3) = \Phi_{pq}(\xi_1,\xi_2)e^{ir\beta\xi_3}.$$
(3.14)

The main advantages of using the Fourier expansion in the third direction become clear when we consider the particular form the differential operators assume in this case:

$$\nabla \Phi_{pqr}(\xi_1, \xi_2, \xi_3) = \begin{bmatrix} \frac{\partial}{\partial \xi_1} \\ \frac{\partial}{\partial \xi_2} \\ ir\beta \end{bmatrix} \Phi_{pqr}(\xi_1, \xi_2, \xi_3),$$

$$\nabla^2 \Phi_{pqr}(\xi_1, \xi_2, \xi_3) = \left(\frac{\partial^2}{\partial \xi_1^2} + \frac{\partial^2}{\partial \xi_2^2} - (r\beta)^2\right) \Phi_{pqr}(\xi_1, \xi_2, \xi_3).$$
(3.15)

The first thing we notice from these equations is that derivatives in the homogeneous direction can be easily calculated. Also, the form of these operators, combined with the orthogonality of the Fourier modes, imply that application of the Galerkin method to linear problems such as the Helmholtz equation result in no coupling between the Fourier modes. Therefore, a three-dimensional problem can be replaced by a set of two-dimensional problems, making paralellization of the method in this situation straightforward.

3.2.2 Operations in standard region

Having defined how the solution can be represented in the standard element, the next step is to determine how to perform operations in this region. The implementation of the Galerkin method requires fields to be integrated and differentiated in the domain. Also, it is important to be able to transform the solution from physical values to the coefficient representation, and vice versa. The operations will be presented only for the quadrilateral expansions, since the operations in the homogeneous direction can be obtained directly from the properties of the Fourier series.

Integration

From equation (3.8), we note that obtaining the matrices required to apply the Galerkin Method involves calculating integrals in the domain, and therefore this is an essential operation for the method. Since we are using polynomial functions, it is natural to evaluate integrals numerically using Gaussian quadrature rules, which lead to exact results for polynomials of degree lower than a certain limit. These rules consist in approximating one-dimensional integrals by

$$\int_{-1}^{1} u(\xi) d\xi \approx \sum_{i=0}^{Q-1} w_i u(\xi_i),$$
(3.16)

where ξ_i represent a set of Q quadrature points and w_i are weights for each of the points. Taking into account the shape of the standard quadrilateral element, it is natural to extend this definition to it using

$$\int_{Q_{st}} u(\xi_1, \xi_2) \, d\Omega_{st} = \int_{-1}^1 \left[\int_{-1}^1 u(\xi_1, \xi_2) d\xi_1 \right] d\xi_2$$

$$\approx \int_{-1}^1 \left[\sum_{i=0}^{Q_1 - 1} w_i u(\xi_{1,i}, \xi_2) \right] d\xi_2 \qquad (3.17)$$

$$\approx \sum_{j=0}^{Q_2 - 1} \sum_{i=0}^{Q_1 - 1} w_i w_j u(\xi_{1,i}, \xi_{2,j}).$$

The simulations of the present work employed the Gauss-Lobatto-Legendre quadrature, which includes both end points as quadrature points. This method is exact for polynomials of order up to 2Q-3, and therefore when the polynomial order is the same in both directions it is usual to employ P+2 quadrature points in each direction in order to obtain exact results for polynomials of order 2P, like the integrands of the matrices of equation (3.8).

Differentiation

Differentiation in the standard element can be performed by expressing the expansion in a nodal representation, using the Lagrange polynomials. This leads to

$$u(\xi_1,\xi_2) = \sum_{i=0}^{Q_1-1} \sum_{j=0}^{Q_2-1} u(\xi_{1,i},\xi_{2,j}) h_i(\xi_1) h_j(\xi_2).$$
(3.18)

Calculating the partial derivative with respect to ξ_1 , for example, results in

$$\frac{\partial u(\xi_1,\xi_2)}{\partial \xi_1} = \sum_{i=0}^{Q_1-1} \sum_{j=0}^{Q_2-1} u(\xi_{1,i},\xi_{2,j}) \frac{dh_i(\xi_1)}{d\xi_1} h_j(\xi_2).$$
(3.19)

If we are only interested in the values at the quadrature points, as is normally the case, this can be represented as

$$\frac{\partial u(\xi_{1,i},\xi_{2,j})}{\partial \xi_1} = \sum_{p=0}^{Q_1-1} u(\xi_{1,p},\xi_{2,j}) \left. \frac{dh_p(\xi_1)}{d\xi_1} \right|_{\xi_1,i},$$
(3.20)

where the fact that $h_j(\xi_{2,q}) = \delta_{jq}$ was employed to eliminate this term. Therefore, numerically evaluating the derivatives of a function was reduced to obtaining the derivatives of the Lagrangian polynomials, which can be calculated from closed expressions.

Backward and forward transformations

The last set of operations described here deal with how to transform the solution between the physical values and the expansion coefficients. Obtaining the physical values from the coefficients is called the backward transformation, and is a straightforward operation considering the definition from equation (3.2). In a two-dimensional space with a tensorial basis, this leads to

$$u(\xi_1,\xi_2) = \sum_{i=0}^{P_1} \sum_{j=0}^{P_2} \hat{u}_{i,j} \Phi_i(\xi_1) \Phi_j(\xi_2).$$
(3.21)

Also, this operation (and the differentiation) can be made more efficient by using the sumfactorization technique, as explained in Karniadakis and Sherwin (2005).

The inverse operation (from physical space to coefficients) is called the forward transformation, and is not so simple. If the number of quadrature points was the same as the number of modes, then this operation would correspond to a polynomial interpolation. However, since there are usually P + 2 quadrature points in each direction, the problem is overspecified, and a projection operation must be defined. One possible way of performing this projection consists in using the Galerkin method, leading to

$$\int_{\Omega} \Phi_{pq} \hat{u}_{ij} \Phi_{ij} \, d\Omega = \int_{\Omega} \Phi_{pq} u(\xi_1, \xi_2) \, d\Omega.$$
(3.22)

From this equation, we note that the forward transformation can be performed by first calculating the inner product of the physical values with respect to the basis functions, and then multiplying by the inverse of the mass matrix.

3.2.3 Local operations

Now that the expansion and operations are defined in a standard region, we have to determine a way to extend the analysis to a general element. The first step is to determine a relation between the local coordinates x_i and the standard element coordinates ξ_i . For a quadrilateral with straight edges connecting the vertices A, B, C and D, the transformation between the two sets of coordinates is

$$x_{i} = \frac{1-\xi_{1}}{2} \frac{1-\xi_{2}}{2} x_{i}^{A} + \frac{1+\xi_{1}}{2} \frac{1-\xi_{2}}{2} x_{i}^{B} + \frac{1+\xi_{1}}{2} \frac{1+\xi_{2}}{2} x_{i}^{C} + \frac{1-\xi_{1}}{2} \frac{1+\xi_{2}}{2} x_{i}^{D}.$$
(3.23)

We note from this expression that x_i is represented as a combination of the vertex modes from the modal basis with boundary-interior decomposition. It is natural to extend this approach to curved side elements by also including the edge modes of the expansion in the transformation. Therefore, in general the coordinates x_i will be represented by an expansion in the standard element using the modified modal basis

$$x_{i} = \sum_{j=1}^{N_{modes}} \hat{x}_{i}^{j} \Phi^{j}.$$
 (3.24)

From this expression, it is clear that derivatives of the form $\frac{\partial x_i}{\partial \xi_j}$ can be obtained using the approach presented in the previous section.

The integral of a field u in a quadrilateral element Q can be obtained based on the integration in the standard element by

$$\int_{Q} u(x_1, x_2) dx_1 dx_2 = \int_{Q_{st}} u(\xi_1, \xi_2) |J| d\xi_1 d\xi_2,$$
(3.25)

where |J| is the Jacobian of the transformation, defined as

$$|J| = \frac{\partial x_1}{\partial \xi_1} \frac{\partial x_2}{\partial \xi_2} - \frac{\partial x_1}{\partial \xi_2} \frac{\partial x_2}{\partial \xi_1}.$$
(3.26)

Differentiation in the local element can be performed by using the chain rule, leading to

$$\frac{\partial}{\partial x_i} = \frac{\partial \xi_1}{\partial x_i} \frac{\partial}{\partial \xi_1} + \frac{\partial \xi_2}{\partial x_i} \frac{\partial}{\partial \xi_2}.$$
(3.27)

Although the terms $\frac{\partial}{\partial \xi_j}$ can be readily obtained in the standard region, the derivatives $\frac{\partial \xi_j}{\partial x_i}$ cannot be obtained directly. Therefore, it is necessary to obtain an expression for these terms. Karniadakis and Sherwin (2005) show that they can be rewritten in terms of $\frac{\partial x_i}{\partial \xi_j}$, with the result being

$$\frac{\partial \xi_1}{\partial x_1} = \frac{1}{J} \frac{\partial x_2}{\partial \xi_2}, \quad \frac{\partial \xi_2}{\partial x_2} = \frac{1}{J} \frac{\partial x_1}{\partial \xi_1}, \quad \frac{\partial \xi_1}{\partial x_2} = -\frac{1}{J} \frac{\partial x_2}{\partial \xi_1}, \quad \frac{\partial \xi_2}{\partial x_1} = -\frac{1}{J} \frac{\partial x_1}{\partial \xi_2}. \tag{3.28}$$

Using these relations, derivatives with respect to x_i can be obtained by knowing the function $x_i(\xi_1, \xi_2)$ and by using differentiation in the standard element.

Finally, the backward transformations can be performed directly in the standard region, and the forward transformation differs from the standard region operation only through the integration procedure already discussed.

3.3 Global assembly and global operations

The previous section described how several operations can be performed in a single element. Now, the process of combining the contributions from all elements in order to obtain a global representation of our problem will be briefly described.

The nature of the equations of interest to the present work suggests that the solution should be C^0 continuous. Considering the boundary-interior decomposition of the local expansions, this can be achieved by matching boundary modes from neighbour elements. Therefore, the combination of these matching local boundary degrees of freedom will correspond to a single global degree of freedom. If we define the vector containing the coefficients of all global degrees of freedom as \hat{u}_g and the concatenation of local coefficients in all elements as \hat{u}_l , this correspondence between global and local modes can be represented in matrix form by a connectivity map \mathcal{A} , and therefore

$$\hat{u}_l = \mathcal{A}\hat{u}_g,\tag{3.29}$$

where the non-zero entries of A are either 1 or -1, with the sign accounting for possible differences in orientation between edges of adjacent elements. In practice, the operation of equation (3.29) can be done more efficiently without assembling the matrix A, since this matrix is very sparse. However, the matrix represents a good illustration of the exact behaviour of the operation.

In the Galerkin method, the entries in the global system matrices correspond to integrals

involving the expansion modes. Therefore, contributions from different local degrees of freedom forming the same global degree of freedom need to be summed when moving from results obtained using the element matrices to the global result, in what is called the global assembly operation. It is easy to see that this will correspond to pre-multiplying the block diagonal concatenation of the elemental matrices by \mathcal{A}^T . Also, the local matrices can only be applied to a local vector \hat{u}_l , and therefore we first have to use equation (3.29) to obtain \hat{u}_l . Following this whole process, a global matrix $\mathbf{M}_{\mathbf{g}}$ can be obtained by

$$\mathbf{M}_{\mathbf{g}} = \mathcal{A}^T \underline{\mathbf{M}}_{\underline{\mathbf{e}}} \mathcal{A}, \tag{3.30}$$

where $\mathbf{M}_{\mathbf{e}}$ is the block diagonal concatenation of the elemental matrices.

The concepts presented thus far illustrate how a simple equation, like the Helmholtz equation of section 3.1 can be solved. Most operations are performed within each elements, and then the operation represented in equation (3.30) is used to obtain a global system, which has to be solved in order to obtain the final solution. The next section extends this discussion, by showing how this method can be used in the particular case of the Navier-Stokes equations.

3.4 Time integration of the Navier-Stokes equations

Since the equations we consider are time-dependent, after applying the spectral/hp spatial discretization it is still necessary to define a time integration scheme to solve the problem. Other points that need to be addressed are the fact that the Navier-Stokes equations are non-linear, and the coupling between pressure and velocity. The approach used here is the splitting scheme of Karniadakis et al. (1991) with a stiffly stable time discretization, which is also called velocity-correction scheme.

We are interested in solving the unsteady Navier-Stokes equations for an incompressible flow, which assuming a unity density can be written as

$$\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u} \cdot \nabla)\mathbf{u} - \nabla p + \nu \nabla^2 \mathbf{u},$$

$$\nabla \cdot \mathbf{u} = 0,$$

(3.31)

where \mathbf{u} is the velocity, p is the pressure, and ν is the kinematic viscosity. Given a reference length L and a reference velocity U, the Reynolds number is defined as $Re = \frac{LU}{\nu}$. Equation (3.31) needs to be solved in the domain Ω subject to a initial condition \mathbf{u}^0 and to the boundary conditions

$$\mathbf{u} = \mathbf{u}_{\mathfrak{D}} \qquad on \ \Gamma_{\mathfrak{D}},$$

$$\frac{\partial \mathbf{u}}{\partial n} = \mathbf{u}_{\mathfrak{N}} \qquad on \ \Gamma_{\mathfrak{N}},$$

(3.32)
Coefficient	1^{st} order	2^{nd} order	3^{rd} order
γ_0	1	3/2	11/6
$lpha_0$	1	2	3
α_1	0	-1/2	-3/2
$lpha_2$	0	0	1/3
β_0	1	2	3
β_1	0	-1	-3
β_2	0	0	1

Table 3.1: Coefficients for stiffly-scheme time integration.

where $\Gamma_{\mathfrak{D}}$ is the part of the boundary where Dirichlet boundary conditions are imposed on the velocity, while on $\Gamma_{\mathfrak{N}}$ Neumann boundary conditions are imposed. It is also convenient to denote the convective terms as

$$\mathbf{N}(\mathbf{u}) = -(\mathbf{u} \cdot \nabla)\mathbf{u} \tag{3.33}$$

and the viscous terms as

$$\nu \mathbf{L}(\mathbf{u}) = \nu \nabla^2 \mathbf{u},\tag{3.34}$$

so that the complete problem we want to solve is:

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbf{N}(\mathbf{u}) - \nabla p + \nu \mathbf{L}(\mathbf{u}) \quad in \ \Omega,$$

$$\nabla \cdot \mathbf{u} = 0 \quad in \ \Omega,$$

$$\mathbf{u} = \mathbf{u}_{\mathfrak{D}} \quad on \ \Gamma_{\mathfrak{D}},$$

$$\frac{\partial \mathbf{u}}{\partial \mathbf{n}} = \mathbf{u}_{\mathfrak{N}} \quad on \ \Gamma_{\mathfrak{N}}.$$
(3.35)

Following the time-integration scheme of Karniadakis et al. (1991), the first step to solving this problem is to discretise the momentum equation in time by approximating the time derivative by a backward differentiation formula (BDF) and by representing the convective term explicitly using a polynomial extrapolation from previous time-steps. Using these approximations, the momentum equation at time-step n + 1 is reduced to

$$\frac{\gamma_0 \mathbf{u}^{n+1} - \sum_{q=0}^{J_i - 1} \alpha_q \mathbf{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J_e - 1} \beta_q \mathbf{N}(\mathbf{u}^{n-q}) - \nabla p^{n+1} + \nu \mathbf{L}(\mathbf{u}^{n+1}),$$
(3.36)

where J_e and J_i are the integration orders of the explicit and implicit terms, respectively, and the coefficients α , β and γ for orders up to 3 are presented in table 3.1. It should be noted that this choice of time discretisation leads to an explicit treatment of the convective terms, which is desired since this term is non-linear, while treating the viscous terms implicitly, avoiding numerical stability problems that would require the use of a very low time step. In order to simplify the notation of equation (3.36), we introduce the following definition of the summations:

$$\sum_{q=0}^{J_i-1} \alpha_q \mathbf{u}^{n-q} = \mathbf{u}^+, \tag{3.37}$$

$$\sum_{q=0}^{J_e-1} \beta_q \mathbf{N}(\mathbf{u}^{n-q}) = \mathbf{N}^*,$$
(3.38)

and therefore equation (3.36) takes the form

$$\frac{\gamma_0 \mathbf{u}^{n+1} - \mathbf{u}^+}{\Delta t} = \mathbf{N}^* - \nabla p^{n+1} + \nu \mathbf{L}(\mathbf{u}^{n+1})$$
(3.39)

Note that the definition from equation (3.38) applies not only for N, but to the extrapolation of any operator.

Karniadakis et al. (1991) propose solving equation (3.39) by splitting the right-hand-side of the equation using the following three steps:

$$\frac{\hat{\mathbf{u}} - \mathbf{u}^{+}}{\Delta t} = \mathbf{N}^{*},$$

$$\frac{\hat{\mathbf{u}} - \hat{\mathbf{u}}}{\Delta t} = -\nabla p^{n+1},$$

$$\frac{\gamma_{0} \mathbf{u}^{n+1} - \hat{\mathbf{u}}}{\Delta t} = \nu \mathbf{L}(\mathbf{u}^{n+1}),$$
(3.40)

where $\hat{\mathbf{u}}$ and $\hat{\mathbf{u}}$ are intermediate fields. After applying the spatial discretisation, the first step would lead to a projection problem, which can be costly since it requires the inversion of the global mass matrix. However, in practice the analytical expression for $\hat{\mathbf{u}}$ from the first step is substituted directly in the second equation, and therefore the first step only involves the direct calculation of $\hat{\mathbf{u}}$ in the physical space, without performing the projection to the solution space. In the second step, the pressure is calculated by taking the divergence of the equation and imposing $\nabla \cdot \hat{\mathbf{u}} = 0$, resulting in

$$\nabla^2 p^{n+1} = \nabla \cdot \left(\frac{\hat{\mathbf{u}}}{\Delta t}\right),\tag{3.41}$$

which is a Poisson equation, requiring appropriate pressure boundary conditions to be solved. The boundary condition proposed by Karniadakis et al. (1991) is

$$\frac{\partial p^{n+1}}{\partial \mathbf{n}} = \mathbf{n} \cdot \left[\mathbf{N}^* - \nu (\nabla \times \nabla \times \mathbf{u})^* \right].$$
(3.42)

Finally, the third step is a Helmholtz equation for the velocity, which can be solved using the formulation from section 3.1.

Although the method described above corresponds precisely to what was employed

in the present work, the way it was obtained is not very enlightening, making it difficult to generalize it when the equations are modified, as will be done in chapter 4. For example, the pressure boundary condition is obtained from the original momentum equation instead of appearing naturally as part of the scheme, and therefore this approach does not show why this boundary condition and the corresponding equation for the pressure are consistent. In order to have a more general approach to obtain the appropriate forms of the equations, the splitting method will be derived again following the framework of Guermond and Shen (2003). The idea behind this method is to first use a modified form of equation (3.39) with the viscous terms treated explicitly, in order to calculate the pressure and obtain an approximation of u satisfying the incompressibility condition. The velocity is then corrected by treating the viscous terms implicitly, hence the name velocity-correction scheme given by Guermond and Shen (2003). Based on this concept, Guermond and Shen (2003) show that the advance in time can be performed by solving the following two steps:

$$\begin{cases} \frac{\gamma_{0}\bar{\mathbf{u}}^{n+1}-\mathbf{u}^{+}}{\Delta t}+\nabla p^{n+1}+\nu(\nabla\times\nabla\times\mathbf{u})^{*}-\mathbf{N}^{*}=0 & in\ \Omega\\ \nabla\cdot\bar{\mathbf{u}}^{n+1}=0 & in\ \Omega\\ \bar{\mathbf{u}}^{n+1}\cdot\mathbf{n}=\mathbf{u}_{\mathfrak{D}}\cdot\mathbf{n} & on\ \Gamma_{\mathfrak{D}}, \end{cases}$$
(3.43)

$$\frac{\gamma_{0}(\mathbf{u}^{n+1}-\bar{\mathbf{u}}^{n+1})}{\Delta t} - \nu \mathbf{L}(\mathbf{u}^{n+1}) - \nu (\nabla \times \nabla \times \mathbf{u})^{*} = 0 \qquad in \ \Omega \\
\mathbf{u}^{n+1} = \mathbf{u}_{\mathfrak{D}} \qquad on \ \Gamma_{\mathfrak{D}} \qquad (3.44) \\
\frac{\partial \mathbf{u}^{n+1}}{\partial \mathbf{n}} = \mathbf{u}_{\mathfrak{N}} \qquad on \ \Gamma_{\mathfrak{N}},$$

where the term $\nabla \times \nabla \times \mathbf{u}$ comes from the decomposition $\nabla^2 \mathbf{u} = \nabla(\nabla \cdot \mathbf{u}) - \nabla \times \nabla \times \mathbf{u}$, and therefore is a valid form for the viscous terms in the continuous case, since $\nabla \cdot \mathbf{u} = 0$. The reason to use this term instead of the Laplacian is to avoid imposing an artificial pressure boundary condition, as was noted by both Karniadakis et al. (1991) and Guermond and Shen (2003).

In the first step (equation (3.43)), we want to solve for the pressure field. If we take the divergence of the equation and then integrate the Laplacian of the pressure by parts to obtain the weak Laplacian matrix (like was done in the second step of equation (3.40)), we will once again need an unknown pressure boundary condition. Instead of that, it is more convenient to dot the equation with gradients of the test functions $\nabla \Phi$ and integrate in the domain, resulting in

$$\int_{\Omega} \nabla p^{n+1} \cdot \nabla \Phi \, d\Omega = \int_{\Omega} \left[\frac{\mathbf{u}^{+} - \gamma_0 \bar{\mathbf{u}}^{n+1}}{\Delta t} + \mathbf{N}^* - \nu (\nabla \times \nabla \times \mathbf{u})^* \right] \cdot \nabla \Phi \, d\Omega.$$
(3.45)

Now, the right-hand-side of the equation can be simplified by using the following form of the divergence theorem:

$$\int_{\Omega} \mathbf{F} \cdot \nabla \Phi \, d\Omega = -\int_{\Omega} \Phi \nabla \cdot \mathbf{F} \, d\Omega + \int_{\Gamma} \Phi \mathbf{F} \cdot \mathbf{n} \, dS.$$
(3.46)

After applying this relation, the equation can be further simplified because $\nabla \cdot \bar{\mathbf{u}}^{n+1} = 0$, and also because the divergence of the curl is zero. Furthermore, based on equation (3.40), it is convenient to define $\hat{\mathbf{u}} = \mathbf{u}^+ + \Delta t \mathbf{N}^*$, leading to

$$\int_{\Omega} \nabla p^{n+1} \cdot \nabla \Phi \, d\Omega = \int_{\Omega} \Phi \nabla \cdot \left(-\frac{\hat{\mathbf{u}}}{\Delta t} \right) \, d\Omega + \int_{\Gamma} \Phi \left[\frac{\hat{\mathbf{u}} - \gamma_0 \bar{\mathbf{u}}^{n+1}}{\Delta t} - \nu (\nabla \times \nabla \times \mathbf{u})^* \right] \cdot \mathbf{n} \, dS,$$
(3.47)

which is equivalent to the weak form of equations (3.41) and (3.42) with the exception of the acceleration term $\frac{\mathbf{u}^{+}-\gamma_{0}\bar{\mathbf{u}}^{n+1}}{\Delta t}$ that is only present when we have time-dependent boundary conditions. This acceleration at timestep n + 1 can be calculated directly when $\bar{\mathbf{u}}^{n+1}$ is known (from a Dirichlet boundary condition), or it can be obtained by extrapolating the acceleration at previous timesteps, as was done with the convective term.

The velocity-correction in the second step (equation (3.44)) can be performed by substituting $\mathbf{\bar{u}}^{n+1}$ from equation (3.43) and defining $\mathbf{\hat{u}} = \mathbf{\hat{u}} - \nabla p^{n+1} \Delta t$, with the result being

$$\frac{\gamma_0 \mathbf{u}^{n+1} - \hat{\mathbf{u}}}{\Delta t} = \nu \mathbf{L}(\mathbf{u}^{n+1}), \qquad (3.48)$$

which is the same as the third step in equation (3.40).

Therefore, equations (3.43) and (3.44) lead to the same splitting-scheme of equation (3.40), but naturally providing the consistent equations with appropriate boundary conditions. This framework will be employed in chapter 4, when the velocity-correction scheme will be generalized to situations when a coordinate system transformation is employed.

4. Treatment of the wavy geometries

The procedure used in the present work to consider the waviness of the wing consists in performing a change of coordinate systems from the Cartesian system (x, y, z) to a transformed system $(\bar{x}, \bar{y}, \bar{z})$. The goal of this transformation is that the waviness of the wing should be removed in the system $(\bar{x}, \bar{y}, \bar{z})$. This approach offers a few advantages. First, wings with waviness of different shapes can be studied using the same mesh simply by changing the function representing the transformation. The second advantage is that, as will be discussed in this chapter, this approach can be easily implemented in existing solvers. Finally, for the infinite wings considered here, the \bar{z} direction becomes homogeneous, and therefore can be discretized using a Fourier series, as explained in section 3.2.1. This leads to a very efficient solver with good parallelization properties, compensating for the time consumed calculating the mapping terms. We are particularly interested in transformations of the form:

$$\begin{cases} x = f(\bar{x}, \bar{z}) \\ y = \bar{y} \\ z = \bar{z}, \end{cases}$$
(4.1)

where by adjusting the function f we can obtain for example a sinusoidal waviness along the span, modulated in the chordwise x direction. An example of how a particular geometry is represented in the two coordinate systems is shown in figure 4.1.

In order to study the flow using the transformed domain represented by the $(\bar{x}, \bar{y}, \bar{z})$ coordinate system, it is necessary to obtain the appropriate form of the Navier-Stokes equations which applies in it. This can be done using tensor calculus, as described in appendix A, with the result being equation (A.21), which can be rewritten as

$$\frac{\partial \mathbf{u}}{\partial t} = \bar{\mathbf{N}}(\mathbf{u}) - \bar{\mathbf{G}}(p) + \nu \bar{\mathbf{L}}(\mathbf{u}),$$

$$D(\mathbf{u}) = 0,$$
(4.2)



Figure 4.1: Geometry of a wing and part of the domain in the different coordinate systems considered.

where:

$$\begin{split} \bar{\mathbf{N}}(\mathbf{u}) &= -\bar{u}^{j}\bar{u}_{,j}^{i} + \bar{V}^{j}\bar{u}_{,j}^{i} - \bar{u}^{j}\bar{V}_{,j}^{i} \\ \bar{\mathbf{G}}(p) &= g^{ij}p_{,j} \\ \nu \bar{\mathbf{L}}(\mathbf{u}) &= \nu g^{jk}\bar{u}_{,jk}^{i} \\ D(\mathbf{u}) &= \frac{1}{J}\nabla \cdot (J\bar{u}^{i}), \end{split}$$

$$(4.3)$$

with g^{ij} representing the inverse of the metric tensor, \bar{u}^i the components of the vector \mathbf{u} in the transformed coordinate system, J the Jacobian of the transformation to the Cartesian system, and a subscript after a comma the covariant derivative. The term \bar{V}^j represents the velocity of the coordinate system, and therefore is only relevant for time-dependent transformations. Although in the present work the transformations of interest are independent of time, the terms containing \bar{V}^j are included to present the formulation in a general context. The particular form of these components for the transformation of equation (4.1) is given in section A.9. It is important to note that although the transformation assumed this particular form throughout this work, the previous equation and the methods presented in this chapter are valid to general transformations, unless otherwise stated. Also, to simplify the notation, throughout this chapter the ∇ operator corresponds to the usual Cartesian operation representing the partial derivatives.

After obtaining the appropriate form of the Navier-Stokes equations, we can consider the domain represented by the $(\bar{x}, \bar{y}, \bar{z})$ coordinates and solve equation (4.2) in it, and the result should be equivalent to solving the Navier-Stokes equations in the original system. However, in order to solve equation (4.2), the time-integration scheme needs to be properly modified to this equation. Sections 4.1 and 4.2 propose two different approaches to achieve this, based on the framework presented in section 3.4. Then, section 4.3 discusses some aspects of post-processing the solution. Section 4.4 presents results of some tests performed to validate the proposed methods, and finally section 4.5 discusses some important points about these methods.

4.1 Explicit treatment of the mapping

In this section a method to solve equation (4.2) treating all the mapping terms explicitly is proposed. First, equation (4.2) is restated as

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbf{N}(\mathbf{u}) - \frac{\nabla p}{J} + \nu \mathbf{L}(\mathbf{u}) + \mathbf{A}(\mathbf{u}, p),$$

$$D(\mathbf{u}) = 0,$$
(4.4)

where $N({\bf u})$ and $L({\bf u})$ are the usual convective and viscous operators from equations (3.33) and (3.34), and

$$\mathbf{A}(\mathbf{u},p) = \left[\bar{\mathbf{N}}(\mathbf{u}) - \mathbf{N}(\mathbf{u})\right] + \left[-\bar{\mathbf{G}}(p) + \frac{\nabla p}{J}\right] + \nu \left[\bar{\mathbf{L}}(\mathbf{u}) - \mathbf{L}(\mathbf{u})\right]$$
(4.5)

is a forcing term that imposes the coordinate transformation. Although it is not explicit in the equations, A also depends on the particular form of the mapping employed. Then, equation (3.39) is modified accordingly, with the forcing term being treated explicitly using an extrapolation from previous time steps in the same manner as the convective term, leading to

$$\frac{\gamma_0 \mathbf{u}^{n+1} - \mathbf{u}^+}{\Delta t} = \mathbf{N}^* + \mathbf{A}^* - \frac{\nabla p^{n+1}}{J} + \nu \mathbf{L}(\mathbf{u}^{n+1}),$$
(4.6)

where the meaning of the superscripts in \mathbf{u}^+ and \mathbf{N}^* follows from the definitions of equations (3.37) and (3.38).

Another aspect that must be modified is the decomposition of the viscous term employed in the velocity-correction scheme. The relation $\nabla^2 \mathbf{u} = \nabla(\nabla \cdot \mathbf{u}) - \nabla \times \nabla \times \mathbf{u}$ that was used in section 3.4 is not very useful here, since the incompressibility condition is not $\nabla \cdot \mathbf{u} = 0$ any more. Thus, it is more convenient to use the identity $\nabla^2 \mathbf{u} = \nabla(D(\mathbf{u})) - \nabla(\frac{\mathbf{u}}{J} \cdot \nabla J) - \nabla \times \nabla \times \mathbf{u}$, since the first term in the right-hand-side vanishes, leading to a behaviour similar to the one that motivates this decomposition. Therefore, it is proposed that in this case the time-integration scheme of equations (3.43) and (3.44) should be replaced by:

$$\begin{cases} \frac{\gamma_{0}\bar{\mathbf{u}}^{n+1}-\mathbf{u}^{+}}{\Delta t} + \frac{\nabla p^{n+1}}{J} + \nu(\nabla(\frac{\mathbf{u}}{J}\cdot\nabla J))^{*} + \nu(\nabla\times\nabla\times\mathbf{u})^{*} - \mathbf{N}^{*} - \mathbf{A}^{*} = 0 & \text{in } \Omega\\ D(\bar{u}^{n+1}) = 0 & \text{in } \Omega\\ \bar{\mathbf{u}}^{n+1}\cdot\mathbf{n} = \mathbf{u}_{\mathfrak{D}}\cdot\mathbf{n} & \text{on } \Gamma_{\mathfrak{D}}, \end{cases}$$

$$(4.7)$$

$$\begin{cases} \frac{\gamma_{0}(\mathbf{u}^{n+1}-\bar{\mathbf{u}}^{n+1})}{\Delta t} - \nu \mathbf{L}(\mathbf{u}^{n+1}) - \nu (\nabla(\frac{\mathbf{u}}{J}\cdot\nabla J))^{*} - \nu (\nabla\times\nabla\times\mathbf{u})^{*} = 0 & \text{in } \Omega\\ \mathbf{u}^{n+1} = \mathbf{u}_{\mathfrak{D}} & \text{on } \Gamma_{\mathfrak{D}} & \text{(4.8)}\\ \frac{\partial \mathbf{u}^{n+1}}{\partial \mathbf{n}} = \mathbf{u}_{\mathfrak{N}} & \text{on } \Gamma_{\mathfrak{N}}. \end{cases}$$

To solve the first step, we multiply the equation by J, dot with gradients of the test functions $\nabla \Phi$ and integrate in the domain. Defining $\hat{\mathbf{u}} = \mathbf{u}^+ + \Delta t (\mathbf{N}^* + \mathbf{A}^*)$, the result is

$$\int_{\Omega} \nabla p^{n+1} \cdot \nabla \Phi \, d\Omega = \int_{\Omega} J \left[\frac{\hat{\mathbf{u}} - \gamma_0 \bar{\mathbf{u}}^{n+1}}{\Delta t} - \nu \left(\nabla \left(\frac{\mathbf{u}}{J} \cdot \nabla J \right) \right)^* - \nu (\nabla \times \nabla \times \mathbf{u})^* \right] \cdot \nabla \Phi \, d\Omega.$$
(4.9)

Applying the divergence theorem to the right-hand-side of the equation, and noting that the divergence of the curl is zero and that from the incompressibility condition $\nabla \cdot (J\bar{\mathbf{u}}^{n+1}) = 0$, we obtain

$$\int_{\Omega} \nabla p^{n+1} \cdot \nabla \Phi \, d\Omega = \int_{\Omega} \Phi \nabla \cdot \left[-\frac{J \hat{\mathbf{u}}}{\Delta t} + \nu \left(\nabla \left(\frac{\mathbf{u}}{J} \cdot \nabla J \right) \right)^* \right] + \Phi \nu \nabla J \cdot (\nabla \times \nabla \times \mathbf{u})^* \, d\Omega \\ + \int_{\Gamma} \Phi J \left[\frac{\hat{\mathbf{u}} - \gamma_0 \bar{\mathbf{u}}^{n+1}}{\Delta t} - \nu \left(\nabla \left(\frac{\mathbf{u}}{J} \cdot \nabla J \right) \right)^* - \nu (\nabla \times \nabla \times \mathbf{u})^* \right] \cdot \mathbf{n} \, dS.$$
(4.10)

This is still a Poisson equation in the weak form for the pressure, and can be solved using the same method as in the original velocity-correction scheme, with the only changes appearing in the forcing term and in the high-order boundary condition term in the righthand-side of the equation.

In the second step of the solution, by changing the definition of $\hat{\mathbf{u}}$ to $\hat{\mathbf{u}} = \hat{\mathbf{u}} - \frac{\nabla p^{n+1}}{J} \Delta t$, we once again obtain equation (3.48), and therefore this is the only change that needs to be done to this step. It is clear that the formulation presented in this section requires only small modifications in the original velocity-correction scheme procedure:

- a forcing term needs to by added to the convective term,
- the forcing terms for the pressure and viscous systems need to be modified,
- the high-order pressure boundary conditions are slightly different.

This characteristic is one of the advantages of this method, since an existing solver can be easily adapted to include the coordinate transformation.

Divergence-free coordinate transformation

An interesting particular case happens when the transformation of equation (4.1) is characterized by $f(\bar{x}, \bar{z}) = \bar{x} + \xi(\bar{z})$. In this case, the deformation is uniform along each xy plane, the Jacobian is J = 1 and we have what is called a divergence-free mapping. In this situation, the only modification that is required in the velocity-correction scheme is adding the forcing term $\mathbf{A}(\mathbf{u}, p)$ to the convective term. The resulting method is exactly the same as the one used by Darekar and Sherwin (2001b). This approach for divergence-free mappings was also applied in a more general context by Newman and Karniadakis (1997), who considered a time-dependent mapping including deformations in both x and y. Therefore, the explicit method proposed here can be considered as a generalized version of this method available in the literature.

In this scenario, instead of using tensor calculus to obtain the appropriate equations, it becomes easier to apply the transformation directly by noting that the velocity components in the two coordinate systems are related by

$$\begin{cases} u = \bar{u} + \bar{w} \frac{\partial \xi}{\partial \bar{z}} \\ v = \bar{v} \\ w = \bar{w}, \end{cases}$$
(4.11)

and the partial derivatives by

$$\begin{cases} \frac{\partial}{\partial x} &= \frac{\partial}{\partial \bar{x}} \\ \frac{\partial}{\partial y} &= \frac{\partial}{\partial \bar{y}} \\ \frac{\partial}{\partial z} &= \frac{\partial}{\partial \bar{z}} - \frac{\partial \xi}{\partial z} \frac{\partial}{\partial \bar{x}}. \end{cases}$$
(4.12)

Applying these relations to the Navier-Stokes equations, we obtain that the components of the forcing term in this case are:

$$A_{x} = \begin{bmatrix} \frac{\partial\xi}{\partial\bar{z}}\frac{\partial p}{\partial\bar{z}} - \left(\frac{\partial\xi}{\partial\bar{z}}\right)^{2}\frac{\partial p}{\partial\bar{x}} - \bar{w}^{2}\frac{\partial^{2}\xi}{\partial\bar{z}^{2}} \end{bmatrix} + \frac{1}{Re}\begin{bmatrix} \frac{\partial^{2}u}{\partial z^{2}} - \frac{\partial^{2}u}{\partial\bar{z}^{2}} + \frac{\partial^{2}}{\partial z^{2}}\left(\bar{w}\frac{\partial\xi}{\partial\bar{z}}\right) - \frac{\partial\xi}{\partial\bar{z}}\frac{\partial^{2}\bar{w}}{\partial z^{2}} \end{bmatrix},$$

$$A_{y} = \frac{1}{Re}\begin{bmatrix} \frac{\partial^{2}\bar{v}}{\partial z^{2}} - \frac{\partial^{2}\bar{v}}{\partial\bar{z}^{2}} \end{bmatrix},$$

$$A_{z} = \frac{\partial\xi}{\partial\bar{z}}\frac{\partial p}{\partial\bar{x}} + \frac{1}{Re}\begin{bmatrix} \frac{\partial^{2}\bar{w}}{\partial z^{2}} - \frac{\partial^{2}\bar{w}}{\partial\bar{z}^{2}} \end{bmatrix}.$$

$$(4.13)$$

Regardless of using equation (4.5) or equation (4.13) to calculate the forcing terms, the divergence-free mapping is obviously simpler and more efficient computationally, since it does not include any extra terms in the forcing and boundary terms of the pressure equation. With this in mind, this characteristic was exploited in all simulations satisfying this condition.

4.2 Semi-implicit treatment of the mapping

This section describes an approach to solve the equations where the mapping terms coming from the convective part of the equation are treated explicitly, while the pressure and viscous terms are treated implicitly, maintaining the characteristics of the original velocitycorrection scheme. This is a modified version of the method used by Carlson et al. (1995), and later restated by Koberg (2007) with a notation similar to equation (3.40). Therefore, their method will be briefly presented before moving to the modified approach, in order to show the motivation for the changes proposed. The main idea is to modify equation (3.40) using the appropriate operators from equation (4.2), leading to:

$$\frac{\hat{\mathbf{u}} - \mathbf{u}^{+}}{\Delta t} = \bar{\mathbf{N}}^{*},$$

$$\frac{\hat{\bar{\mathbf{u}}} - \hat{\mathbf{u}}}{\Delta t} = -\bar{G}(p^{n+1}),$$

$$\frac{\gamma_{0}\mathbf{u}^{n+1} - \hat{\bar{\mathbf{u}}}}{\Delta t} = \nu \bar{\mathbf{L}}(\mathbf{u}^{n+1}).$$
(4.14)

In the first step, all the convective terms (including the ones imposed by the mapping) are treated explicitly. To obtain the pressure in the second step, the generalized divergence operator D() is applied to the equation, which after imposing $D(\bar{\mathbf{u}}) = 0$ results in

$$D(\bar{G}(p^{n+1})) = D\left(\frac{\hat{\mathbf{u}}}{\Delta t}\right).$$
(4.15)

This equation can be solved iteratively by

$$\nabla^2 p_{s+1}^{n+1} = D\left(\frac{\hat{\mathbf{u}}}{\Delta t}\right) + \nabla^2 p_s^{n+1} - D(\bar{G}(p_s^{n+1})), \tag{4.16}$$

where *s* is the iteration counter. Note that the first term in the right-hand-side of the equation is independent of *s*, and therefore only needs to be computed once at each time step.

Similarly, the velocity equation can be solved using the iteration

$$\frac{\gamma_0 \mathbf{u}_{s+1}^{n+1}}{\Delta t} - \nu \mathbf{L}(\mathbf{u}_{s+1}^{n+1}) = \frac{\hat{\mathbf{u}}}{\Delta t} - \bar{G}(p^{n+1}) + \nu \bar{\mathbf{L}}(\mathbf{u}_s^{n+1}) - \nu \mathbf{L}(\mathbf{u}_s^{n+1}),$$
(4.17)

where once again only the terms which depend on s need to be calculated on every iteration.

The main problem with this time integration scheme is that it is not clear what the boundary conditions should be in equation (4.16). This is similar to the problem with the original splitting scheme of Karniadakis et al. (1991), but while in that case we could use the momentum equation to obtain a consistent boundary condition, equation (4.16) is an iterative procedure which is not directly related to a momentum equation. Therefore, a modified version of this approach leading to equations with consistent boundary conditions will be presented now. The first step is to note that the operator \bar{L} can be decomposed as

$$\bar{\mathbf{L}}(\mathbf{u}) = g^{jm} \bar{u}^i_{,mj} = g^{ji} (\bar{u}^m_{,m})_{,j} - \varepsilon^{imn} \varepsilon^{ljk} g_{nl} g_{kp} \bar{u}^p_{,jm},$$
(4.18)

where g_{ij} is the metric tensor and $\varepsilon^{ijk} = g^{-1/2} \epsilon^{ijk}$, with ϵ^{ijk} being the permutation symbol, is a generalization of the permutation symbol. Noting that $\bar{u}_{,m}^m = D(\mathbf{u}) = 0$, we conclude that this is a generalized form of the decomposition employed in the Cartesian case. Defining $\mathbf{Q} = \varepsilon^{imn} \varepsilon^{ljk} g_{nl} g_{kp} \bar{u}_{,jm}^p$, it is proposed to restate the velocity-correction scheme of equations (3.43) and (3.44) as

$$\begin{cases} \frac{\gamma_{0}\bar{\mathbf{u}}^{n+1}-\mathbf{u}^{+}}{\Delta t} + \bar{G}(p^{n+1}) + \nu \mathbf{Q}^{*} - \bar{\mathbf{N}}^{*} = 0 & in \ \Omega\\ D(\bar{\mathbf{u}}^{n+1}) = 0 & in \ \Omega\\ \bar{\mathbf{u}}^{n+1} \cdot \mathbf{n} = \mathbf{u}_{\mathfrak{D}} \cdot \mathbf{n} & on \ \Gamma_{\mathfrak{D}}, \end{cases}$$
(4.19)

$$\begin{cases} \frac{\gamma_{0}\left(\mathbf{u}^{n+1}-\bar{\mathbf{u}}^{n+1}\right)}{\Delta t}-\nu\bar{\mathbf{L}}(\mathbf{u}^{n+1})-\nu\mathbf{Q}^{*}=0 & in \ \Omega\\ \mathbf{u}^{n+1}=\mathbf{u}_{\mathfrak{D}} & on \ \Gamma_{\mathfrak{D}}\\ \frac{\partial\mathbf{u}^{n+1}}{\partial\mathbf{n}}=\mathbf{u}_{\mathfrak{N}} & on \ \Gamma_{\mathfrak{N}}. \end{cases}$$
(4.20)

The first step is solved using the following iteration:

$$\nabla p_{s+1}^{n+1} = \nabla p_s^{n+1} + J \left[\frac{\mathbf{u}^+ - \gamma_0 \bar{\mathbf{u}}^{n+1}}{\Delta t} - \nu \mathbf{Q}^* + \bar{\mathbf{N}}^* - \bar{G}(p^{n+1}) \right].$$
(4.21)

Dotting the equation with $\nabla \Phi$ and integrating to obtain the weak form, and after using the identities $\nabla \cdot (J\bar{\mathbf{u}}^{n+1}) = 0$ and $D(\mathbf{Q}) = 0$, this equation becomes

$$\int_{\Omega} \nabla p_{s+1}^{n+1} \cdot \nabla \Phi \, d\Omega = \int_{\Omega} \Phi \left[JD\left(\frac{-\hat{\mathbf{u}}}{\Delta t}\right) + JD(\bar{G}(p_s^{n+1})) - \nabla^2 p_s^{n+1} \right] \, d\Omega \\ + \int_{\Gamma} \Phi \left[J\left(\frac{\hat{\mathbf{u}} - \gamma_0 \bar{\mathbf{u}}^{n+1}}{\Delta t}\right) - \nu J \mathbf{Q}^* - J\bar{G}(p_s^{n+1}) + \nabla p_s^{n+1} \right] \cdot \mathbf{n} \, dS,$$
(4.22)

where $\hat{\mathbf{u}}$ follows the definition from equation (4.14).

An important difference in this case is that the iteration is defined as an equation for the pressure gradient instead of an equation for the Laplacian of the pressure. This allows for the determination of the complete weak equation for the pressure, including the boundary conditions. The velocity is again solved using equation (4.17), completing the time integration scheme.

4.3 Post-processing and forces calculation

The results obtained using the methods presented in sections 4.1 and 4.2 contain the representation of the velocity in the transformed coordinate system $(\bar{x}, \bar{y}, \bar{z})$, and not the physical velocities in the Cartesian system (x, y, z). Also, the computational domain corresponds to the transformed system, and thus we can only directly obtain the partial derivatives with respect to $(\bar{x}, \bar{y}, \bar{z})$. Therefore, it is important to be careful when processing these results, to guarantee that all results obtained are represented in the Cartesian system. First, the velocity is a contravariant vector, and according to equation (A.4) must be transformed by

$$u^{i} = \frac{\partial x^{i}}{\partial \bar{x}^{j}} \bar{u}^{j}, \qquad (4.23)$$

where the repeated index summation convention applies. In addition, the calculation of velocity gradients needs to be performed by the covariant derivatives, which are later transformed to the Cartesian system by

$$\frac{\partial u^{i}}{\partial x^{j}} = \frac{\partial x^{i}}{\partial \bar{x}^{k}} \frac{\partial \bar{x}^{l}}{\partial x^{j}} \bar{u}_{,l}^{k}, \qquad (4.24)$$

Alternatively, the velocity gradients can be obtained by first transforming the velocity to the Cartesian system, and then calculating the partial derivatives using the chain rule.

Another important aspect that needs to be addressed is the calculation of the forces acting on the wing. We begin by noting that in the usual Cartesian case, the forces can be obtained by

$$F^{i} = \int_{\Gamma} \sigma^{ij} dS_{j}, \qquad (4.25)$$

where

$$\sigma^{ij} = -\delta^{ij}p + \nu \left[\frac{\partial u^i}{\partial x^j} + \frac{\partial u^j}{\partial x^i}\right]$$
(4.26)

is the stress tensor in Cartesian coordinates and $dS_j = n_j dS$ is the surface element oriented by the normal vector.

This can be extended to the transformed system by noting that the appropriate form of the stress tensor and of the surface element are

$$\bar{\sigma}^{ij} = -g^{ij}p + \nu \left[g^{pj}\bar{u}^i_{,p} + g^{pi}\bar{u}^j_{,p}\right],\tag{4.27}$$

$$d\bar{S}_i = J\bar{n}_i d\bar{S},\tag{4.28}$$

where \bar{n}_i and $d\bar{S}$ are the normal vector and the area element calculated directly from the computational domain. This last relation can be obtained by noting that the surface element is defined by a cross product, which when applied in a general coordinate system to vec-

tors a and b is $\varepsilon_{ijk}a^jb^k = J(\epsilon_{ijk}a^jb^k)$, where ϵ_{ijk} is the permutation symbol. The term in parentheses leads to the usual surface element as calculated in the computational domain, which only needs to be multiplied by the Jacobian. Before integrating, we must also convert the integrand to the Cartesian system. This can be easily done since it is a contravariant vector, and after integrating the result we finally obtain the forces acting on the body.

4.4 Test cases

This section presents results of tests employing the previous methods in three different types of flows, the first two validating the formulation, with the third demonstrating how this approach is able to deal with very general problems with time-dependent transformations. The first case considered is the Kovasznay flow, for which there is an exact analytical solution, allowing for the accuracy of the methods to be evaluated. The second case is the flow around an infinite square cylinder deformed along the span direction. This is a more complex flow with results at low Reynold numbers available in the literature, allowing one to assess how the method performs in a more practical scenario. The last simulation presented is the two-dimensional flow around two circular cylinders in tandem, with the upstream cylinder subject to forced oscillations, while the downstream cylinder is held fixed.

4.4.1 Kovasznay flow

The Kovasznay flow, presented by Kovasznay (1948), consists in a steady analytical solution for the two-dimensional Navier-Stokes equations with a periodic direction, which can be viewed as a representation of the flow behind a two-dimensional grid. Since this flow has an analytical solution, it serves as a benchmark for the accuracy of the methods, being an important first test for the proposed methods. Considering a three-dimensional domain with the flow in the xz plane, with z being the periodic direction, the exact solution is

$$u = 1 - e^{kx} \cos(2\pi z),$$

$$v = 0,$$

$$w = \frac{k}{2\pi} e^{kx} \sin(2\pi z),$$

$$p = \frac{1}{2} (1 - e^{2kx}),$$

(4.29)

where the constant k is defined as

$$k = \frac{Re}{2} - \sqrt{\frac{Re^2}{4} + 4\pi^2},$$
(4.30)

with the Reynolds number in this case being defined as $Re = \frac{1}{\nu}$.



Figure 4.2: Convergence on the infinity norm of the u velocity error for the Kovasznay flow with different mappings.

Simulations for this flow were performed in four different situations. The first did not employ any mapping, and therefore serves as a reference to compare the methods proposed here. The second case considered a divergence-free mapping, using the transformation

$$x = \bar{x} + \xi(\bar{z}) = \bar{x} - 0.05 \cos(2\pi \bar{z}).$$
(4.31)

The third scenario considered an explicit treatment of the mapping, while the fourth dealt with the semi-implicit case, both using the mapping

$$x = f(\bar{x}, \bar{z}) = \bar{x} - 0.05 \frac{\tanh(1 - \bar{x})}{\tanh(1)} \cos(2\pi \bar{z}).$$
(4.32)

In all cases the y and z coordinates were not modified.

It is clear that the study of this simple flow does not benefit from the use of these transformations, and therefore the only motivation for using them here is for testing the accuracy of the methods.

The simulations consisted of advancing the equations in time for 15 time units using a time step of 0.001 with a first-order time integration, with Re = 40. The computational domain extends from -0.5 to 1.0 in \bar{x} , from -0.5 to 1.5 in \bar{y} and from 0 to 2 in \bar{z} , with the $\bar{x}\bar{y}$ plane represented by a uniform mesh consisting of 12 quadrilateral elements, with the \bar{z} direction discretized using a Fourier expansion. The initial conditions consisted of a uniform flow with u = 1.0, v = 0.1 and w = 0.1. The reference resolution considered 10 modes per direction in each element of the $\bar{x}\bar{y}$ plane (9th order polynomials) with 64 degrees of freedom for the Fourier expansion in the \bar{z} direction. From this reference, either the polynomial order or the number of Fourier modes were reduced, in order to determine how these parameters affect the error.

Figure 4.2 presents the results in term of the infinity norm (maximum value in the quadra-



Figure 4.3: Contours of the w velocity component of the Kovasznay flow in the xz plane for the simulations with no mapping and with the mapping of equation (4.32).

ture points) of the error for the u velocity. We note that for the polynomial order, the mapping $x = f(\bar{x}, \bar{z})$ has a slight loss of accuracy, but all cases still show a good convergence. In the case of the z direction resolution, the use of the mappings results in higher errors when the number of modes is small, although this is likely due to the fact that in the absence of any mapping the solution can be exactly represented by the Fourier expansion with just the mean mode and the second harmonic, corresponding to $N_z = 6$. As the number of modes increases, once again the errors are comparable in all cases. It is also clear that the choice between the explicit or semi-implicit treatment of the mappings has little effect on the accuracy of these results.

Figure 4.3 presents contours of the w velocity for the first case (no mapping) and the fourth case (mapping of equation (4.32)), both for the highest spatial resolution considered. We note that although the physical domain is different (since the computational domain is the same), the results are equivalent, showing how the exact solution is obtained even in the deformed geometry.

4.4.2 Square cylinder with waviness

The flow around an infinite square cylinder with spanwise sinusoidal waviness was studied numerically by Darekar and Sherwin (2001a,b), for a Reynolds number Re = 100. As noted in section 4.1, the method they employed is a particular case of the explicit method of section 4.1. Therefore, besides being closer to a practical application, this is also a good test to check if the current implementation is consistent with theirs. Four cases considered by them with different waviness amplitudes h were reproduced here, all of them with a waviness of wavelength $\frac{\lambda}{H} = 2.8$, where H is the height of the cylinder. For each case, a



Figure 4.4: Contours of spanwise vorticity for square cylinder with and without spanwise waviness with Re = 100. The waviness has a wavelength $\frac{\lambda}{H} = 2.8$ and amplitude $\frac{h}{\lambda} = 0.2$.

coordinate transformation

$$x = \bar{x} + \xi(\bar{z}) = x = \bar{x} - \frac{h}{2} \cos\left(\frac{2\pi}{\lambda}\bar{z}\right),$$
(4.33)

is employed to remove the waviness from the cylinder, allowing for the use of a Fourier discretization in the span direction. Figure 4.4 shows contours of spanwise vorticity for the cylinder without any waviness and for a waviness with amplitude $\frac{h}{\lambda} = 0.2$, illustrating the geometry being considered, and how vortex shedding is suppressed in the modified geometry.

The spatial discretization consisted of 827 quadrilateral elements extending from -10 to 20 in \bar{x} and -15 to 15 in \bar{y} , with the solution in each one of them represented by 6^{th} order polynomials, with the cylinder having a unit height. The spanwise length was $L_z = 2.8$, discretized using a Fourier series with 16 degrees of freedom. The equations were integrated in time for 200 time units, using a first order scheme with time step $\Delta t = 0.002$, with the last 10 time units being used to calculate the results.

The results for different waviness amplitudes are presented in figure 4.5, showing a good agreement. Although there is a discrepancy in the drag coefficient, this is probably due to different resolutions in both sets of simulations, since their convergence studies indicate errors around 2%, while the simulations presented here were not preceded by any rigorous convergence studies. The fact that this systematic error persists even for a zero amplitude, where no transformation is employed, further supports the claim that the discrepancies are not caused by the treatment of the waviness.



Figure 4.5: Comparison of results for a square cylinder with waviness of wavelength $\frac{\lambda}{H} = 2.8$ at Re = 100, with the results from Darekar and Sherwin (2001b). Each case consists of a different amplitude h.

4.4.3 Flow around moving cylinders

To demonstrate the possibility of using the techniques presented here to time-dependent transformations, a simulation of the two-dimensional flow around a pair of moving circular cylinders was performed, with Reynolds number Re = 100. The centre-to-centre distance is 3 diameters, the downstream cylinder is held fixed, and a forced oscillation in the y direction with non-dimensional frequency 0.3 and amplitude of 0.75 diameter was imposed on the upstream cylinder. Instead of using a moving mesh to solve this problem, a fixed mesh where the displacement of the cylinders is zero was used, with a mapping accounting for the movement of the upstream cylinder. Before each time step, the displacements on the boundaries were used as boundary conditions to solve a Laplace equation, leading to a global representation of the mapping which was used to solve the equations with the semi-implicit method of section 4.2.

The domain extended from -10 to 20 in \bar{x} and from -15 to 15 in \bar{y} , with both cylinders having a unit diameter and the origin being located in the centre of the upstream cylinder. This domain was divided into 702 quadrilateral elements, with the solution in each of them represented by 8^{th} order polynomials. The equations were integrated in time for 200 time units, using a first order scheme with time step $\Delta t = 0.0002$. The iterative parts of the semi-implicit method were solved using a tolerance 10^{-8} and relaxation parameters equal to 0.15.

Figure 4.6 shows instantaneous contours of vorticity for this case, exhibiting a behaviour that is compatible with what could be expected for this flow, although these results were not compared to any other investigations on this topic. The effect of the transformation is illustrated in figure 4.7, showing the computational mesh in the region close to the cylinders, along with the corresponding representations in the physical domain at two different time



(b) $t = t_0 + \frac{T}{2}$

Figure 4.6: Instantaneous contours of vorticity for flow around two circular cylinder in tandem with Re = 100. The downstream cylinder is fixed, while the upstream cylinder oscilates with frequency 0.3 and amplitude 0.75.



Figure 4.7: Detail of computational mesh used in the simulation of the flow around two circular cylinders, and the same mesh after applying the coordinate transformation at two different time instants.

instants.

4.5 Discussion

This chapter presented two methods of including coordinate transformations in the velocitycorrection scheme for solving the incompressible Navier-Stokes equations. The computational cost of advancing the equations by one time step using the semi-implicit method of section 4.2 is clearly higher than the one for the explicit approach of section 4.1. Therefore, it is important to address the question of under what circumstances the former should be employed. Although section 4.4.1 showed that both schemes maintain the accuracy of the underlying discretisation, the explicit scheme becomes unstable as the coordinate transformation becomes more energetic. Therefore, it is reasonable to anticipate that the explicit method is only applicable to less energetic or smaller coordinate changes, while the semi-implicit method with its higher computational cost is more capable of handling more complex or energetic deformations. In these more complex cases, the iterative procedures employed in the semi-implicit method can also become unstable as the deformation of the mapping becomes even larger; however, this issue can be moderated by including a relaxation technique.

To illustrate the limitation of the explicit and semi-implicit schemes, the performance of

both methods for the square cylinder case of section 4.4.2 is considered. For this particular case, the maximum amplitude that could be stably handled using the explicit method is $h = 0.25\lambda$ (where $\lambda = 2.8H$), while the semi-implicit method remained stable for a significantly higher amplitude of $h = 0.75\lambda$, after which convergence becomes increasingly difficult. In terms of computational cost, in the implementation employed in the present work the explicit method takes approximately 1.7 times the time required for the straight cylinder case. For the semi-implicit method, the computational cost depends significantly on the number of iterations required in the iterative solutions, with typical cost between 5 and 10 times that of the undeformed case. Obviously, these increases in cost have to be weighted against the gains obtained by being able to employ a quasi-3D discretization instead of a full three-dimensional simulation.

When considering the computational costs, it is important to note that the cost is influenced by the particular problem being considered and by the characteristics of the numerical implementation being employed, especially by the balance between the costs of the implicit and explicit parts of the algorithm. In the simulations presented here, the implicit equations were solved using an efficient direct solver with multi-level static condensation, as described in Karniadakis and Sherwin (2005). Because of this, the time spent calculating the advection terms in the original solver was significant (around 30%), emphasizing increases in the cost of the explicit calculations. If the implicit equations were solved using a less efficient iterative solver, as is often the case in large massively-parallel simulations, the solution of the implicit systems would dominate the cost, and thus the performance of our explicit method the ratio between the cost of computing the mapping terms and computing the advection terms might be a more representative measure of the performance of the method. Determining this ratio, which is related to the number of derivatives that need to be calculated, we observe it to be close to 2.5.

When considering the case of wavy wings, which is the topic of interest for the present work, we observe that the explicit approach is stable until amplitudes of approximately $h = 0.2\lambda$. All simulations of part III stayed within this limit, and therefore it was possible to use the cheaper explicit approach. However, if we wished to extend the parameter space to larger amplitudes, it would be necessary to switch to the implicit approach, increasing the computational cost. We also observe in this case that increasing the Reynolds number does not significantly affect the stability of these methods.

Another important point is that the two methods can be combined, leading to a formulation where only the pressure or the viscous terms are treated implicitly. This can be useful if evidence indicates that the convergence problems caused by mappings with high amplitude are related to only one of these terms. In this situation, only the term leading to stiffness of the system needs to be treated implicitly, while all the other terms remain explicit, reducing the computational cost. Finally, it should be noted that although the focus of the present work is on fixed coordinate changes, this approach can be very useful when time-dependent transformations are required. The example of section 4.4.3 shows how the methods proposed here can lead to a general framework to dealing with fluid-structure interaction without employing moving meshes, which can be computationally expensive in spectral/hp methods, due to the requirement of reassembling the pressure and velocity systems to account for a new geometry.

5. Numerical considerations for high Reynolds number simulations

For the lowest Reynolds number considered in this work, Re = 1,000, the flow is laminar and the computational cost of the simulations is reasonably low. In this case, the traditional spectral/hp element method presented in chapter 3, coupled with the mapping techniques of chapter 4, was sufficient to perform the study. However, as the Reynolds number increases and the flow starts transitioning to a turbulent state, some difficulties arise. This chapter briefly describes the problems encountered in these situations, and discusses what was done to circumvent each of them.

5.1 The spectral-vanishing viscosity

The high-order methods presented in chapter 3 have low numerical diffusion, as discussed in Karniadakis and Sherwin (2005). As a consequence of that, simulations with marginal spatial resolution become numerically unstable, especially in the presence of turbulence at high Reynold numbers. One way of circumventing this problem is to use what is known as spectral-vanishing viscosity (SVV), which is a technique that adds artificial dissipation to the smallest scales of the solution, in order to stabilize the simulation. Although for the lowest Reynolds number considered in the present work, Re = 1,000, this sort of stabilization was not required, this technique had to be employed in the other cases. The formulation considered in the present work will be briefly described in this section, and it is based on Kirby and Sherwin (2006b), which also presents a good review on the historical development of SVV.

The main idea of SVV consists in expanding the Navier-Stokes equations (equation (3.31)) to include an artificial dissipation operator, leading to

$$\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u} \cdot \nabla)\mathbf{u} - \nabla p + \nu \nabla^2 \mathbf{u} + S_{VV}(\mathbf{u}),$$

$$\nabla \cdot \mathbf{u} = 0$$
(5.1)

$$S_{VV}(u) = \epsilon \sum_{i=1}^{N_{dim}} \frac{\partial}{\partial x_i} \left[Q_i \star \frac{\partial u}{\partial x_i} \right],$$
(5.2)

with N_{dim} being the spatial dimension of the problem, ϵ a constant coefficient, and \star representing the application of the filter Q_i through a convolution operation. Therefore, the implementation of SVV depends on obtaining the matrix form of the operator from equation (5.2), and defining the filter Q_i . After obtaining the SVV matrix, it can be incorporated into the velocity-correction scheme by adding it to the Laplacian matrix in the solution of the velocity fields.

Kirby and Sherwin (2006b) show that the SVV operator in a single element can be obtained by

$$\mathbf{S}_{VV}^{e} = \epsilon \sum_{i=1}^{N_{dim}} \tilde{\mathbf{S}}_{i}^{T} \hat{\mathbf{Q}} \tilde{\mathbf{M}}^{-1} \tilde{\mathbf{S}}_{i},$$
(5.3)

with

$$\widetilde{\mathbf{S}}_i = \widetilde{\mathbf{B}}^T \mathbf{W} \mathbf{D}_i \mathbf{B},$$
 (5.4)

where for quadrilateral elements the matrices with a tilde correspond to the basis formed by the tensor product of Legendre polynomials, while the other matrices correspond to the basis employed in the solution. In these equations, \mathbf{M} is the mass matrix, \mathbf{W} is a diagonal matrix containing the quadrature weights, \mathbf{B} contains the values of the basis functions at the quadrature nodes, \mathbf{D}_i is the matrix representing numerical differentiation in the i - thdirection, $\hat{\mathbf{Q}}$ represents the filter in modal space, and $\tilde{\mathbf{B}}$ is the analogous of \mathbf{B} using a basis formed by the tensor product of Legendre polynomials. Reading this equation from right to left, this can be interpreted as the following sequence of operations: backward transformation (\mathbf{B}), differentiation in physical space (\mathbf{D}_i), projection to Legendre polynomials basis ($\tilde{\mathbf{M}}^{-1}\tilde{\mathbf{B}}^T\mathbf{W}$), application of the filter ($\hat{\mathbf{Q}}$), transformation to physical space ($\tilde{\mathbf{B}}$), and inner-product with respect to derivatives of the basis functions (($\mathbf{D}_i\mathbf{B}$)^T \mathbf{W}).

Although the matrices \tilde{B} and \tilde{M} could correspond to any polynomial basis spanning the same space as the one representing the solution (including using that same basis), one of the motivations for using the Legendre basis is to make the operator symmetric positive-definite, similarly to the original Laplacian operator, as demonstrated in Kirby and Sherwin (2006b). However, this is only true when \tilde{M} is diagonal, which is only true if the element has a constant Jacobian (a parallelogram in the case of quadrilaterals). For elements that do not satisfy this condition, the diagonal nature of the mass matrix corresponding to the tensor-product Legendre basis can be recovered by performing the projection to this basis in the standard element, but this also breaks the symmetry of the operator because the quadrature weights for this operation would be different. To understand this last point, we denote the quadrature weights in the standard element as W_{std} , and thus the full form of

Reynolds number	Angle of attack	P_{cut}	ϵ
10,000 50,000	all cases 6°	$\begin{array}{c} 0.5P \\ 0.5P \end{array}$	$0.1 \\ 0.1$
)	15°	0.5P	1.0

Table 5.1: Parameters for the spectral-vanishing viscosity employed in the present work.

the operator in this case becomes

$$\mathbf{S}_{VV}^{e} = \epsilon \sum_{i=1}^{N_{dim}} \left(\tilde{\mathbf{B}}^{T} \mathbf{W} \mathbf{D}_{i} \mathbf{B} \right)^{T} \hat{\mathbf{Q}} \left(\tilde{\mathbf{B}}^{T} \mathbf{W}_{std} \tilde{\mathbf{B}} \right)^{-1} \left(\tilde{\mathbf{B}}^{T} \mathbf{W}_{std} \mathbf{D}_{i} \mathbf{B} \right).$$
(5.5)

Although in this case the mass matrix $\tilde{\mathbf{B}}^T \mathbf{W}_{std} \tilde{\mathbf{B}}$ is diagonal, the operator is not symmetric because the terms in the first and last brackets use different forms of \mathbf{W} . To recover the symmetry of the operator, the previous formulation was modified using the element Jacobian. By noting that $\mathbf{W} = \mathbf{J}\mathbf{W}_{std}$, where \mathbf{J} is the diagonal matrix with the absolute value of the Jacobian in the quadrature points, we can obtain a symmetric positive-definite operator by using

$$\mathbf{S}_{VV}^{e} = \epsilon \sum_{i=1}^{N_{dim}} \left(\tilde{\mathbf{B}}^{T} \mathbf{J}^{-\frac{1}{2}} \mathbf{W} \mathbf{D}_{i} \mathbf{B} \right)^{T} \hat{\mathbf{Q}} \left(\tilde{\mathbf{B}}^{T} \mathbf{W}_{std} \tilde{\mathbf{B}} \right)^{-1} \left(\tilde{\mathbf{B}}^{T} \mathbf{J}^{\frac{1}{2}} \mathbf{W}_{std} \mathbf{D}_{i} \mathbf{B} \right).$$
(5.6)

This corresponds to a multiplication by the square root of the Jacobian before projecting to the Legendre basis, and a division by the same factor after projecting the filtered values back to physical space. Also, we note that when the Jacobian is constant, this form becomes equivalent to equation (5.3). By modifying the exponents in the newly introduced Jacobian terms, we could also obtain a formulation of the SVV which scales with a power of the Jacobian in each element.

The other point that needs to be defined is the form of the filter determined by \hat{Q} . The form of the filter employed here is

$$\hat{\mathbf{Q}}[i(p,q)][i(p,q)] = \begin{cases} 0, & p+q \le P_{cut} \\ e^{-\frac{(p+q-P)^2}{(p+q-P_{cut})^2}}, & p+q > P_{cut}, \end{cases}$$
(5.7)

where i(p,q) represents the numbering of the modes $\Phi_i = L_p(\xi_1)L_q(\xi_2)$, and P_{cut} represents a threshold after which the modes start to be filtered, and is usually given as a percentage of the polynomial order P. The values of P_{cut} and ϵ employed in the simulations of part III are presented in table 5.1.

In the case of the homogeneous direction, the procedure for applying SVV is more straightforward. Considering the properties of the Fourier series, and since the Helmholtz equation corresponding to the viscous step is solved independently for each Fourier mode,

the only change that needs to be done is to add a factor to the viscosity of each affected mode. This factor is obtained in a similar fashion to equation (5.7), with p + q replaced by the Fourier mode number and P with the total number of modes.

5.2 Parallelization strategy

In order to perform simulations of complex flows within a reasonable time, it is essential to have efficient parallel algorithms that allow us to make appropriate use of the resources available in modern supercomputers. In the context of the Fourier-spectral/hp element method, first presented by Karniadakis (1990), the most natural approach in terms of parallelism is to have different processes responsible for different Fourier planes. This idea, which is illustrated in figure 5.1b, has been explored for example by Evangelinos and Karniadakis (1996), who showed that it leads to a very efficient parallel algorithm. This occurs because the Fourier modes are not coupled in the linear part of the Navier-Stokes equations, and therefore communication is only required when performing the FFTs for explicitly computing the non-linear terms in physical space.

Despite the fact that the parallelization in the homogeneous direction is efficient, it has a significant limitation: the maximum number of processes in this approach is equal to half the number of degrees of freedom in the Fourier expansion. For the simulations with Re = 1,000 presented in chapter 6, this limit was eight processes. Considering the computational cost of this case, this allowed running each simulation in around 40 hours, what was considered acceptable for the purposes of the present work. However, for the simulations with higher Reynolds numbers, relying solely on this strategy would cause the wall-clock time of each simulation to increase in a manner which would compromised the feasibility of the study. As an example, each simulation with Re = 10,000 would take around one and a half months.

The alternative to parallelism in the the Fourier modes is applying a mesh decomposition in the xy plane, as represented in figure 5.1c. Although in most cases this allows for a larger number of processes to be used (equal to the number of elements in the twodimensional mesh), this approach might not retain the efficiency of the previous method. This happens because in this case the coupling between the partitions occurs in the implicit part of the velocity-correction scheme, and therefore it relies on the solution of linear systems in parallel, which is not a trivial task.

In order to use more processes than the traditional Fourier parallelism allows, while attempting to keep its high efficiency, the simulations of the present work employed the hybrid approach of Bolis et al. (2016). In this method, illustrated in figure 5.1d, the two-dimensional mesh is divided in P_{XY} partitions and the Fourier planes are divided in P_Z groups, leading to a total of $P_{TOT} = P_{XY}P_Z$ processes. Bolis et al. (2016) showed that for a large number of processes, this approach can be more efficient than a pure mesh



Figure 5.1: Illustration of different parallelization strategies using four processes. Extracted from Bolis (2013).

decomposition.

The present work extends that of Bolis et al. (2016) by using a slightly modified approach to the hybrid parallelism and performing an extensive scaling study. The main difference between the approach adopted here and that of Bolis et al. (2016) is that here the linear systems were solved by a direct solver based on the XX^T matrix decomposition, as presented by Tufo and Fischer (2001), instead of using an iterative solution based on the conjugate gradient method, as was done by them. Although the direct solver scales poorly for a large number of partitions, for a small number of partitions (as was the case here), it is expected to be much faster, since we avoid the loses incurred when switching to an iterative solver. The reduced number of partitions can be attributed to the hybrid parallelism, since we can transfer most of the burden to the homogeneous parallelism. Therefore, the possibility of using direct solvers for the linear systems of large problems is a peculiarity of Fourier-spectral/hp methods, and it is unlikely to occur in fully three-dimensional simulations.

Having determined the parallel implementation, the question of what values of P_Z and P_{XY} should be used still remained. In particular, it is desirable to have a large P_{TOT} (to obtain results in a shorter time), while avoiding a low efficiency which would lead to a waste of computational resources. To determine appropriate values for these parameters, an extensive study of their effect on the performance of the simulations for concrete cases of the current work was performed.

The simulations were based in the cases with Re = 10,000 presented in chapter 7, using one configuration of a straight wing and another of a wavy wing. For each combination of P_{XY} and P_Z , the equations were solved for 50 time steps, and the median computation time per time step was taken as a representative value to evaluate the computational cost. The simulations were performed in the Archer national supercomputer, where each node has 24 cores. Therefore, P_{TOT} is always a multiple of 24, with a minimum value of 48 since a single node did not have enough memory to perform the simulation. Although it would be possible to obtain other values of P_{TOT} by leaving idle cores, these situations are not considered here as they tend to overestimate the performance.

Figure 5.2 presents results obtained for the straight wing. In figure 5.2a the scaling with respect to P_{TOT} is shown for all cases. The results seem to follow the theoretical P_{TOT}^{-1} slope scaling until $P_{TOT} \approx 500$, after which there is a large loss of efficiency. Figures 5.2b and 5.2c consider the separate effects of P_{XY} and P_Z . In these figures, the y-axis shows the product of the time per time step by P_{TOT} , which is a measure of the cost of the simulation. First, we note that by changing P_Z for a constant P_{XY} , the more pronounced increase in cost comes for $P_Z = 64$, with lower values showing a good scaling. On the other hand, by increasing P_{XY} the computational cost increases progressively.

Figure 5.3 shows the same results for the tests with a wavy wing. In this case, the points in the overall scaling are closer to the theoretical slope, and only deviate around



Figure 5.2: Scaling of the parallel simulations for the straight wing with Re = 10,000.

 $P_{TOT} \approx 1000$. When considering the effect of P_Z , once again the main losses appear when moving from $P_Z = 32$ to $P_Z = 64$. However, the scaling with P_{XY} is now a lot better, showing good results even for 18 partitions. This can be understood by considering that these simulations use the explicit mapping of section 4.1. This method increases the cost of the explicit part of the solution, and therefore makes inefficiencies in the linear system solver less important, favouring the mesh decomposition approach.

Although the previous results provide a good description of the effect of the individual parameters in the computation cost, they are not of much use in choosing these parameters appropriately. In order to do so, first the following efficiency is defined:

$$\eta_p = \frac{(P_{TOT}T)_{ref}}{P_{TOT}T},\tag{5.8}$$

where *T* is the time per time step and the subscript ref refers to an arbitrary reference configuration, here considered as $P_Z = 8$ and $P_{XY} = 6$. Note that this efficiency can be larger than 1, since there is no guarantee that the reference is the most efficient configuration.

The efficiency defined by equation (5.8) is plotted against P_{TOT} in figure 5.4, with each curve representing a different value of P_Z . By considering this figure, we can first choose the maximum P_{TOT} that can be used based on an acceptable efficiency level. Then, for this P_{TOT} the best combination of P_Z and P_{XY} will be given by the highest curve. For the straight wing, this usually involves using $P_Z = 32$ or 64. On the other hand, for the wavy wing other combinations are equally acceptable for moderate number of processes, as there is not a clear trend in this case.

The previous efficiency results can also be represented by the isocontours of figure 5.5. This plot needs to be considered with caution, since only the points in the intersection of the grid lines were computed, with the rest consisting of an interpolation. Taking this into account, the process is similar to the previous figure. As a first step, we can define regions where the efficiency is acceptable under some specified criterion, and determine the maximum P_{TOT} in this region. Then, we can also consider a line of constant P_{TOT} (the solid lines in the figure) to determine the best combinations of P_Z and P_{XY} .

It is important to highlight that the results presented in this section depend on the particular case considered and on the hardware being used. In fact, significant differences in the timings can be observed even while keeping these factors constant. Therefore, considering that the computational cost of the tests presented here are low compared to that of the actual simulations, performing this type of study is recommended when running simulations that consume a large amount of computational time. However, when a quick choice of P_Z and P_{XY} is required, some simple guidelines can be obtained based on the previous results:

• Start with a pure homogeneous direction parallelization, choosing *P*_Z such that the number of Fourier planes per process is 4.



Figure 5.3: Scaling of the parallel simulations for wavy wing with Re = 10,000.



(b) Wavy wing

Figure 5.4: Efficiency of the parallelization, compared to the reference $P_Z = 8$, $P_{XY} = 6$. Simulations with Re = 10,000.



Figure 5.5: Isocontours of the efficiency of the parallelization (colours), for simulations with Re = 10,000. The solid lines represent a constant number of processes P_{TOT} .

- If more processes are required, switch to a hybrid parallelization using the same P_Z , and choose a P_{XY} that will make the number of local degrees of freedom (in the xy plane) per partition $\frac{N_{el}(P+1)^2}{P_{XY}}$ more than 20,000 (10,000 for a wavy geometry). For example, in the current case the number of elements in the mesh is $N_{el} = 1042$ and the polynomial order is P = 10, resulting in 6 partitions for the straight wing and 12 for the wavy wing.
- If a higher P_{TOT} is still desired, tests should be performed to determine the effect on the computational cost of doubling P_Z or increasing P_{XY} . If these changes lead to an unacceptable loss of efficiency, then the current method is not appropriate for the particular case, and other possibilities need to be considered (like using an iterative solver for the linear systems).

These are rough estimates, and therefore should be used with caution, as they can be affected significantly by changes in hardware or in the case considered.

To conclude this section, note that the parameters used in the simulations of chapter 7 ($P_Z = 64$ and $P_{XY} = 6$ for the straight wing, $P_Z = 64$ and $P_{XY} = 9$ for the wavy wing) lead to efficiencies around 0.9. Therefore, the use of the hybrid parallelization strategy allowed the simulation to be accelerated by a factor of almost 9 (compared to the parallelization only in the homogeneous direction), without large increases in the computational cost. This allowed simulations that would take around one and a half months to be performed in less than one week, illustrating the profound effect this parallelism study had on the viability of the simulations.

5.3 Adaptive order refinement

In the spectral/hp methods, the spatial resolution can be increased by refining the mesh (usually called h refinement), by increasing the polynomial order (p refinement), or by a combination of both. Even though the computational mesh allows us to cluster more degrees of freedom in regions where we expect the solution to contain higher gradients, this approach has limitations, since the details of the flow are not known a priori. This problem becomes more significant as the Reynolds number is increased towards the laminar-turbulent transition range, since the exact extent of the region of turbulent flow can vary greatly, with turbulence restricted to a thin boundary layer in the pre-stall regime, while in the post-stall it spreads over a large wake. Therefore, it is desirable to have a method of adjusting the spatial discretisation to the actual solution, leading to a more efficient distribution of the degrees of freedom and reducing the computational cost.

An adaptive refinement procedure was developed and used in the simulations of chapter 8. In this procedure, the mesh is held fixed, and serves the usual purpose of providing a good distribution of degrees of freedom based on the expected behaviour of the flow. During the solution, an estimate of the spatial discretisation error is calculated in each element, and based on this error the polynomial order in the element is adjusted to obtain a more efficient distribution of degrees of freedom. This estimate for the error (sometimes called *sensor*) was based on the one used for shock capture in Persson and Peraire (2006). In the present work, this is defined as

$$S_e = \frac{||u_P - u_{P-1}||_{2,e}^2}{||u_P||_{2,e}^2},$$
(5.9)

where u_P is the solution obtained for the u velocity using the current polynomial order P, u_{P-1} is the projection of this solution to a polynomial of order P-1, $|| \cdot ||_2$ is the L_2 norm and the subscript e indicates that this refers to a single element. In the homogeneous direction, S_e was taken as the maximum value of equation (5.9) in the element among the two-dimensional planes.

Considering this estimate for the discretisation error, the following adaptive procedure was used:

- 1. Advance the equation for a number of time steps n_{steps}
- 2. Calculate the sensor using equation (5.9)
- 3. Modify the polynomial order in each element using the following rule:
 - If $S_e \ge \epsilon_u$ and $P < P_{max}$, increase P by 1
 - If $S_e \leq \epsilon_l$ and $P > P_{min}$, decrease P by 1
 - Maintain same P if none of the above is true
- 4. Project the solution to the new polynomial space
- 5. Repeat the procedure for n_{runs}

where ϵ_u is the tolerance above which the polynomial order is increased, $\epsilon_l \leq \epsilon_u$ is the tolerance below which the polynomial order is decreased and P_{min} and P_{max} are the minimum and maximum polynomial orders imposed on the procedure. The tolerances can be functions of the spatial coordinates, in which case their value in each element is obtained by averaging the function inside the element. This is useful for imposing more strict tolerances close to the wing, allowing for larger errors far from it.

It is important to note that restarting the solution after the polynomial orders have been changed is an expensive operation, since the linear systems for the implicit part of the problem need to be assembled and decomposed again. Therefore, n_{steps} needs to be chosen carefully, since a low value will lead to an inefficient solution, while with a high value the refinement might not capture accurately changes in the flow structures.

As an example, figure 5.6 shows the distribution of polynomial order for a simulation similar to those of chapter 8. In this case the Reynolds number is Re = 50,000 and the



Figure 5.6: Polynomial order distribution obtained with adaptive order procedure. Simulation with Re = 50,000 and $\alpha = 15^{\circ}$.

angle of attack is $\alpha = 15^{\circ}$, using $P_{min} = 2$ and $P_{max} = 9$. It is clear that the boundary layers and the regions of turbulent separated flow are represented by high order polynomials, while lower orders are used in regions of laminar flow far from the wing. In this case, the average number of degrees of freedom per element is approximately 49, which is equivalent to the value for a constant P = 6 simulation. In terms of computational cost, table 5.2 compares the cost of this simulation using the adaptive procedure with $n_{steps} = 4,000$ with the cost for several different values of constant polynomial order, and with using the same variable polynomial order distribution without performing the adaptive procedure. The only difference between the simulations is in regard to the polynomial order, with the time step and all other parameters assuming always the same value. We note that for this value of n_{steps} , the refinement procedure corresponds to 5% of the computational cost. This is more than offset by the gains obtained from using a more efficient distribution of degrees of freedom, with the adaptive case presenting roughly the same cost as the P = 7 case, and being 35% faster than the P = 9 case. Therefore, if we consider that the accuracy achieved with a variable order is similar to that obtained using a constant order equals to P_{max} , then in this example the adaptive refinement procedure allows us to increase the order by two without increasing the cost, or to achieve the same accuracy with a much lower cost.
Table 5.2: Comparison of the computational cost of adaptive order case of figure 5.6 with constant uniform polynomial order and with variable order without adaptive procedure. The computational costs are normalized with respect to the adaptive order case.

Case	Cost	$\frac{1}{Cost}$
P = 5	0.60	1.66
P = 6	0.72	1.39
P = 7	1.08	0.93
P = 8	1.19	0.84
P = 9	1.53	0.65
variable order (static)	0.95	1.05
adaptive order	1.00	1.00

Part III

Results

6. Numerical simulations with $\mathbf{Re} = \mathbf{1},000$

This chapter presents the results for a Reynolds number based on the mean chord and on the freestream velocity Re = 1,000. This very low value of Re leads to a low computational cost, allowing for the use of fully resolved direct numerical simulations of the flow in a extensive parametric study. Moreover, the flow at low Re does not contain small-scale structures typical of turbulent flow, and therefore it is easier to identify the main structures causing the observed aerodynamic behaviour. In later chapters, the analysis is extended to higher values of Reynolds number, with the findings of the present chapter serving as a starting point in the analysis of the flow.

The profile chosen for the wing in all simulations is a NACA0012, which is given by the following equation:

$$\frac{y}{c} = \pm \frac{0.12}{0.2} \left[0.2969 \sqrt{\frac{x}{c}} - 0.1260 \left(\frac{x}{c}\right) - 0.3516 \left(\frac{x}{c}\right)^2 + 0.2843 \left(\frac{x}{c}\right)^3 - 0.1015 \left(\frac{x}{c}\right)^4 \right]$$
(6.1)

where c is the chord. However, this results in a profile with a trailing edge with a thickness of approximately 0.25% of the chord. This trailing edge represents a difficulty, since it demands that the mesh have very small elements to capture it. A simulation with such small elements would require a very small time step because of the CFL number restriction, and therefore would have a high computational cost. To avoid this problem, the following modified form of the NACA0012 profile was adopted:

$$\frac{y}{c} = \pm \frac{0.12}{0.2} \left[0.2969 \sqrt{\frac{x}{c}} - 0.1260 \left(\frac{x}{c}\right) - 0.3516 \left(\frac{x}{c}\right)^2 + 0.2843 \left(\frac{x}{c}\right)^3 - 0.1036 \left(\frac{x}{c}\right)^4 \right]$$
(6.2)

This modified profile is almost identical to the original, but it has a sharp trailing edge with zero thickness.

First, simulations were carried out for two-dimensional and three-dimensional straight infinite wings, which we will refer to as the baseline case. Subsequently, wings with waviness represented by the transformation $x = \bar{x} - \frac{h}{2}cos\left(\frac{2\pi}{\lambda}\bar{z}\right)$ were considered, where *h* is the waviness peak-to-peak amplitude and λ is its wavelength. Figure 6.1 shows an example of a geometry obtained through this transformation, identifying the waviness parameters and



Figure 6.1: Geometry of a wavy wing with h/c = 0.1 and $\lambda/c = 0.5$.

the orientation and origin of the coordinate system. This transformation leads to a uniform displacement in each xy plane, with both leading and trailing edges being deformed. Also, it corresponds to a divergence-free mapping, reducing the computational cost, as described in section 4.1. For this type of mapping, nine geometric configurations were considered consisting of different wavelengths λ and amplitudes h. Then, non-divergence-free mappings were used to study four configurations where the waviness is restricted to only the leading or the trailing edge of the wing, isolating each of the effects. Finally, two configurations were considered with different forms of defining a wing with deformations only in the leading edge, the first with a constant absolute thickness and the second with a constant profile. The naming convention adopted for the cases considered is provided in table 6.1.

For each case considered, simulations were performed for angles of attack between 0 and 21 degrees, with increments of 3 degrees between each simulation. All simulations were performed using the mesh and simulation parameters obtained through the convergence study presented in appendix B. In all simulations the equations were solved for 100non-dimensional time units using a time step $\Delta t = 0.0005$, with the last 10 units being used to calculate the results. The spatial discretization consisted of a mesh with 549 guadrilateral elements in the xy plane employing 8^{th} order polynomials. This mesh extends from -10 to 10 in the x directions and from -15 to 15 in the y directions, with the origin positioned at the leading edge of the airfoil and the unit-length chord aligned with the x axis. For the three-dimensional simulations, the periodic length in the span direction was equal to the chord, and was discretized using a Fourier expansion with 16 degrees of freedom. Also, the three-dimensional simulations were started using the two-dimensional result with the same angle of attack as the initial condition. Sections 6.1 to 6.4 present the results obtained using the divergence-free mappings, while section 6.5 show how different types of waviness, like restricting the waviness to only the leading edge or to the trailing edge, affect the results.

Case	λ/c	h/c	Notes
baseline	_	0.0	
L025h0125	0.25	0.0125	
L025h025	0.25	0.025	
L025h05	0.25	0.05	
L05h025	0.5	0.025	
L05h05	0.5	0.05	
L05h10	0.5	0.1	
L10h05	1.0	0.05	
L10h10	1.0	0.1	
L10h20	1.0	0.2	
L05h10_LE	0.5	0.1	deformation only on leading edge
L05h10_TE	0.5	0.1	deformation only on trailing edge
L10h10_LE	1.0	0.1	deformation only on leading edge
L10h10_TE	1.0	0.1	deformation only on trailing edge
L05h05_LE	0.5	0.05	deformation only on leading edge,
(constant thickness)			constant absolute thickness
L05h05_LE	0.5	0.05	deformation only on leading edge,
(constant profile)			constant profile

Table 6.1: Parameters of the waviness for the cases analyzed.

6.1 Aerodynamic forces

Figure 6.2 shows the results obtained for a straight infinite wing, comparing the results of two-dimensional and three-dimensional simulations. The results in this figure represent the mean drag coefficient ($\overline{C_d}$), the mean lift coefficient ($\overline{C_l}$), the lift-to-drag ratio (L/D), and the rms value of the lift coefficient fluctuations (C'_l). It can be noted that for angles of attack lower than $\alpha = 12^\circ$, the results of the two sets of simulations are almost identical. Starting at $\alpha = 15^\circ$, the three-dimensional effects in the flow lead to a reduction in $\overline{C_d}$ and $\overline{C_l}$ that result in a slight reduction of L/D, and there is also a significant reduction in C'_l .

Figure 6.3 presents the effect of the different geometries on the L/D value, with the cases separated by wavelength and the results for the three-dimensional simulations of a straight wing presented as a reference. For the shortest wavelength $\lambda/c = 0.25$, the change in the results is negligible for all cases, indicating that at this value of Re this scale is too short to have a significant effect on the flow. For the cases with $\lambda/c = 0.5$ and $\lambda/c = 1.0$, a significant reduction in L/D is observed. This reduction becomes more notable and occurs for a wider range of angles of attack as the waviness amplitude is increased, making the L/D value less responsive to variations in the angle of attack.

Additional results for the simulations with wavelength $\lambda/c = 0.5$ are presented in figure 6.4. It can be noted that the reduction in L/D is a consequence of a combination of re-



Figure 6.2: Results obtained for a straight wing using two-dimensional and three-dimensional simulations.



Figure 6.3: Comparison between L/D results obtained for wavy wings with different wavelengths with the results from the reference straight wing.



Figure 6.4: Results for wavy wings with wavelength $\lambda/c = 0.5$.

ductions in both lift and drag forces. Also, the waviness leads to a suppression of the lift coefficient fluctuations, such that in the case L05h05 C'_l remain close to zero until $\alpha = 12^\circ$, while in the case L05h10 it remains at low values until $\alpha = 15^\circ$.

6.2 Separation and recirculation zones

In order to obtain a better understanding of the characteristics of the flow which lead to the behaviour described above, figure 6.5 shows skin-friction lines on the surface of the wing for the baseline, L05h05 and L05h10 geometries. These figures represent top views for $\alpha = 9^{\circ}$, $\alpha = 12^{\circ}$ and $\alpha = 15^{\circ}$, with the flow oriented from left to right. Also, to make visualization easier, the colours represent the orientation of the skin-friction in the chord direction, with blue regions corresponding to reversed flow associated with boundary layer separation. For the baseline wing, the flow is approximately two-dimensional in the range of angles of attack considered in figure 6.5, with an increase in the angle of attack leading to the flow separating closer to the leading edge. In the other cases, the use of waviness

tends to prevent separation in the regions corresponding to the peaks of the waviness. In the case L05h05, the separation in this regions is only reduced, while in the case with higher amplitude it is completely eliminated, leading to the formation of isolated separation regions behind the troughs. Also, in the case L05h05, there is separation along the entire span for $\alpha = 15^{\circ}$, what is consistent with the increase in C'_l and the recovery of L/D observed at this angle of attack. Therefore, it can be concluded that the low values of C'_l and the reduction in L/D observed in figures 6.3 and 6.4 at moderate values of angles of attack are related to a three-dimensional flow pattern where separation is suppressed in the region behind the peaks of the waviness.

To further illustrate the behaviour discussed above using skin-friction lines, figure 6.6 shows the recirculation regions for the same cases of figure 6.5, with $\alpha = 12^{\circ}$. These figures were obtained considering the regions where the chordwise velocity becomes negative. From this figure, it is clear that the baseline wing has a recirculation region extending the whole span, while as the waviness amplitude is increased, the boundary layer separation becomes restricted to the waviness troughs.



Figure 6.5: Instantaneous skin-friction lines on the wall for the baseline, L05h05 and L05h10 cases. The flow is from left to right and the colours represent the orientation of the skin-friction in the chord direction, with blue corresponding to reversed flow.



Figure 6.6: Recirculation zones for different cases with $\alpha = 12^{\circ}$.

6.3 Surface pressure distribution

A next reasonable step to understand the previous results is to consider the pressure distribution on the wing surface, as it is related to flow separation and to the aerodynamic forces. Figure 6.7a shows the pressure coefficient distribution on the wing surface for the case L05h10 with $\alpha = 12^{\circ}$ along two different cross-sections, one at the waviness peak and the other on the trough, and also the distribution on the baseline wing as a reference. From this plot, it is clear that there is a strong reduction in the leading edge suction for the section on the waviness peak. The section containing the waviness valley also experiences a reduction in the suction peak, although it is much less severe. The reduction in the aerodynamic forces observed earlier can be attributed to this loss of suction on the upper surface of the wing, as becomes clear when we look at the areas inside the $-C_p$ curves. Thus, even though there is a reduction in the separation region, the loss of leading edge suction prevents this from being converted into an increase of lift.

Another key result from the pressure distribution of figure 6.7a is that, as the flow progresses along the chord, there is a tendency for the pressure to equalize along the sections, eliminating spanwise pressure gradients. These gradients are almost entirely eliminated at x = 0.4, and since the suction peak on the trough section occur behind the corresponding point on the peak section, the former has a shorter distance to recover the pressure when compared to the latter. This, combined with the strongest suction in the trough makes this section experience stronger adverse pressure gradients, suggesting why separation is far more likely to occur in the trough sections rather than the peak sections. Figure 6.7b presents the streamwise tangential pressure gradients corresponding to the previous pressure distribution. While the gradients on the trough section are comparable to the baseline case, the peak section faces weaker gradients, confirming the previous argument.

Also, it should be emphasized that, at least for the low Reynolds number considered here, the fact that the valley section has a shorter distance to recover the pressure is related to the elimination of spanwise pressure gradients, and not to a physical restriction imposed by a shorter chord as has been proposed in the literature, for example in van Nierop et al. (2008). In fact, in the simulations corresponding to figure 6.7 the chord is constant along the span, and therefore there is no chord length effect present.

6.4 Physical mechanism

The previous sections presented results showing that for moderate angles of attack the presence of waviness in the wing leads to a reduction of L/D, which is associated with a flow regime where separation is restricted to the regions behind the waviness valleys. Also, this can be explained by considering the pressure distribution around the wing surface. Therefore, the next step to obtain a better understanding of the flow is to explain how this pressure distribution is formed. This section addresses this issue, proposing a physical description which can justify the observed behaviour.

The main point to understanding how the waviness affects the flow at this Re is to note that, by deforming the wing, spanwise pressure gradients are created. To illustrate this, figure 6.8a shows the spanwise pressure gradients contours at z = 0.125, which is a plane halfway between a peak and a valley of the waviness, for the case L05h10 with $\alpha = 12^{\circ}$. The closest peak is at z = 0.0 and the closest valley at z = 0.25, and thus a positive gradient is oriented from peak to valley and a negative gradient from valley to peak. Close to the leading edge, where x < 0, the flow has already reached the wing in the peak of the waviness, but not in the trough. Therefore, the pressure gradient is negative on the lower surface of the wing (due to the stagnation point) and positive on the upper surface of the wing (due to the suction peak). This pressure gradient accelerates the flow in the span direction, generating a spanwise flow moving away from the protuberance on the lower surface and towards it on the upper surface, as shown in figure 6.8b. This behaviour can be interpreted as an attempt by the flow to move around the protuberance, as illustrated in figure 6.9, which shows contours of w at a plane x = -0.03, which is close to the location of the leading edge in the peak of the waviness at x = -0.05. The effect of the flow towards the suction peak is to increase the pressure at that point, causing the reduction of $-C_p$ at this section discussed in figure 6.7. Also, this spanwise flow around the leading edge protuberance is somewhat compatible with the behaviour of the streamlines obtained by Skillen et al. (2015) at a higher value of Re. However, although they observed the flow being deflected by the lower surface of the leading edge, their results do not show any flow towards the suction peak in the upper surface. They suggest that this deflection led to



Figure 6.7: Pressure coefficient and streamwise tangential pressure gradient on the wing surface for the case L05h10, with $\alpha = 12^{\circ}$. The two-dimensional baseline result at the same angle of attack is also presented as a reference.





Figure 6.8: Contours of spanwise pressure gradient and velocity at z = 0.125. The z direction points out of the page. Case L05h10, with $\alpha = 12^{\circ}$.

the acceleration of the flow behind the troughs, generating an improved suction peak. For the conditions considered here, the suction peak in these regions is still weaker than the baseline case, and therefore it seems that this mechanism is not very efficient.

The previous argument can also be interpreted in a more intuitive manner. In the peak section, by extending the leading edge towards the incoming flow, we hinder the wing's ability to force the flow to accelerate around the leading edge. Although this is not optimal in the sense of generating the maximum possible lift, it prevents the flow from separating.

If we consider now the flow at downstream positions, we note that as the flow reaches the rest of the wing, at x > 0, the spanwise pressure gradients are reversed. The spanwise flow then acts to eliminate the spanwise pressure gradients, leading to the behaviour for the surface pressure described in the previous session. Clearly, the existence of gradients of the w velocity is only possible in the presence of streamwise vortices. Figure 6.10 presents contours of w and of streamwise vorticity at x = 0.2, showing how the existence of spanwise velocity leads to the formation of counter-rotating vortices. However, these vortices are lifted away from the wing and then are quickly dissipated. Therefore, they are not expected to have a significant impact on the observed aerodynamic behaviour, being rather simply a consequence of the spanwise flow.



Figure 6.9: Contours of spanwise velocity (colours) at x = -0.03 and pressure on the wing surface (grayscale). Case L05h10, with $\alpha = 12^{\circ}$.

In conclusion, we have used simple arguments based on the spanwise flow to explain how the pressure distribution is affected by the use of waviness. The pressure distribution, in turn, is responsible for the separation patterns and the loss of lift caused by the waviness.

6.5 Effect of waviness with different shapes

In order to isolate the effect of deforming the leading edge or the trailing edge of the wing, further simulations were performed for these cases. This was accomplished by modulating the wing deformation in the chord direction, using transformations of the form:

$$x = \bar{x} - f(\bar{x})\frac{h}{2}\cos\left(\frac{2\pi}{\lambda}\bar{z}\right)$$
(6.3)

When deformations were desired only on the leading edge, a function $f(\bar{x}) = \frac{\tanh(\kappa(1-\bar{x}))}{\tanh(\kappa)}$ was employed, while the function $f(\bar{x}) = \frac{\tanh(\kappa\bar{x})}{\tanh(\kappa)}$ was used to deform only the trailing edge, with κ being a constant which here assumed the value $\kappa = 1$. Although this choice of functions leads to the desired geometrical properties, it is clearly not unique; indeed, using linear functions could make more sense in terms of affecting the original profile only through a scaling. However, the hyperbolic tangent functions were chosen because they limit the deformations of the domain in the farfield, while the difference in the profile when compared to the linear deformation is minimal. Also, we note that this type of functions allows for a wide range of modifications to be studied, like having out-of-phase deformations in the leading and trailing edges. However, the focus here was restricted to deforming one of them, while leaving the other unmodified. We should also note that in these cases, the NACA0012 profile serves only as the base for defining the wing, with particular sections having deformed forms of this profile, always with a thickness of 12% of the mean chord.



Figure 6.10: Contours of spanwise velocity and streamwise vorticity(colours) at x = 0.2 and pressure on the wing surface (grayscale). Case L05h10, with $\alpha = 12^{\circ}$.



Figure 6.11: Comparison of results for L/D obtained using waviness only on leading edge or trailing edge with the previous results.

Four geometric configurations were considered. They consist of the combinations of wavelength and amplitude corresponding to the cases L05h10 and L10h10, each one of them in the two scenarios where only the leading edge is deformed (denoted by the suffix LE) and where only the trailing edge is deformed (denoted by the suffix TE).

Figure 6.11 presents the lift to drag ratio for these cases, together with the corresponding case of the previous sections and with the baseline. From these results, it is clear that the main determinant of the changes in the aerodynamic performance is the shape of the leading edge. This is reinforced by comparing the recirculation regions presented in figure 6.12 with those of figure 6.6. Therefore, we can conclude that using a uniform deformation along the span (as was the case in the previous sections) is a good approximation for deforming only the leading edge, which is the situation usually encountered in the literature. It is also clear from these results that for the conditions considered here, deforming the trailing edge has little effect on the aerodynamic performance of the wing, contrasting with the significant drag reduction observed in the literature for blunt trailing edges. Considering that the use of serrated trailing edges is a possible way of reducing noise emissions (Howe, 1991), this suggests that both mechanisms can be combined, with the leading edge waviness designed to affect the aerodynamic performance of the wing, and the serrated trailing edge designed to obtain noise reductions.

In the previous results, all sections of the wing had a thickness of 12% of the mean chord, and therefore the profile changes along the chord as the relative thickness is changed. Another approach encountered in the literature is to maintain the same profile throughout the span. The difference between these two geometries was tested for the case L05h05_LE, with the later being obtained by keeping the transformation of equation (6.3) and modifying the *y* coordinate by

$$y = \bar{y} + \frac{\tanh(\kappa \bar{y})}{\tanh(\kappa)} \frac{h}{2} \cos\left(\frac{2\pi}{\lambda}\bar{z}\right)$$
(6.4)



Figure 6.12: Recirculation zones for cases with waviness only on leading edge or on trailing edge, with $\alpha = 12^{\circ}$.



Figure 6.13: Comparison between lift to drag ratio of case L05h05_LE keeping a constant thickness with the same case keeping the profile constant.

In this case, the constant was reduced to $\kappa = 0.5$ in order to reduce the distortions in the original profile.

Figure 6.13 compares the lift to drag ratio in these two configurations. The results are very similar, and therefore we conclude that this distinction between a constant thickness wing and a constant profile wing is not important in this case, although at higher Reynolds numbers the scales on which these geometries differ might become more relevant. Also, since the variations in thickness depend on the waviness amplitude, it is also possible that differences would be observed between the two configurations for larger values of h.

7. Numerical simulations with $\mathbf{Re} = \mathbf{10},000$

This chapter presents results for simulations performed at Re = 10,000. Following the naming convention of table 6.1, the geometries considered were the baseline wing and three different wavy geometries: L025h05, L05h05 and L05h10. Each case was analysed at four angles of attack: 6°, 12°, 15° and 18°, and an additional angle of attack $\alpha = 9^{\circ}$ was considered for the baseline wing.

Based on the tests presented in appendix B, the simulations consisted of 140 time-units, with the last 100 units being used to compute time-averages. The spatial discretisation consisted of a mesh with 1041 quadrilateral elements in the xy plane employing 10^{th} order polynomials, and 128 degrees of freedom in the Fourier expansion of the span direction. This corresponds to approximately 16 million local degrees of freedom. The spectral-vanishing viscosity presented in section 5.1 was used to stabilize the simulations, with parameters $P_{cut} = 0.5P$ and $\epsilon = 0.1$. These parameters were chosen to obtain a stable simulation without using an excessive amount of artificial dissipation.

In order to maintain a systematic approach to the analysis of the results, the presentation of this chapter follows a similar structure to that of chapter 6, with each section describing one aspect of the flow.

7.1 Aerodynamic forces

Figure 7.1 presents the time-averaged force coefficients for all the simulations, including the lift coefficient, the drag coefficient, and the lift-to-drag ratio. Note that for $\alpha = 6^{\circ}$, the waviness leads to a significant increase in the lift, approaching the theoretical 2π slope, without significant changes to the drag. On the other hand, for higher angles of attack the waviness leads to a decrease in both drag and lift. This behaviour at higher angles of attack is similar to what was observed with Re = 1,000. However, in this case the combination of these effects does not result in a reduction of the lift-to-drag ratio.

The changes in the aerodynamic forces are not restricted to their average value, but also to the qualitative characteristics of their time evolution. Figure 7.2 illustrate this by comparing the results of the lift coefficient of the baseline and of the L05h10 geometries



Figure 7.1: Time-averaged force coefficients for simulations with Re = 10,000.



Figure 7.2: Comparison between time series and spectra of the lift coefficient of the L05h10 geometry with the baseline wing, for simulations with Re = 10,000.



Figure 7.3: Comparison between spectra of the lift coefficient of different geometries, for $\alpha = 18^{\circ}$ and Re = 10,000.

for different angles of attack, both in terms of time-series and of frequency contents, represented by the power spectral density (PSD). For $\alpha = 6^{\circ}$, the lift coefficient of the baseline wing has a clear periodic behaviour, with a distinct frequency close to St = 2, where the Strouhal number is defined as $St = \frac{fc}{U_{\infty}}$. On the other hand, at this same angle of attack the lift coefficient of the wavy wing has a broad band spectrum, with no clear periodicity. In other words, the behaviour of the baseline wing is typical of laminar flows while that of the wavy wing is closer to what is expected of turbulent flows, suggesting that the waviness induces the transition of the flow to a more turbulent state, as will be confirmed by flow visualizations later in this chapter. As the angle of attack is increased, the baseline wing also moves towards a broad band spectrum, although a distinct peak is still present, even though it becomes less sharp and at a lower frequency. Also, comparing the results of the time-series in a qualitative sense, we note that by suppressing the tonal component of the lift force, the waviness leads to less severe lift fluctuations. The other wavy geometries considered here present a similar behaviour, with the exception of the angle of attack $\alpha = 18^{\circ}$. The spectra of all cases for this angle of attack are presented in figure 7.3. At this high angle of attack the wing's behaviour is closer to that of a bluff body, and the leading edge waviness is not so effective at preventing the tonal component of the spectrum, specially in the case L05h05. Figures 7.2 and 7.3 also show a change in the slope of the spectrum around St = 10 for the wavy wings at the largest angles on attack. After this change the slope of the spectrum is close to -5/3, suggesting this could be related to turbulence in the flow.

7.2 Separation and recirculation zones

Following the same idea as in chapter 6, the effect of the waviness on the flow will now be analysed in detail, in order to explain the results described in the previous section for the aerodynamic forces. Once again, a good starting point is to consider the separation of the flow and some general aspects of the flow topology.

Figure 7.4 shows the time-averaged skin-friction lines on the wall of the wing for several cases. These are similar to figure 6.5, with the flow directed from left to right and regions in blue representing recirculating flow. Note that the time interval employed in the averaging procedure (100 time units) was chosen to obtain convergence to the mean value in the forces, and therefore it is possible that the local flow is not completely converged to the mean flow. In any case, the time averaging removes all but the lowest frequencies of the flow, and therefore it accomplishes the goal of filtering turbulence and vortex shedding fluctuations. As expected, the time-averaged flow for the baseline wing is almost two-dimensional, with the exception of small three-dimensional artefacts which are most likely caused by the already noted limited extent of the time-averaging. Also, both angles of attack present extensive separation regions. Although this figure suggests the presence of reattachment for the higher angle of attack, this is actually a secondary flow inside the separation region. The wavy wings present a behaviour similar to that of Re = 1,000, in which the flow tends to remain attached behind the peaks, with more extensive separation behind the valleys. Another interesting aspect of these results is that in some cases the timeaveraged flow around the wavy wings is not symmetric between successive wavelengths, indicating a sub-harmonic behaviour.

Figures 7.5 and 7.6 respectively present instantaneous and time-averaged recirculation regions for all simulations with $\alpha = 6^{\circ}$ and $\alpha = 12^{\circ}$. The clear similarity between both results for the flow around the baseline wing with $\alpha = 6^{\circ}$ is further evidence of the laminar character of the flow in this case. Although in varying degrees, in all other cases the instantaneous flow presents a break down into small-scale structures typical of turbulence. From this figures it is also clear that the asymmetry in the wavy wings is characterized by recirculation regions of significant different extent between successive valleys, specially in the case L05h10, and L025h05 with $\alpha = 12^{\circ}$. It can also be noted that although the general tendency observed for Re = 1,000 of separation being favoured behind the troughs remains, the separation patterns are significantly different in this case. While at Re = 1,000 the flow exhibited distinct separation regions behind the troughs with completely attached flow behind the peaks, for Re = 10,000 at $\alpha = 12^{\circ}$ this pattern is only observed in the upstream half of the wing. After that, the separation extends to the whole span, in some cases preceded by a region with reattachment of the flow.

It is interesting to note that the structures generated by the wavy leading edge are still present even at the higher angle of attack $\alpha = 18^{\circ}$, when the wing is already well into the



Figure 7.4: Time-averaged skin-friction lines on the wall for simulations with Re = 10,000. The flow is from left to right and the colours represent the orientation of the skin-friction in the chord direction, with blue corresponding to reversed flow.



Figure 7.5: Instantaneous recirculation regions for simulations with Re = 10,000.

stall regime. This is illustrated in figure 7.7, which compares the time-average recirculation regions for the baseline and L025h05 geometries at this angle of attack. Even though at this high angle of attack the effect of the waviness is limited, with both configurations presenting separation over most of the suction side of the wing, the effect of the waviness can still be observed near the leading edge. In this region, we note that the flow remains attached for a longer extent behind the peaks. Also, a sub-harmonic behaviour can be observed for the separation behind the troughs, although this is more subtle than in other configurations.

7.3 Surface pressure distribution

In this section, the impact of the waviness on the pressure coefficient (C_p) on the wing surface will be discussed, always considering time-averaged results. In this case, it is im-



Figure 7.6: Time-averaged recirculation regions for simulations with Re = 10,000.

portant to be careful when referring to the peak or trough section, since the results of section 7.2 show that at this Reynolds number the symmetry between successive wavelengths is broken in some cases. Taking this into account, two consecutive peaks and troughs were analysed for each case. It was observed that discrepancies between different peaks are small in all cases, while the valleys present significantly different behaviour. Therefore, results are always presented for a single peak section, and for two consecutive troughs. The exact station along the *z* direction for each result is always indicated, in order to allow comparisons with the flow visualizations (for example, in figure 7.4 the top corresponds to z = 0 and the bottom to z = 1).

Considering each angle of attack separately, figure 7.8 contains the results for the pressure coefficient distribution for the simulations with $\alpha = 6^{\circ}$. As discussed in the previous paragraph, for each case three sections are considered (two troughs and one peak), and



Figure 7.7: Time-averaged recirculation regions for simulations with $\alpha = 18^{\circ}$ and Re = 10,000.

the result for the baseline geometry is presented as a reference. First, note that in this case the pressure coefficient of the baseline geometry has a thin shape, leading to a low lift coefficient. This behaviour is similar to the one observed for Re = 1,000, for example that of figure 6.7a. Another important point is that on all cases the peak section has a weaker suction peak, which is consistent with the discussion of section 6.4 about the results for Re = 1,000. We also note that downstream of x = 0.2, one of the valleys closely follows the pressure profile of the peak section, with the pressure remaining constant for a long extent before recovering near the trailing edge. On the other hand, the other valley has a different behaviour, with pressure recovery starting further upstream.

The results for $\alpha = 12^{\circ}$ are presented in figure 7.9. In the baseline case, the C_p curve is much broader than in the lower angle of attack, most likely as a consequence of the transition to a turbulent state. The constant pressure on the suction side is compatible with the fully separated flow observed in figure 7.6. The L05h05 geometry is not capable of significantly affecting the main characteristic of the pressure distribution, due to the low amplitude of the waviness compared with the wavelength. In this case, the only effect is a weakening of the suction peak in all sections, leading to a much lower lift coefficient. In the other two cases, the effects observed for $\alpha = 6^{\circ}$ are intensified. In particular, we note that the trough with the smaller recirculation region (z = 0.25 for L05h10 and z = 0.375 for L025h05) develops a clear low pressure plateau before beginning a pressure recovery. This behaviour is typical of a flow regime containing a laminar separation bubble (Galbraith and Visbal, 2008), which is characterized by separation of the laminar boundary layer followed by transition in the shear layer and turbulent reattachment.

For $\alpha = 18^{\circ}$, all cases present similar results in terms of C_p distribution. Therefore, only the result for L05h10 will be discussed, in figure 7.10. At this high angle of attack, the wing is stalled in all scenarios. The pressure on the suction side is almost constant in all geometries, although the particular level of pressure varies from case to case, affecting the value of the lift coefficient. To be specific, the L05h05 case has little impact on the flow due to the small amplitude, and therefore the pressure distribution and lift coefficient are close to the baseline; on the other hand, L05h10 and L025h05 present higher pressures on the suction side (as shown in figure 7.10), and therefore have lower lift forces.



Figure 7.8: Time-averaged surface pressure coefficient distribution for simulations with $\alpha = 6^{\circ}$ and Re = 10,000.





Figure 7.9: Time-averaged surface pressure coefficient distribution for simulations with $\alpha = 12^{\circ}$ and Re = 10,000.



Figure 7.10: Time-averaged surface pressure coefficient distribution for case L05h10 with $\alpha = 18^{\circ}$ and Re = 10,000.

7.4 Velocity profiles

In order to obtain further information regarding the flow features causing the different behaviours observed for the pressure distribution, this section analyses chordwise velocity profiles along the y direction. The discussion will focus on comparing the baseline with the L025h05 geometry, since this is the waviness which leads to a more pronounced impact on the flow.

Figure 7.11 shows the results for $\alpha = 6^{\circ}$, including the same cross sections as in the C_p results of figure 7.8. Both instantaneous and time-averaged profiles are presented, with a spacing of 0.1 in the chord direction between profiles, starting at x = 0.1 (measured from the leading edge). As expected from the previous discussion, at this angle of attack the flow around the baseline wing separates without any reattachment, remaining in a laminar state. This last point can be observed by comparing the instantaneous and time-averaged results. As for the wavy wing, as usual we observe the flow separating close to the leading edge in the troughs, while in the peak section the flow remains attached for a longer extent. Also, as expected the instantaneous profiles show the turbulent characteristic of this case. Although the previous observations were already known from the other results, this figure provides new information when we consider the patterns that can be observed for the reattachment of the flow. While for z = 0.125 and z = 0.250 the reattachment can be observed only around x = 0.8, in the z = 0.375 section it occurs much earlier, at x = 0.5. This may well explain the different C_p curves, with the earlier pressure recovery for z = 0.375 being related to the reattachment of the flow.

Figure 7.12 shows the results for $\alpha = 12^{\circ}$. Similarly to the $\alpha = 6^{\circ}$ case, the baseline wing presents a fully separated flow. However, the instantaneous flow now contains fluctuations, specially after x = 0.5. For the wavy wing case, once again the flow reattaches earlier for z = 0.375, at x = 0.3. However, from x = 0.6 onwards the time-averaged profiles are very similar in all sections. Comparing with figure 7.9c, we note that after this location the pressure is also the same in all sections. Therefore, in this case there is a equalization of the flow along the spanwise direction, much like what was observed for Re = 1,000.

7.5 Discussion

The results from the previous sections show that the impact of the waviness on the aerodynamic forces is quite different when comparing the lower angle of attack $\alpha = 6^{\circ}$ with the other cases. For $\alpha = 6^{\circ}$, the waviness causes a large increase in the lift coefficient. This was observed to be related to a distinction between laminar and turbulent flows. At low Reynolds numbers (for example, the Re = 1,000 cases of chapter 6) the flow is laminar and the lift coefficient is significantly lower than the high Reynolds number value. By triggering the transition to turbulence, the waviness prevents the performance degradation caused by



Figure 7.11: Comparison between u velocity profiles along the y direction for baseline and L025h05 geometries with $\alpha = 6^{\circ}$ and Re = 10,000.

the laminarization of the flow as the Reynolds number decreases. For this angle of attack, this appears to be the dominant effect of the waviness. Therefore, we can conclude that the use of spanwise waviness can offer the advantage of delaying the loss of lift as the Reynolds number is decreased.

Considering now the higher values of angle of attack, some aspects of the effect of the waviness are similar to what was obtained at Re = 1,000, illustrating how the flow structures observed at that case can persist at higher Reynolds numbers. First, at this Re the waviness also cause reductions in the lift and drag coefficients, although the combined effect in this case leads to a stable lift-to-drag ratio, compared to the reductions of L/D for the lower Re. The separation patterns also follow the trend of separation behind the troughs with attached flow behind peaks, related to a lower suction peak in the latter section. The



Figure 7.12: Comparison between u velocity profiles along the y direction for baseline and L025h05 geometries with $\alpha = 12^{\circ}$ and Re = 10,000.

occurrence of this reduced suction peak can be explained by the same argument based on spanwise flow used in chapter 6, which can be summarized by the waviness deflecting the flow along the spanwise direction, hindering the ability of the leading edge to force the flow to accelerate as it move around the leading edge, and therefore weakening the suction peak. This is supported by figure 7.13, which presents spanwise velocity contours at a xyplane between a peak and a trough. This figure is the Re = 10,000 equivalent of figure 6.8b. These two figures exhibit similar patterns for the spanwise velocity in the leading edge region, further supporting that similar flow features are responsible for the reduction in the suction peak in both cases.

One aspect of the flow that is significantly different from the lower Reynolds number results is that now the flow presents a sub-harmonic behaviour, with a loss of symmetry be-



Figure 7.13: Contours of time-averaged spanwise velocity at z = 0.125. Case L05h10, with $\alpha = 12^{\circ}$ and Re = 10,000.

tween adjacent wavelengths of the waviness. This asymmetry is more notable in the trough sections, and is characterized by different separation patterns and pressure distributions. We also note that no switching between the patterns of the two wavelengths of the wing was observed, although this may occur over periods longer than the 100 time units considered here. It was observed that the pressure distribution in some of the trough sections present a low pressure plateau, followed by a pressure recovery region. This is a typical behaviour associated with laminar-separation bubbles. This observation is consistent with the velocity profiles from figures 7.11 and 7.12, which show that in the sections presenting this behaviour there is a reattachment of the flow close to where the pressure recovery begins. To reinforce the distinction between different xy sections, figure 7.14 shows time-averaged streamlines projected in the xy plane for different sections, considering the geometries L05h10 and L025h05 with $\alpha = 12^{\circ}$. The colours in the figure indicate the orientation of the flow in the chordwise direction. This figure illustrates some interesting characteristics of the flow: in the peak section the flow remains attached for longer and then separates with no reattachment (or with a small reattached region); the trough section with pressure distribution similar to the peak section has an extensive separation region; in the other trough section the flow separates near the leading edge but then reattaches, leading to the laminar separation bubble behaviour observed in the pressure distribution.



(e) L025h05, z = 0.375 (valley)



Figure 7.14: Time-averaged streamlines projected in the xy planes at different sections for the L025h05 and L05h10 geometries with $\alpha = 12^{\circ}$ and Re = 10,000. The colours represent the orientation of the chordwise velocity.
8. Numerical simulations with $\mathbf{Re} = 50,000$

This chapter discusses the results of a series of simulations at Re = 50,000. Simulations were carried out for the baseline and L05h10 geometries, in both cases for $\alpha = 6^{\circ}$ and $\alpha = 15^{\circ}$. As was mentioned previously, being able to perform highly resolved simulations at this Reynolds number is of significant importance, since it is close to what might be expected in some applications, and it also allows comparisons to the experiments presented in Paula (2016).

All simulations used a time step of 5×10^{-5} . Due to the reduced number of simulations, it was possible to choose in each case a different duration for the initial transient required to eliminate the influence of the initial conditions, and the time-averaging interval required to obtain converged results. Similarly to the simulations with Re = 10,000, these intervals lead to converged results in the aerodynamic forces as will be discussed later, although the locally time-averaged flow may not be completely converged to the mean flow. The intervals considered in each simulation are presented in table 8.1.

The span of the wing was also chosen specifically for each particular simulation, with the baseline wing using $L_z = 0.5c$, while the wavy wing had $L_z = 1.0c$ for $\alpha = 6^{\circ}$ and $L_z = 2.0c$ for $\alpha = 15^{\circ}$. The reason for the longer span in the last case is that flow visualizations from Paula (2016) for a condition similar to this indicate a subharmonic behaviour with a pattern which is only repeated after 4 wavelengths. For the baseline wing, it would be possible to determine the appropriate spanwise length by calculating two-point correlations along lines aligned with the *z* axis. If the L_z is sufficiently large, the correlation should drop to zero, before rising to 1 again to satisfy the periodic boundary condition. However, this procedure

Case	Initial transient time	Averaging interval	L_z	ε
baseline, $\alpha = 6^{\circ}$	25	25	0.5c	0.1
baseline, $\alpha = 15^{\circ}$	50	25	0.5c	1.0
L05h10, $lpha=6^\circ$	20	15	1.0c	0.1
L05h10, $\alpha = 15^{\circ}$	pprox 60	30	2.0c	1.0

Table 8.1: Parameters used in simulations with Re = 50,000. L_z is the spanwise length and ϵ is the SVV diffusion coefficient.

is not straightforward in the case of the wavy wing, since in the presence of the waviness the span direction is not homogeneous. Therefore, for the baseline wing L_z was chosen considering the results from Skillen et al. (2015), while for the wavy wing choosing L_z based on the visualizations of Paula (2016) was preferred over the correlation method.

In terms of spatial resolution, all simulations were performed in a mesh consisting of 4674 quadrilateral elements in the xy plane, with the adaptive procedure presented in section 5.3 being used to adjust the order between $P_{min} = 2$ and $P_{max} = 11$. The adaptive procedure was only employed during the initial transient phase, with the time-averaging being performed over a fixed distribution of polynomial order. The number of Fourier modes was adjusted to maintain a resolution of 256 degrees of freedom per chord length in the span direction.

Similarly to the previous chapter, the spectral-vanishing viscosity presented in section 5.1 was used to stabilize the solution. All simulations used $P_{cut} = 0.5P$, while the diffusion coefficient was $\epsilon = 0.1$ for $\alpha = 6^{\circ}$ and $\epsilon = 1.0$ for $\alpha = 15^{\circ}$. Once again, these parameters were chosen to stabilize the simulation without adding an excessive amount of artificial dissipation. This set of simulations also used dealiasing techniques to consistently integrate the non-linear terms. The spectral/hp expansion on the xy plane was dealiased using the technique described in Kirby and Sherwin (2006a), while for the Fourier expansion the 3/2 padding rule was used. The reasoning for using dealiasing in addition to SVV is briefly discussed in appendix C using simulations of the Taylor-Green Vortex. The main conclusion from that discussion is that although dealiasing techniques are not always sufficient to stabilize the solution, they can potentially improve the accuracy of the solution by consistently integrating the non-linear terms and by reducing the amount of SVV required in the simulation.

The following sections will present in detail the most important results obtained from these simulations, with each section describing one particular feature of the flow.

8.1 Aerodynamic forces

Figure 8.1 presents the results of aerodynamic forces for all the simulations. The timeaveraged lift coefficient is compared with the experimental results from Paula (2016), which were also performed at Re = 50,000 for a NACA0012 profile. We observe a good agreement for the baseline wing, while the comparison is more difficult for the wavy wing since the parameters of the waviness considered here are not identical to the ones used in the experiments. For $\alpha = 6^{\circ}$, the waviness causes a small reduction in both \overline{C}_l and \overline{C}_d , with the resulting effect being a higher lift-to-drag ratio. For $\alpha = 15^{\circ}$, the wavy wing has a significantly higher lift and a lower drag. In this last case, there is also a reduction in C'_l , similar to what was observed at lower Re.

Considering now the transient behaviour of the flow, figure 8.2 presents the spectra of



Figure 8.1: Results of force coefficients for simulations with Re = 50,000. For the timeaveraged lift coefficient, the solid lines are the experimental results from Paula (2016) and the dashed line is the 2π slope.



Figure 8.2: Spectra of lift coefficient for simulations with Re = 50,000.



Figure 8.3: Moving averages of the lift coefficient for simulations with Re=50,000 and $\alpha=15^{\circ}$.

the lift coefficient for the simulations. At the lower angle of attack, the fluctuations are very small and there are no predominant characteristic frequencies, with the waviness having little effect on the frequency composition. On the other hand, at $\alpha = 15^{\circ}$ the baseline wing presents a distinct peak around a Strouhal number St = 0.7, which is suppressed by the waviness.

Still considering the time evolution of the lift force, some interesting behaviour is observed for $\alpha = 15^{\circ}$. To illustrate this, figure 8.3 presents moving averages of different lengths for the lift coefficient of both geometries at this angle of attack. Considering $T_{ma} = 0.01$, which is almost identical to the instantaneous lift, we notice how both geometries have a significant low frequency behaviour, although for the baseline wing the fluctuations associated with the Strouhal number are much more intense. The second moving average ($T_{ma} = 5$) shows that the low frequency behaviour is very different in both cases. While the baseline wing exhibits a trend of increasing lift which suddenly breaks down, the

wavy wing slowly alternates between low and high lift configurations. Finally, the moving average with $T_{ma} = 25$ indicates that despite these low frequency components, the averaging intervals used for these simulations are sufficient to obtain reasonably converged results for the lift coefficient.

8.2 Separation and recirculation zones

We now consider the separation patterns on the wing, in order to obtain some information on the flow structures which might explain the changes in aerodynamic forces described in the previous section. Figure 8.4 shows the time-averaged skin-friction lines on the surface of the wing with the flow oriented from top to bottom, while figure 8.5 presents the corresponding recirculation regions.

Starting with the baseline wing, the simulation with $\alpha = 6^{\circ}$ clearly shows a relatively long laminar separation bubble, with reattachment around x = 0.54. For $\alpha = 15^{\circ}$, the skin-friction alone would suggest there is a region of reattached flow. However, the recirculation regions clarify that this is actually a secondary flow inside the region of fully separated flow.

Considering now the wavy wing with $\alpha = 6^{\circ}$, several observations can be made. First, the behaviour observed at lower Re where the flow tends to remain attached behind the troughs is still present. However, separation is very limited in this case compared to the lower Re (figure 7.4, for example), with the flow behind peaks remaining completely attached. Another interesting aspect of this result is that both wavelengths show identical results, and therefore this configuration does not exhibit any sub-harmonic behaviour. Finally, we observe that the separation bubbles are much shorter than they were in the baseline wing, suggesting that transition to turbulence occurs earlier. This last point will be confirmed in later sections.

The results for the wavy wing with $\alpha = 15^{\circ}$ show a much more complicated separation pattern. There is a significant distinction between different wavelengths along the span, justifying the need to use four wavelengths in this simulation. We note that the first two wavelengths (starting from z = 0) have a widespread separation with a large separation bubble. On the other hand, the rest of the wing has shorter separation regions concentrated behind the troughs. We can also observe some small separation regions moving diagonally close to the wing surface.

These results can be compared with the oil-visualizations from Paula (2016) for similar conditions, which are reproduced in figure 8.6. In these experiments, the geometry of the wavy wing is slightly different, and the lower angle of attack is $\alpha = 5^{\circ}$ instead of $\alpha = 6^{\circ}$. For the baseline wing and for the wavy wing with the lower angle of attack, there is clearly a good qualitative agreement with the experiments. The results for the wavy wing with $\alpha = 15^{\circ}$ also seem consistent, although in this case it is harder to identify all structures and match the numerical and experimental results.



(d) L05h10, $\alpha=15^\circ$

Figure 8.4: Time-averaged skin-friction lines on the wall for simulations with Re = 50,000. The flow is from top to bottom.



Figure 8.5: Time-averaged recirculation regions for simulations with Re = 50,000.

8.3 Surface pressure distribution

We will now analyse the results for pressure coefficient distribution on the wing surface, which are presented in figure 8.7.

For $\alpha = 6^{\circ}$, the baseline wing presents a typical profile for a flow with a laminar separation bubble, with a suction plateau in the separated region followed by a pressure recovery after the flow transitions to turbulence and reattaches. The transition can be clearly identified by comparing the instantaneous and the time-averaged plots. In the wavy wing, the valley section also presents a laminar separation bubble behaviour, but with separation occurring earlier and with a much shorter bubble. These characteristics of the bubble cause the pressure plateau to have a stronger suction than in the baseline wing. The peak section, on the other hand, presents a completely distinct characteristic. In this case, the suction peak is much weaker than in the baseline wing, as was also the case for lower Reynolds numbers. The pressure then undergoes a slow recovery, without any strong adverse gradients, therefore allowing the flow to remain attached. Since there is no separation, the flow in this section does not experience transition to turbulence, which at this Reynolds numbers is associated with a Kelvin-Helmholtz instability in the shear layer. Therefore, behind the trough section there are only fluctuations near the trailing edge, which as will be discussed later are associated to the turbulence generated in the shear layer of the peak section spreading to the whole span.

The situation is quite different for the higher angle of attack. The baseline wing presents



(a) baseline, $\alpha = 5^{\circ}$



(b) baseline, $\alpha = 15^{\circ}$



(c) L04h11, $\alpha = 5^{\circ}$



(d) L04h11, $\alpha = 15^{\circ}$

Figure 8.6: Oil visualization from experiments of Paula (2016) with Re = 50,000. The flow is from top to bottom.



Figure 8.7: Surface pressure coefficient distribution for simulations with Re = 50,000.



Figure 8.8: Location of the probes on the xy plane.

	Probe									
	1	2	3	4	5	6	7	8		
\overline{x}	0.01	0.10	0.30	0.50	0.70	0.85	1.03	1.20		
y	0.025	0.090	0.065	0.120	0.200	0.080	0.000	0.100		

Table 8.2: Coordinates of the probes.

a C_p consistent with the post-stall regime, and the sections of the wavy wing in the first half of the span ($0 \le z < 1$) follow a similar pattern. On the second half of the span, it is possible to observe stronger suction, with plateaus that are similar to those present in the pre-stall regime. However, the distinctions between the two halves of the span are much clearer than the distinctions between peaks and troughs, contrasting with what is observed in most of the other simulations.

8.4 Spectra of velocity fluctuations

In order to demonstrate how these simulations are capable of resolving the turbulent structures, and also obtain a better understanding of some aspects of the flow, this section presents the spectra of chordwise velocity fluctuations in specific points of the domain. Figure 8.8 shows the position on the xy plane of all probes that were monitored during the simulations, although results will only be presented for some of these probes. Their exact coordinates in the xy plane are detailed in table 8.2. For the baseline wing the third coordinate was z = 0.25, while for the wavy wing one probe was used for the trough section (z = 0.25) and another for the peak (z = 0.5). Note that for the wavy wing these correspond to coordinates in the computational domain.

Figure 8.9 presents the power spectral density of the fluctuations of *u*-velocity at different probes for the baseline wing, with the dashed line corresponding to the $-\frac{5}{3}$ slope. In all of these situations, the flow is turbulent, and we observe that the simulation captures the inertial range of turbulence. For probe 8 with $\alpha = 15^{\circ}$, there is also a peak which is likely related to vortex shedding.

Figure 8.10 contains the same type of results, this time for the wavy wing with $\alpha = 6^{\circ}$. In the valley section, once again the flow is turbulent and resolved well into the inertial range.



Figure 8.9: Spectra of *u*-velocity fluctuations for simulations of the baseline wing with Re = 50,000. The dashed line corresponds to the $-\frac{5}{3}$ slope.



Figure 8.10: Spectra of *u*-velocity fluctuations for simulations of the wavy wing with $\alpha = 6^{\circ}$ and Re = 50,000. The dashed line corresponds to the $-\frac{5}{3}$ slope.

In the peak section, however, the flow is laminar with much lower fluctuations, leading to a spectrum almost 5 orders of magnitude lower than the other cases.

These energy spectra are a good indication that these simulations were performed with a sufficiently high resolution. Combined with the good agreement observed between numerical and experimental results, they confirm the high accuracy of the simulations, validating our results.

8.5 Discussion

The previous sections presented several characteristics of the flow for the simulations performed at Re = 50,000. Now, we briefly provide a general description of the flow, with some additional results used where necessary.

For the simulations with $\alpha = 6^{\circ}$, there are two main effects caused by the waviness. The first is the reattachment of the flow behind the waviness peaks, leading to a separation pattern with distinct separation regions behind the troughs. This is similar to what was observed at lower Re, and once again can be attributed to a weaker suction peak behind the waviness peaks, which in turn leads to lower adverse pressure gradients. In order to understand if the cause for this lower suction peak is similar to what was described in earlier chapters, figure 8.11 shows contours of spanwise velocity for a section between a peak and a trough. Near the leading edge, we observe the same pattern which were seen for Re = 1,000 and Re = 10,000 in figures 6.8 and 7.13. Therefore, it is reasonable to assume that the same mechanism is also valid here. This is characterized by the waviness allowing the flow to move around the leading edge in this spanwise direction, hindering the ability of the wing to generate a strong suction peak.

The second effect of the waviness at $\alpha = 6^{\circ}$ is a shortening of the laminar separation



Figure 8.11: Time-averaged contours of spanwise velocity at z = 0.125 for L05h10 wing with $\alpha = 6^{\circ}$ and Re = 50,000.

bubble behind the waviness valleys. Due to the nature of the separation bubble, this is expected to be related to an earlier transition to turbulence. To illustrate this, figure 8.12 shows contours of spanwise vorticity and of turbulence kinetic energy (TKE) for the baseline wing and for different sections of the wavy wing. Note that for the wavy wing the TKE was calculated based on fluctuations of the velocity on the transformed computational domain. Although this could be slightly different from the exact value in the physical domain, qualitatively they are the same. This figure confirms that the earlier reattachment is related to an earlier transition to turbulence. Although the causes for this are not clear, they could be related to the shape of the shear layer, which is curved for the wavy wing and straight for the baseline wing. It is possible that some characteristics associated to this curved shear layer make it less stable, allowing it to break down into turbulent structures more easily. In this figure, we also observe some turbulence in the peak section. This is unexpected from a purely sectional point of view, since from figure 8.5 we know that the flow remains attached on this section and the Reynolds number is too low for transition to occur in the boundary layer. However, considering the iso-surface of TKE of figure 8.13, we note that what actually happens is that after the flow transitions in the shear layer behind the trough, the turbulence spreads to the rest of the span. This confirms that there is no transition in the attached boundary layer, and illustrate the importance of considering three-dimensional effects in order to obtain a complete understanding of this flow.

Moving on to $\alpha = 15^{\circ}$, one distinct characteristic of the flow for both the baseline and the wavy wing is the low frequency component observed in the lift coefficient. Although a conclusive explanation for this phenomenon requires further investigation, we will now present some results which indicate possible causes for this behaviour.

Considering first the baseline wing, we observed earlier that the flow averaged over the entire 25 time units interval considered here is completely separated. However, this may not be the case for shorter periods of time. Figure 8.14 presents the recirculation region of the flow filtered using an exponential moving average with time constant $\tau = 1.5$ for different final times. This filter was used to remove the high-frequency turbulent fluctuations, while



(c) L05h10, peak

Figure 8.12: Slices with contours of instantaneous spanwise vorticity (on the left) and of turbulence kinetic energy (on the right) for simulations with $\alpha = 6^{\circ}$.



Figure 8.13: Iso-surface of TKE = 0.015 for wavy wing with $\alpha = 6^{\circ}$ and Re = 50,000.

maintaining the distinctions which might explain the low-frequency behaviour. It is possible to see that the recirculation region of the high lift configuration (t = 37.3) is ruptured near the centre of the span, while this does not happen in the low lift case. This seems to be caused by an interaction with the secondary flow inside the separation bubble, which is present to different extents in both cases, as can be seen from the side view. Therefore, we can hypothesize that the lift initially increases as this secondary flow region grows and disturbs the separation bubble. This structure then bursts, leading to the sharp drop in the C_l observed in figure 8.3 around t = 40.

Following an analogous idea to that of the previous paragraph, figure 8.15 presents recirculation regions of the flow averaged over different intervals for the wavy wing with $\alpha = 15^{\circ}$. This figure shows a significant distinction between the two intervals in the second half of the span (1 < z < 2). In this region, separation is much more extensive in the lower C_l configuration, while in the higher C_l configuration there is flow reattachment in most of this region. Therefore, we associate the changes in separation patterns in some portions of the span as being responsible for the low-frequency variation of the lift, although the mechanisms leading to these changes are not clear yet.

As a final note, it is important to discuss that for $\alpha = 15^{\circ}$ the use of the waviness did not decrease the suction peak in the peak sections of the waviness, as was the case for most of the simulations of the present work. As mentioned in chapter 6, the spanwise



Figure 8.14: Recirculation region filtered using an exponential moving average with different final times. Simulation with $\alpha = 15^{\circ}$ and Re = 50,000.

flow permitted by the waviness creates a tendency to reduce the suction peaks in these sections. In order to further our understanding on this point, we can consider, as a first approximation, that this effect is similar to what might be expected from a reduction in the angle of attack. In the pre-stall regime, or for lower Re where the lift coefficient curve is monotonic, this invariably leads to an actual decrease in the suction peak, accompanied by a reduction in C_l . However, in the post-stall regime, it is possible that by attempting to reduce the suction peak, this mechanism causes the flow to reattach, as was observed for parts of the wavy wing with $\alpha = 15^{\circ}$. Therefore, despite the tendency for the suction peak to be reduced, the overall effect in the post-stall regime is hard to quantify since there are competing effects at play.



(b) $83.75 \le t \le 91.25$

Figure 8.15: Recirculation regions for flow obtained using different time-averaging intervals. Wavy wing with $\alpha = 15^{\circ}$ and Re = 50,000.

Part IV

Conclusions and future work

9. Conclusions

This thesis presented an extensive numerical study of the effect of spanwise waviness on the flow around infinite wings at low Reynolds numbers. The simulations employed the spectral/hp method available through nektar++, including a number of improvements developed during the course of this research. Therefore, the conclusions of the present work can roughly be divided into those pertaining to the numerical methods and those drawn from the simulations.

9.1 Numerical methods

Considering first the numerical methods, the main contribution of this work is the development of new techniques for introducing coordinate transformations into the velocitycorrection scheme, which is the time-integration method commonly employed in the solution of the incompressible Navier-Stokes equations using spectral/hp methods. Although the spectral/hp discretisation offers geometric flexibility, in this case using a coordinate transformation was desirable because it allows us to employ a guasi-3D approach, leading to a more efficient solution of the governing equations. The new methods presented here can be understood as generalizations of two approaches available in the literature: an explicit method restricted to constant Jacobian transformation (Darekar and Sherwin, 2001b); and a semi-implicit method employed previously in pseudo-spectral methods, for which the appropriate high-order pressure boundary conditions were not readily available (Carlson et al., 1995). Based on these methods, an explicit and a semi-implicit method were obtained, both of them suited for general transformations (including those with a non-constant Jacobian) and compatible with the framework of the velocity-correction scheme. Through numerical examples, it was shown that these methods preserve the exponential convergence of the underlying discretisation, and that they are flexible enough to consider very general time-dependent transformations. The computational cost of the semi-implicit scheme is considerably larger than that of the explicit scheme, but the former is more stable than the latter. Therefore, the explicit scheme is best suited for less energetic transformations, while the semi-implicit one should be considered in situations where the explicit scheme is not stable.

Although the new coordinate transformation technique was sufficient to tackle the very

low Reynolds number Re = 1,000 case, other considerations had to be made for higher Re. First, the parallelism strategy had to be improved, so that the full potential of modern supercomputers could be exploited. Previously, two main parallelism options were available through Nektar++ for these quasi-3D simulations. The first consisted of partitioning only the homogeneous direction, using a direct solver in the implicit portion of the solution. This is very efficient, but the maximum number of processes that can be used is relatively small, preventing us from performing massively parallel simulations. The second approach consisted in partitioning both the homogeneous direction and the xy plane, using an iterative conjugate gradient method in the implicit solves. This approach scales to more processes, but is significantly less efficient, therefore increasing the overall computational speed. The strategy we followed here consisted in using the hybrid parallelism combined with direct solves for the linear system. It is expected that this approach will perform poorly with a large number of partitions in xy, but we showed that for a moderate number of partitions the direct solver still performs very well. Therefore, this extended the number of processes for which the highly efficient direct solver strategy is applicable, allowing us to leverage the available computational resources in order to significantly speed up the simulations. For example, in the simulations of a wavy wing with Re = 10,000, 9 partitions were employed in xy plane, maintaining an efficiency close to 90% when compared to a single partition in the xy plane. This allowed a simulation that would take almost two months using only the homogeneous parallelism to be performed in under one week, illustrating the importance of these considerations to the feasibility of the present study.

In order to further reduce the computational cost of the simulations, an adaptive polynomial order procedure was also developed. Although changing the polynomial order in the middle of the simulation is expensive, due to the requirement to reassemble and decompose the global matrices, an efficient approach was obtained by performing this refinement only after a few thousand time steps, therefore diluting the refinement cost. It was shown that in a practical situation the refinement process consumes around 5% of the time. This is more than offset by the gains from using a better distribution of polynomial order, which in that case corresponded to a reduction of 35% in the computational cost when compared with the simulation with a constant order equal to the maximum order of the adaptive procedure.

Summing up, the numerical methods presented in this thesis led to significant reductions in the computational cost and in the time required to run a simulation. Although in principle it was possible to run the simulations without any of these techniques, it is expected that the computational cost in this case would be many times larger. This would likely make some of the simulations of the present study impracticable (specially the ones at higher *Re*), demonstrating the importance of the developments to the numerical methods presented here.

9.2 Effect of spanwise waviness on the flow

The present study considered infinite wings with a NACA0012 profile. Three sets of simulations were performed, with the Reynolds number being progressively increased between each one of them.

First, simulations with Re = 1,000 were carried out. Due to the low computational cost of these simulations, it was possible to perform an extensive parametric study, considering several geometries with different combinations of waviness wavelength and amplitude, and also several angles of attack for each geometry. This first set of simulations also facilitated the task of analysing the flow in detail, since in this case the flow is laminar and the large scale structures are easily observed. The results show that the waviness leads to a reduction in the lift-to-drag ratio, which is the consequence of the combined effect of both lower lift and lower drag forces. There is also a significant reduction in the lift coefficient fluctuations, in some cases with a complete suppression of these fluctuations. It was observed that this behaviour is associated with a flow regime where the flow remains attached behind the peaks of the waviness, while separation persists behind the troughs. This separation pattern is caused by a weakening of the suction peaks on the section corresponding to the waviness peaks, which leads to weaker adverse pressure gradients, therefore allowing the flow to remain attached in these sections. It was proposed that the main physical mechanism responsible for the observed behaviours is related to the spanwise flow near the leading edge. By extending the leading edge towards the flow in the waviness peak region, we allow the flow to be deflected in the spanwise direction, instead of being forced to accelerate around the leading edge. This hinders the wing ability to create a strong suction peak in these sections, eventually causing the separation patterns that were observed. The weaker suction peak is also responsible for the lower aerodynamic forces observed for the wavy wings. The simulations with Re = 1,000 also showed that the main impact on the flow is caused by the deformation of the leading edge, with the trailing edge shape having a negligible effect. Therefore, the simulations performed here with a uniform deformation affecting both the leading and trailing edges should be comparable to the results usually encountered in the literature, where only the leading edge is modified.

The second set of simulations considered Re = 10,000. In this case, only four angles of attack were considered: $\alpha = 6^{\circ}$, $\alpha = 12^{\circ}$, $\alpha = 15^{\circ}$ and $\alpha = 18^{\circ}$. In the lowest angle of attack the waviness leads to a strong increase in the lift coefficient, without significant changes to the drag. On the other hand, for the other angles of attack the waviness leads to decreases in both lift and drag, just like the Re = 1,000 case. The large impact of the waviness for $\alpha = 6^{\circ}$ is attributed to a distinction between a laminar flow for the baseline wing and transitional flows for the wavy wings. In all cases, the waviness also caused reductions of lift fluctuations, with a suppression of the tonal component observed in the baseline wing. Another interesting finding for these simulations is that the separation patterns follow the same trend from the Re = 1,000 case, although in this case there is separation along the whole span in the downstream half of the wing. This also seems to be caused by a similar mechanism as that of the lower Re, with the spanwise flow near the leading edge playing a significant role by weakening the suction peaks. In this case, we also observe that the waviness leads to a sub-harmonic behaviour, with the flow being significantly different in subsequent wavelengths. It is also worth mentioning that in some sections the pressure distribution is typical of laminar-separation bubbles, and the flow visualizations support that this feature occurs in these sections.

Finally, the third set of simulations covered the Reynolds number Re = 50,000, where we considered the baseline and one wavy wing for angles of attack $\alpha = 6^{\circ}$ and $\alpha = 15^{\circ}$. This has the importance of allowing comparisons with experimental results, and also being closer to the conditions which might be encountered in some practical applications, like micro-air vehicles. The results show a good agreement with the experimental data, with the waviness increasing the lift in the post-stall regime, with little effect in this respect before the stall. For $\alpha = 6^{\circ}$ the baseline wing displays a typical laminar separation bubble behaviour, with a long bubble extending approximately half of the chord. In this case, the waviness causes the flow to reattach behind the peak, while there is a substantial shortening of the bubble behind the troughs. For the higher angle of attack we observe a low frequency behaviour in the lift coefficient. Although it was not possible to obtain a definite explanation to this phenomenon, some results were presented suggesting it is caused by changes in the separation patterns.

In summary, our results show that the waviness may lead to the following benefits at low Reynolds numbers:

- Suppression of lift coefficient fluctuations, especially tonal components
- Making the wing more resistant to the performance deterioration caused by laminarization of the flow, as was seen for Re = 10,000 and $\alpha = 6^{\circ}$
- In the highest Reynolds number considered here, Re = 50,000, the waviness leads to an increase in the lift coefficient in the post-stall regime.

Also, it was demonstrated that one important mechanism by which the waviness affects the flow is the weakening of the suction peak behind the peak section of the waviness. This effect is present in all Reynolds numbers considered here.

10. Suggestions for future work

Despite the fact that the work presented here is quite extensive, there is plenty of possibilities for complementing it. Furthermore, some of the findings presented open new opportunities for future research. Similarly to the conclusions presented earlier, this opportunities can be divided into those arising from the developments of the numerical methods and those related to the simulations of the wavy wings.

Starting with the numerical methods, an interesting issue concerning the coordinate transformation techniques is its robustness. It would be important to understand in a more precise manner what dictates the numerical stability of those methods, and also search for ways to improve it, allowing us to use them in simulations using transformations with larger deformations. These methods also open several opportunities in terms of applications. A particularly interesting one would be in fluid-structure interaction, where the possibility of using time-dependent mappings seems advantageous over the use of deformable meshes. Considering the limits of these transformations in terms of stability, a hybrid approach could be explored, where the body movement would be treated using a coordinate transformations until the deformation reaches a threshold or for a fixed number of time steps, at which point a re-meshing would be performed. This idea is in some sense analogous to the adaptive polynomial order presented here: if the re-meshing is only performed after several time steps, its high cost is diluted, leading to an efficient algorithm.

Another line of research from this thesis which can be extended is the parallelism strategy for quasi-3D simulations using the spectral/hp methods. As discussed in chapter 5, the use of a direct solver using only a few partitions in the *xy* plane leads to a very efficient solution. However, this is limited to a moderate number of processes, with the scaling becoming very poor when using more partitions. A more scalable approach would probably rely on iterative solvers, but the currently available conjugate gradient implementation has been observed to be no match to the direct solver in terms of computational cost. Therefore, it is important to ascertain whether it would be possible to obtain a scaling to more processes while remaining competitive to the strategy adopted here.

Considering now the study of wavy wings, one characteristic of the flow that deserves further investigation is the sub-harmonic behaviour observed for Re = 10,000 and Re = 50,000. As noted in chapters 7 and 8, subsequent wavelengths can present very different behaviours, including changes in the separation pattern and pressure distribution. Under-

standing the causes for this might provide useful information on how the waviness affect the flow. Obtaining an explanation for this phenomenon would possibly involve a study including stability analysis of the flow. Another feature observed in the present work that could be further explored is the low frequency behaviour for Re = 50,000 and $\alpha = 15^{\circ}$.

In more general terms, the study could be extended by performing more simulations, either at higher Reynolds numbers, extending the parameter space covered here, or by also considering finite wings. However, it is important to note that the limited availability of experimental data makes it difficult to validate the numerical results, since most of the experimental works only present measurements of time-averaged forces. Therefore, future numerical investigation would definitely benefit from experimental works exploring the time evolution of the forces and visualization of the flow. These would possibly confirm the low-frequency and sub-harmonic behaviours discussed in the previous paragraph.

Bibliography

Arai, H.; Doi, Y.; Nakashima, T.; Mutsuda, H. A study on stall delay by various wavy leading edges. **Journal of Aero Aqua Bio-mechanisms**, v. 1, p. 18–23, 2010.

Aris, R. Vectors, Tensors, and the Basic Equations of Fluid Mechanics. Dover Publications, Inc., 1989.

Barkley, D.; Henderson, R. D. Three-dimensional Floquet stability analysis of the wake of a circular cylinder. **Journal of Fluid Mechanics**, v. 332, p. 215–241, 1996.

Bearman, P. W. Investigation of the flow behind a two-dimensional model with a blunt trailing edge and fitted with splitter plates. **Journal of Fluid Mechanics**, v. 21, n. 2, p. 241–255, 1965.

Bearman, P. W.; Owen, J. C. Reduction of bluff-body drag and suppression of vortex shedding by the introduction of wavy separation lines. **Journal of Fluids and Structures**, v. 12, p. 123–130, 1998.

Bearman, P. W.; Tombazis, N. The effects of three-dimensional imposed disturbances on bluff body near wake flows. **Journal of Wind Engineering and Industrial Aerodynamics**, v. 49, p. 339–350, 1993.

Bolis, A. Fourier spectral/hp element method: investigation of time-stepping and parallelisation strategies. PhD thesis, Imperial College London, 2013.

Bolis, A.; Cantwell, C. D.; Moxey, D.; Serson, D.; Sherwin, S. J. An adaptable parallel algorithm for the direct numerical simulation of incompressible turbulent flows using a Fourier spectral/hp element method and MPI virtual topologies. **Computer Physics Communica-tions**, v. 206, p. 17–25, 2016.

Cai, J.; Chng, T. L.; Tsai, H. M. On vortical flows shedding from a bluff body with a wavy trailing edge. **Physics of Fluids**, v. 20, p. 064102, 2008.

Cantwell, C. D.; Moxey, D.; Comerford, A.; Bolis, A.; Rocco, G.; Mengaldo, G.; De Grazia, D.; Yakovlev, S.; Lombard, J.-E.; Ekelschot, D.; Jordi, B.; Xu, H.; Mohamied, Y.; Eskilsson, C.; Nelson, B.; Vos, P.; Biotto, C.; Kirby, R. M.; Sherwin, S. J. Nektar++: An open-source

spectral/ element framework. **Computer Physics Communications**, v. 192, p. 205 – 219, 2015.

Carlson, H. A.; Berkooz, G.; Lumley, J. L. Direct numerical simulation of flow in a channel with complex, time-dependent wall geometries: A pseudospectral method. **Journal of Computational Physics**, v. 121, p. 155–175, 1995.

Carmichael, B. H. Low Reynolds number airfoil survey. **NASA Contractor Report**, n. 165803, 1981.

Custodio, D. The effect of humpback whale-like leading edge protuberances on hydrofoil performance. Master's thesis, Worcester Polytechnic Institute, 2007.

Darekar, R. M.; Sherwin, S. J. Flow past a bluff body with a wavy stagnation face. **Journal** of Fluids and Structures, v. 15, n. 3, p. 587–596, 2001a.

Darekar, R. M.; Sherwin, S. J. Flow past a square-section cylinder with a wavy stagnation face. **Journal of Fluid Mechanics**, v. 426, p. 263–295, 2001b.

Dugan, D. W. Effects of three types of blunt trailing edges on the aerodynamic characteristics of a plane tapered wing of aspect ratio 3.1, with a 3-percent-thick biconvex section. **NACA Research Memorandum**, n. A52E01, 1952.

Evangelinos, C.; Karniadakis, G. E. Communication performance models in Prism: A spectral element-Fourier parallel Navier-Stokes solver. In **Proceedings of the 1996 ACM/IEEE Conference on Supercomputing**, 1996.

Favier, J.; Pinelli, A.; Piomelli, U. Control of the separated flow around an airfoil using a wavy leading edge inspired by humpback whale flippers. **Comptes Rendus Mecanique**, v. 340, p. 107–114, 2012.

Fish, F. E.; Battle, J. M. Hydrodynamic design of the humpback whale flipper. **Journal of Morphology**, v. 225, p. 51–60, 1995.

Gad-el Hak, M. Modern developments in flow control. **Applied Mechanics Reviews**, v. 49, p. 365–379, 1996.

Gad-el Hak, M. Flow Control: Passive, Active, and Reactive Flow Management. Cambridge University Press, 1st edition, 2000.

Gai, S. L.; Sharma, S. D. Measurements in the wake of segmented blunt trailing edge aerofoil in subsonic flow. In **Proc. 8th Australasian Fluid Mechanics Conference**, v. 13C, p. 9–11, 1983.

Galbraith, M.; Visbal, M. Implicit large eddy simulation of low Reynolds number flow past the SD7003 airfoil. In **Proceedings of the 46th AIAA Aerospace Sciences Meeting and Exhibit**, 2008.

Guermond, J. L.; Shen, J. Velocity-correction projection methods for incompressible flows. **SIAM J. Numer. Anal.**, v. 41, n. 1, p. 112–134, 2003.

Guerreiro, J. L. E.; Sousa, J. M. M. Low-Reynolds-number effects in passive stall control using sinusoidal leading edges. **AIAA Journal**, v. 50, n. 2, p. 461–469, 2012.

Hansen, K. L.; Kelso, R. M.; Dally, B. B. An investigation of three-dimensional effects on the performance of tubercles at low Reynolds numbers. In **17th Australasian Fluid Mechanics Conference**, 2010.

Hieser, G. An investigation at transonic speeds of the effects of fences, drooped nose, and vortex generators on the aerodynamic characteristics of a wing-fuselage combination having a 6-percent-thick, 45° sweptback wing. **NACA Research Memorandum**, n. L53B04, 1953.

Howe, M. S. Aerodynamic noise of a serrated trailing edge. **Journal of Fluids and Structures**, v. 5, n. 1, p. 33 – 45, 1991.

Jaworski, J. W.; Peake, N. Aerodynamic noise from a poroelastic edge with implications for the silent flight of owls. **Journal of Fluid Mechanics**, v. 723, p. 456–479, 2013.

Johari, H.; Henoch, C.; Custodio, D.; Levshin, A. Effects of leading edge protuberances on airfoil performance. **AIAA Journal**, v. 45, n. 11, p. 2634–2641, 2007.

Karniadakis, G. E. Spectral element-Fourier methods for incompressible turbulent flows. **Computer Methods in Applied Mechanics and Engineering**, v. 80, n. 1–3, p. 367–380, 1990.

Karniadakis, G. E.; Sherwin, S. J. **Spectral/hp Element Methods for Computational Fluid Dynamics**. Oxford University Press, 2nd edition, 2005.

Karniadakis, G. E.; Israeli, M.; Orszag, S. A. High-order splitting methods for the incompressible Navier-Stokes equations. **Journal of Computational Physics**, v. 97, n. 2, p. 414–443, 1991.

Kirby, R. M.; Sherwin, S. J. Aliasing errors due to quadratic nonlinearities on triangular spectral/hp element discretisations. **Journal of Engineering Mathematics**, v. 56, p. 273–288, 2006a.

Kirby, R. M.; Sherwin, S. J. Stabilisation of spectral/hp element methods through spectral vanishing viscosity: Application to fluid mechanics modelling. **Computer Methods in Applied Mechanics and Engineering**, v. 195, n. 23–24, p. 3128–3144, 2006b.

Koberg, H. **Turbulence control for drag reduction with active wall deformation**. PhD thesis, Imperial College of Science, Technology and Medicine, 2007.

Kovasznay, L. I. G. Laminar flow behind a two-dimensional grid. **Mathematical Proceedings of the Cambridge Philosophical Society**, v. 44, n. 1, p. 58–62, 1948.

Levshin, A.; Custodio, D.; Henoch, C.; Johari, H. Effects of leading edge protuberances on airfoil performance. In **36th AIAA Fluid Dynamics Conference and Exhibit**, 2006.

Luo, H.; Bewley, T. R. On the contravariant form of the Navier-Stokes equations in timedependent curvilinear coordinate systems. **Journal of Computational Physics**, v. 199, p. 355–375, 2004.

Malipeddi, A. K. Numerical analysis of effects of leading-edge protuberances on aircraft wing performance. Master's thesis, Wichita State University, 2011.

Malipeddi, A. K.; Mahmoudnejad, N.; Hoffmann, K. A. Numerical analysis of effects of leading-edge protuberances on aircraft wing performance. **Journal of Aircraft**, v. 49, n. 5, p. 1336–1344, 2012.

Miklosovic, D. S.; Murray, M. M.; Howle, L. E.; Fish, F. E. Leading-edge tubercles delay stall on humpback whale (Megaptera novaeangliae) flippers. **Physics of Fluids**, v. 16, p. L39, 2004.

Miklosovic, D. S.; Murray, M. M.; Howle, L. E. Experimental evaluation of sinusoidal leading edges. **Journal of Aircraft**, v. 44, n. 4, p. 1404–1407, 2007.

Mueller, T. J.; DeLaurier, J. D. Aerodynamics of small vehicles. **Annual Review of Fluid Mechanics**, v. 35, n. 1, p. 89–111, 2003.

Neuhart, D. H.; Pendergraft Jr., O. C. A water tunnel study of Gurney flaps. **NASA Techni**cal Memorandum, n. 4071, 1988.

Newman, D. J.; Karniadakis, G. E. A direct numerical simulation study of flow past a freely vibrating cable. **Journal of Fluid Mechanics**, v. 344, p. 95–136, 1997.

Paula, A. A. The airfoil thickness effects on wavy leading edge phenomena at low **Reynolds number regime**. PhD thesis, University of São Paulo, 2016.

Pedro, H. T. C.; Kobayashi, M. H. Numerical study of stall delay on humpback whale flippers. In **46th AIAA Aerospace Sciences Meeting and Exhibit**, 2008. Persson, P.-O.; Peraire, J. Sub-cell shock capturing for discontinuous Galerkin methods. **AIAA paper**, v. 112, 2006.

Petrusma, M.; Gai, S. L. Investigation into the wake of blunt trailing edge airfoils at low Reynolds numbers. In **Proc. 10th Australasian Fluid Mechanics Conference**, v. 13, p. 35–38, 1989.

Rostamzadeh, N.; Kelso, R. M.; Dally, B. B.; Hansen, K. L. The effect of wavy leading edge modifications on NACA 0021 airfoil characteristics. In **18th Australasian Fluid Mechanics Conference**, 2012.

Serson, D.; Meneghini, J. R. Numerical study of wings with wavy leading and trailing edges. In **Proceedings of the IUTAM-ABCM 8th Symposium on Laminar Turbulent Transition**, 2014.

Serson, D.; Meneghini, J. R.; Sherwin, S. J. Direct numerical simulations of the flow around wings with spanwise waviness at a very low Reynolds number. **Computers & Fluids (sub-mitted to publication)**, 2016a.

Serson, D.; Meneghini, J. R.; Sherwin, S. J. Extension of the velocity-correction scheme to general coordinate systems. In **Proceedings of the International Conference on Spectral and High Order Methods**, 2016b.

Serson, D.; Meneghini, J. R.; Sherwin, S. J. Velocity-correction schemes for the incompressible Navier-Stokes equations in general coordinate systems. **Journal of Computational Physics**, v. 316, p. 243–254, 2016c.

Serson, D.; Meneghini, J. R.; Sherwin, S. J. Direct numerical simulations of the flow around wings with spanwise waviness. **Journal of Fluid Mechanics (submitted to publication)**, 2016d.

Shun, S.; Ahmed, N. A. Wind turbine performance improvements using active flow control techniques. **Procedia Engineering**, v. 49, p. 83 – 91, 2012. International Energy Congress 2012.

Skillen, A.; Revell, A.; Pinelli, A.; Piomelli, U.; Favier, J. Flow over a wing with leading-edge undulations. **AIAA Journal**, v. 53, n. 2, p. 464–472, 2015.

Stanway, M. J. Hydrodynamic effects of leading-edge tubercles on control surfaces and in flapping foil propulsion. Master's thesis, Massachusetts Insitute of Technology, 2008.

Storms, B. L.; Jang, C. S. Lift enhancement of an airfoil using a Gurney flap and vortex generators. **Journal of Aircraft**, v. 31, n. 3, p. 542–547, 1994.

Summers, J. L.; Page, W. A. Lift and moment characteristics at subsonic Mach numbers of four 10-percent-thick airfoil sections of varying trailing-edge thickness. **NACA Research Memorandum**, n. A50J09, 1950.

Tombazis, N.; Bearman, P. W. A study of three-dimensional aspects of vortex shedding from a bluff body with a mild geometric disturbance. **Journal of Fluid Mechanics**, v. 330, p. 85–112, 1997.

Tufo, H.; Fischer, P. Fast parallel direct solvers for coarse grid problems. **Journal of Parallel and Distributed Computing**, v. 61, n. 2, p. 151–177, 2001.

van Nierop, E. A.; Alben, S.; Brenner, M. P. How bumps on whale flippers delay stall: an aerodynamic model. **Physical Review Letters**, v. 100, p. 054502, 2008.

Watts, P.; Fish, F. E. The influence of passive, leading edge tubercles on wing performance. In **Proc. Twelfth Intl. Symp. Unmanned Untethered Submers. Technol**. Auton. Undersea Syst. Inst. Durham New Hampshire, 2001.

Weber, P. W.; Howle, L. E.; Murray, M. M.; Miklosovic, D. S. Computational evaluation of the performance of lifting surfaces with leading-edge protuberances. **Journal of Aircraft**, v. 48, n. 2, p. 591–600, 2011.

Williams, M. D.; Reeder, M. F.; Maple, R. C.; Solfelt, D. A. Modeling, simulation, and flight tests for a T-38 Talon with wing fences. **Journal of Aircraft**, v. 47, n. 2, p. 423 – 433, 2010.

Yoon, H. S.; Hung, P. A.; Jung, J. H.; Kim, M. C. Effects of the wavy leading edge on hydrodynamic characteristics for flow around low aspect ratio wing. **Computers & Fluids**, v. 49, p. 276–289, 2011.

Zhang, M. M.; Wang, G. F.; Xu, J. Z. Aerodynamic control of low-Reynolds-number airfoil with leading-edge protuberances. **AIAA Journal**, v. 51, n. 8, p. 1960–1971, 2013.

Appendices

A. Tensor calculus

This appendix provides a brief introduction to tensor calculus. The goal here is to present the main elements needed to understand how the Navier-Stokes equations can be represented in a general coordinate system. For a more detailed presentation of this topic, the reader is referred to Aris (1989), which served as a basis for this chapter.

A.1 Coordinate systems transformations

In the physical three-dimensional Euclidean space, the most common way to describe the position of a point in space is by its components in a Cartesian coordinate system. This consists in representing the point by three numbers (x, y, z), which correspond to the distance from the point to the orthogonal coordinate planes passing through the origin of the system. However, it is possible to represent the position of a point using other coordinates, for example the cylindrical polar coordinate system.

Given two coordinate systems (x^1, x^2, x^3) and $(\bar{x}^1, \bar{x}^2, \bar{x}^3)$ related by

$$\bar{x}^i = \bar{x}^i (x^1, x^2, x^3), \qquad i = 1, 2, 3,$$
(A.1)

we say that the transformation is proper if the Jacobian of the transformation,

$$J = \left| \frac{\partial \bar{x}}{\partial x} \right| \tag{A.2}$$

exists and is non-zero. This implies that the transformation can be inverted, and therefore we can go back to the original coordinate system, making both representations equivalent.

Also, note that in the previous discussion both x^i and \bar{x}^i can represent general coordinate systems. In some occasions, it is useful to distinguish the Cartesian coordinate system, since many properties are well-known for this system. For this end, from now on the notation y^i will be used to refer specifically to the Cartesian coordinates.

A.2 Contravariant vectors

The behaviour of the differentials of the coordinates under transformation can be obtained using the chain rule, resulting in

$$d\bar{x}^{i} = \frac{\partial \bar{x}^{i}}{\partial x^{j}} dx^{j}, \tag{A.3}$$

where the summation convention is that in a product, an index repeated once in an upper position and once in a lower position is summed (the index j in the denominator counts as a lower index).

Based on this behaviour, a contravariant vector ${\bf u}$ is defined as an entity represented by components which transform according to

$$\bar{u}^i = \frac{\partial \bar{x}^i}{\partial x^j} u^j = c^i_j u^j, \tag{A.4}$$

where the last identity is a definition to simplify the notation. It is a convention to represent a contravariant vector by an upper index.

A.3 Covariant vectors

If we take a scalar function $f(x^1, x^2, x^3)$ and calculate its partial derivatives in the system $O\overline{1}\overline{2}\overline{3}$, we obtain

$$\frac{\partial f}{\partial \bar{x}^i} = \frac{\partial x^j}{\partial \bar{x}^i} \frac{\partial f}{\partial x^j},\tag{A.5}$$

which is a different type of behaviour under transformation when compared to equation (A.4).

Extending upon this particular case, a covariant vector is defined as an entity whose components transform as

$$\bar{u}_i = \frac{\partial x^j}{\partial \bar{x}^i} u_j = \bar{c}_i^j u_j, \tag{A.6}$$

and is by convention represented by a lower index.

Like in the typical case of Cartesian vectors, the contravariant and covariant vectors defined here are physical entities which do not depend on the coordinate system. Only the coordinates which are used to represent the vectors change when we make a coordinate system transformation.

A.4 Higher order tensors

The previous definitions can be naturally generalized to higher orders. A tensor of contravariant order m and covariant order n (or simply order m + n) is defined as an entity with
$3^{(m+n)}$ components that transform according to

$$\bar{A}_{q_1\cdots q_n}^{p_1\cdots p_m} = \frac{\partial \bar{x}^{p_1}}{\partial x^{i_1}}\cdots \frac{\partial \bar{x}^{p_m}}{\partial x^{i_m}}\frac{\partial x^{j_1}}{\partial \bar{x}^{q_1}}\cdots \frac{\partial x^{j_n}}{\partial \bar{x}^{q_n}}A_{j_1\cdots j_n}^{i_1\cdots i_m}.$$
(A.7)

When the right-hand-side of equation (A.7) also contains a factor J^w , the tensor is called a relative tensor of weight w. When it is necessary to make a distinction, a tensor satisfying equation (A.7) can be called an absolute tensor.

A.5 The metric tensor

In Cartesian coordinates, the distance between two points y^i and $y^i + dy^i$ is

$$ds^{2} = \sum_{k=1}^{3} dy^{k} dy^{k}.$$
 (A.8)

In this equation, the summation has to be written explicitly, since the repeated indices are both in the upper position, and therefore the summation convention does not apply. In order to obtain the distance in term of differentials in the system x^i , the terms in the right-hand-side can be transformed using equation (A.3), leading to

$$ds^{2} = \sum_{k=1}^{3} \left(\frac{\partial y^{k}}{\partial x^{i}} \frac{\partial y^{k}}{\partial x^{j}} \right) dx^{i} dx^{j} = g_{ij} dx^{i} dx^{j},$$
(A.9)

where the last identity contains the definition of the metric tensor g_{ij} . It is important to note that the metric tensor is defined based on the transformation to a Cartesian coordinate system, since this is the case where the distance between two points is given by equation (A.8). Also, instead of following this argument, in a more general sense the metric tensor can be directly defined as the tensor relating coordinate differentials and distances, without requiring the existence of the Cartesian system.

The inverse of the metric tensor is denoted g^{ij} , and by construction $g^{ij}g_{jk} = \delta^i_j$ is the Kronecker delta. Also, if we define g as the determinant of the metric tensor, then it can be shown that $J = g^{1/2}$ is the Jacobian of the transformation to the Cartesian system.

Another important aspect of the metric tensor is the notion of associated tensors. Starting with a contravariant vector u^i , we can obtain the associated covariant vector by $u_i = g_{ij}u^j$. This operation is called lowering the index. Similarly, we can obtain a contravariant vector associated with the covariant vector v_i by $v^i = g^{ij}v_j$, and this operation is called raising the index. Since $g^{ij}g_{jk} = \delta^i_j$, by lowering the index of a vector and subsequently raising the index, we obtain the original vector. Therefore, there is a unique relation between a pair of associated vectors, and any vector can be represented either in covariant or contravariant form by using these operations. It is possible to demonstrate that a con140

travariant representation corresponds to the components of the vector in terms of a basis of vectors tangent to the coordinate lines, while the covariant representation corresponds to the components of the vector in terms of a basis formed by vectors normal to coordinate surfaces.

The metric tensor can also be used to obtain the length of a vector and the angle between vectors. The length of a contravariant vector can be obtained by generalizing equation (A.9), leading to

$$|A|^2 = g_{ij} A^i A^j. (A.10)$$

Using the definition of associated vectors, this equation can be restated in the following forms:

$$|A|^{2} = g_{ij}A^{i}A^{j} = A^{i}A_{i} = g^{ij}A_{i}A_{j}.$$
(A.11)

Also, the angle between two vectors A^i and B^i is

$$\cos\theta = \frac{g_{ij}A^iB^j}{|A||B|}.$$
(A.12)

From this equation, it is easy to note that when g_{ij} is diagonal the coordinate lines are orthogonal, and therefore the coordinate system is called orthogonal.

A.6 The covariant derivative

In a Cartesian coordinate system, the partial derivatives of a tensor form another tensor with increased order. However, this is not the case in a general system. Although by equation (A.5) the partial derivatives of a scalar form a covariant vector, the partial derivatives of a vector or of higher order tensors are not themselves tensors. Therefore, it is necessary to generalize the differentiation process in a way which yields tensors and reduces to the typical partial differentiation in the particular case of a Cartesian system. The operation which results from this generalization is the covariant derivative.

Before proceeding to the definition of the covariant derivative, it is important to introduce the Christoffel symbols. The Christoffel symbols of the first kind are:

$$[jk,i] = \frac{1}{2} \left(\frac{\partial g_{ij}}{\partial x^k} + \frac{\partial g_{ik}}{\partial x^j} - \frac{\partial g_{jk}}{\partial x^i} \right), \tag{A.13}$$

and the Christoffel symbols of the second kind are defined as

$$\begin{cases} i \\ j k \end{cases} = g^{ip}[jk, p]. \tag{A.14}$$

Using this definition, the covariant derivative of a contravariant vector is defined as

$$A_{,k}^{j} = \frac{\partial A^{j}}{\partial x^{k}} + \begin{cases} j\\ i & k \end{cases} A^{i},$$
(A.15)

where a comma before the subscript k indicates the covariant differentiation. Similarly, for a covariant vector:

$$A_{j,k} = \frac{\partial A_j}{\partial x^k} - \begin{cases} i\\ j & k \end{cases} A_i.$$
(A.16)

These expressions can be easily extended to higher order tensors, with one extra term appearing in the right-hand-side of the equation for each index of the tensor.

In equations (A.15) and (A.16), the extra terms are related to the variation of the basis vectors. In the case of a Cartesian system, the Christoffel symbols vanish and therefore the relations reduce to the partial differentiation. Also, the covariant derivatives of the metric tensor and of its inverse are always zero.

A.7 Divergence and curl of vectors

The divergence and curl operators can be generalized to retain the tensorial character of vectors. Using the covariant differentiation, the divergence can be defined as

$$div(A^{i}) = A^{i}_{,i} = \frac{1}{J} \frac{\partial}{\partial x^{i}} \left(JA^{i} \right).$$
(A.17)

The curl is defined as

$$curl(A^i) = \varepsilon^{ijk} A_{k,j} = \varepsilon^{ijk} g_{kp} A^p_{,j},$$
 (A.18)

where $\varepsilon^{ijk} = g^{-1/2} \epsilon^{ijk}$ (with ϵ^{ijk} being the permutation symbol) is a third order tensor which reduces to the permutation symbol in the Cartesian system. This tensor can also be represented in covariant form as $\varepsilon_{ijk} = g^{1/2} \epsilon_{ijk}$, where ϵ_{ijk} corresponds again to the permutation symbol.

A.8 The generalized Navier-Stokes equations

In a Cartesian coordinate system, the Navier-Stokes equations for an incompressible flow can be represent in non-dimensional form as

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \frac{1}{Re}\nabla^2 \mathbf{u}$$

$$\nabla \cdot \mathbf{u} = 0,$$
(A.19)

where \mathbf{u} is the velocity, p is the pressure and Re is the Reynolds number.

This is not a valid tensor identity, and therefore is not valid in a general coordinate system. However, it can be transformed into a tensor relation by substituting the partial derivatives by covariant derivatives and by raising the indices of covariant vectors, in order to make all terms of the equation into contravariant vectors. This procedure results in

$$\frac{\partial u^{i}}{\partial t} + u^{j}u^{i}_{,j} = -g^{ij}p_{,j} + \frac{1}{Re}(g^{jk}u^{i}_{,j})_{,k}$$

$$u^{i}_{,i} = 0.$$
(A.20)

In this equation, all terms are contravariant vectors, having the same behaviour under a coordinate system transformation. Therefore, the identity is independent of the coordinate system. Also, for a Cartesian system it reduces to the typical Navier-Stokes equations. From these considerations, it can be deduced that equation (A.20) is the appropriate form of the Navier-Stokes equations in a general time-independent coordinate system. When the coordinate system is time-dependent, we must also account for differences in the partial derivative with respect to time. Luo and Bewley (2004) showed that in this case, the appropriate form of the equations is

$$\frac{\partial u^{i}}{\partial t} + u^{j}u^{i}_{,j} - V^{j}u^{i}_{,j} + u^{j}V^{i}_{,j} = -g^{ij}p_{,j} + \frac{1}{Re}(g^{jk}u^{i}_{,j})_{,k}$$

$$u^{i}_{,i} = 0,$$
(A.21)

where $V^i = -\frac{\partial x^i}{\partial t}$ is the velocity of the coordinates, represented in the transformed system. This term can be obtained by first calculating the contravariant vector $\bar{V}^i = \frac{\partial y^i}{\partial t}$ representing the velocity of the coordinates in the Cartesian system, and then converting it to the transformed system.

A.9 Application to wings with waviness

The formulation presented in the previous sections is very general, and because of that it can be very inefficient if applied directly for simple transformations. The calculation of the covariant derivatives, for example, require the computation and storage of the Christoffel

symbols, which in general consists of 18 components (since they are symmetric with regard to the j and k indices of equation (A.13)). In this section, the appropriate form of the several elements presented above are described for the particular transformation considered in the present work.

From now on, we denote the Cartesian system by (x, y, z) and the transformed system by $(\bar{x}, \bar{y}, \bar{z})$. For most of the transformations used in the present work, these systems are related by

$$x = f(\bar{x}, \bar{z})$$

$$y = \bar{y}$$

$$z = \bar{z}.$$

(A.22)

Denoting the partial derivatives of f with respect to \bar{x} and \bar{z} by f_x and f_z , we obtain the following results:

• Transformation from transformed to Cartesian system:

$$\bar{c}_i^j = \begin{bmatrix} f_x & 0 & f_z \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$
(A.23)

• Metric tensor:

$$g_{ij} = \begin{bmatrix} f_x^2 & 0 & f_x f_z \\ 0 & 1 & 0 \\ f_x f_z & 0 & 1 + f_z^2 \end{bmatrix}$$
(A.24)

• Inverse of the metric tensor:

$$g^{ij} = \begin{bmatrix} \frac{1+f_z^2}{f_x^2} & 0 & -\frac{f_z}{f_x} \\ 0 & 1 & 0 \\ -\frac{f_z}{f_x} & 0 & 1 \end{bmatrix}$$
(A.25)

• Jacobian:

$$J = f_x \tag{A.26}$$

• Christoffel symbols of the second kind:

$$\begin{cases} 1\\ j & k \end{cases} = \frac{1}{f_x} \begin{bmatrix} f_{xx} & 0 & f_{xz} \\ 0 & 0 & 0 \\ f_{xz} & 0 & f_{zz} \end{bmatrix}$$

$$\begin{cases} 2\\ j & k \end{cases} = \begin{cases} 3\\ j & k \end{cases} = 0$$
(A.27)

These expressions show how in this particular case the transformation becomes much simpler. It is only necessary to store five terms, corresponding to the first and second order partial derivatives of f, and only four of the Christoffel symbols are non-zero.

B. Mesh generation and convergence tests

In order to guarantee the accuracy of the numerical results, it is important to perform tests to determine an appropriate mesh and some parameters to be used in the simulations.

As discussed in chapter 6, the airfoil profile used in the simulations is given by:

$$\frac{y}{c} = \pm \frac{0.12}{0.2} \left[0.2969 \sqrt{\frac{x}{c}} - 0.1260 \left(\frac{x}{c}\right) - 0.3516 \left(\frac{x}{c}\right)^2 + 0.2843 \left(\frac{x}{c}\right)^3 - 0.1036 \left(\frac{x}{c}\right)^4 \right]$$
(B.1)

which is a slightly modified NACA0012 with zero-thickness trailing-edge.

The procedure adopted here is to first generate an initial mesh representing the computational domain with the distribution of elements reflecting the expected local resolution requirements. Then, the parameters of interest were varied one at a time to verify how they affect the results. The preferred approach to increasing the spatial resolution was to increase the polynomial order of the basis function, leading to what is called a p-refinement.

Each of the following sections presents the convergence tests for a given value of Reynolds number. For the lower value Re = 1,000, the simulations are relatively cheaper, and therefore several tests were conducted. For Re = 10,000 and Re = 50,000, the number of tests had to be reduced, due to the increased computational cost, and therefore they focused on the spatial convergence.

B.1 Convergence tests for Re = 1,000

The convergence tests for Re = 1,000 considered initially the mesh of figure B.1. This mesh is formed by 449 quadrilateral elements, extending from -7 to 10 in the *x* direction and from -8 to 8 in the *y* direction, with the airfoil having a unit chord aligned with the *x* axis and the coordinate system origin coinciding with the leading edge.

Following that, several convergence tests were performed to define the following parameters: simulation time, basis functions polynomial order, mesh size, time-step, and number of Fourier modes in the span direction. The tests were carried out with angles of attack $\alpha = 5^{\circ}$ and $\alpha = 18^{\circ}$, where the angle of attack is adjusted through the inflow boundary condition. Also, most tests were performed using two-dimensional simulations, with



Figure B.1: Initial mesh used in the convergence tests for Re = 1,000. On the left the entire domain is shown, while on the right only the region close to the airfoil is presented.

the exception being the tests for the number of Fourier modes in the span direction. The conditions for this last case will be detailed in the section related to them.

For the case with $\alpha = 5^{\circ}$, it was observed that the simulation reaches a steady state. Therefore, it is possible to check the accuracy of the solution considering only the final flow field. The simulations with $\alpha = 18^{\circ}$, however, do not reach a steady state, and therefore the last 10 time units were considered to calculate the results. The results were compared with respect to the mean drag and lift coefficients. Also, for $\alpha = 5^{\circ}$ the infinity norm (maximum value at the quadrature points) of the deviation with respect to the most accurate case of each section was computed, what is easily done since the flow reaches a steady state, while for $\alpha = 18^{\circ}$ the *rms* value of the lift coefficient fluctuations was calculated. Throughout the convergence tests, it was considerable acceptable the parameter values which lead to relative errors inferior to 0.2% in the lift and drag coefficients for the case with $\alpha = 5^{\circ}$, with the results for $\alpha = 18^{\circ}$ being used as estimates for the error occurring in the more critical cases.

In the following, the test results will be presented, with each section describing one parameter.

B.1.1 Simulation time

In the numerical simulations, it is necessary to consider a sufficiently long simulation time in order to eliminate the effects of the initial transient caused by the initial condition. To determine for how long the flow should be analysed, simulation with final time T equal to 100, 200 and 400 non-dimensional time units were performed. All of them used basis functions with 9^{th} order polynomials and a time step 0.0005, therefore guaranteeing that the only parameter affecting the results is T.

Table B.1 contains the results for $\alpha = 5^{\circ}$, while table B.2 presents the results for $\alpha = 18^{\circ}$. From this results and considering the adopted convergence criteria, it is clear that

Time	$\overline{C_d}$	$\overline{C_l}$	$ u - u_{T=400} _{\infty}$	$ v - v_{T=400} _{\infty}$
100	0.1280	0.2457	2.51E-11	1.65E-11
200	0.1280	0.2457	2.56E-11	1.68E-11
400	0.1280	0.2457		

Table B.1: Convergence tests for the simulation time with $\alpha = 5^{\circ}$.

Table B.2: Convergence tests for the simulation time with $\alpha = 18^{\circ}$.

Time	$\overline{C_d}$	$\overline{C_l}$	$C_l^\prime(r.m.s)$
100	0.3908	0.8718	0.1645
200	0.3914	0.8755	0.1660
400	0.3915	0.8755	0.1647

performing the simulations until T = 100 is enough to eliminate the initial transients.

B.1.2 Basis functions polynomial order

In the spectral/hp element method, it is possible to increase the spatial resolution by mesh refinement (also called h refinement), by increasing the basis functions polynomial order (also called p refinement), or by a combination of both approaches. The procedure followed here was to use the mesh from figure B.1 to obtain a higher refinement in the regions where high velocity gradients are expected to occur, and guarantee spatial convergence by increasing the polynomial order.

Tests were performed using polynomial orders between 6 and 9, and a simulation with order 14 was done to act as a reference for the results. For the calculation of the norm of the error, all solutions were projected to the quadrature points of the solution with order 14. This was done in order to make sure that all results are compared with respect to the same set of points; otherwise, low order solutions could benefit from the fact that they are originally obtained in fewer points.

Tables B.3 and B.4 contain the results for these tests, being presented in the same form as in the previous section. From the results, it can be observed that 7th order polynomials are sufficient to obtain an appropriate resolution. However, order 8 was chosen as a measure of precaution, and also to allow the flow field visualizations to have higher quality.

B.1.3 Domain size

Another factor that might compromise the accuracy of the results is the size of the computational domain. If the domain is not large enough, the boundary conditions will be imposed in a location where they are not valid, causing a blockage phenomenon which can affect the results.

Order	$\overline{C_d}$	$\overline{C_l}$	$ u - u_{P=14} _{\infty}$	$ v - v_{P=14} _{\infty}$
6	0.1281	0.2462	1.04E-01	1.29E-01
7	0.1280	0.2457	5.37E-03	1.29E-02
8	0.1280	0.2457	4.07E-03	1.01E-02
9	0.1280	0.2457	4.05E-03	8.69E-3
14	0.1279	0.2456		

Table B.3: Convergence tests for basis functions polynomial order with $\alpha = 5^{\circ}$.

Table B.4: Convergence tests for basis functions polynomial order with $\alpha = 18^{\circ}$.

Order	$\overline{C_d}$	$\overline{C_l}$	$C_l^\prime(r.m.s)$
6	0.3918	0.8726	0.1640
7	0.3912	0.8720	0.1644
8	0.3911	0.8722	0.1650
9	0.3908	0.8718	0.1645
14	0.3895	0.8686	0.1632

Tests were performed using 3 domains with different inlet length (L_i) and lateral sizes (L_h) . The outflow length was kept constant at 10 chords, since results available in the literature, like the ones in Barkley and Henderson (1996), show that for similar simulations with a circular cylinder the effect of this parameter is secondary compared to the inlet and lateral sizes, at least when the outflow length has this order of magnitude. Therefore, the domains of this section extend from $-L_i$ to 10 in the *x* direction and from $-L_h$ to L_h in the *y* direction.

Tables B.5 and B.6 show the results for $\alpha = 5^{\circ}$ and $\alpha = 18^{\circ}$, respectively. In this case the norm of the error was not considered, since this calculation becomes difficult because of the fact that the simulations use different meshes. However, this should not be a problem, since errors in the local values of the solution are expected to be caused by a low spatial resolution, and this has been dealt with in tests for the polynomial order. Based on the results, the second domain was chosen, with $L_h = 15$ and $L_i = 10$. The mesh for this domain, composed of 549 quadrilateral elements, was used in the simulations of chapter 6, and is shown in figure B.2.

L_h/c	L_i/c	$\overline{C_d}$	$\overline{C_l}$
8	7	0.1280	0.2457
15	10	0.1276	0.2438
20	15	0.1274	0.2434

Table B.5: Convergence tests for domain size with $\alpha = 5^{\circ}$.

L_h/c	L_i/c	$\overline{C_d}$	$\overline{C_l}$	$C_l'(r.m.s)$
8	7	0.3911	0.8722	0.1650
15	10	0.3861	0.8599	0.1614
20	15	0.3835	0.8574	0.1592

Table B.6: Convergence tests for domain size with $\alpha = 18^{\circ}$.



Figure B.2: Final mesh obtained from the convergence tests for Re = 1,000.

B.1.4 Time-step

The method employed in the simulations requires a low time-step Δt in the time integration, due to the explicit treatment of the convective term. This introduces a numerical stability restriction based on the CFL number, which Karniadakis and Sherwin (2005) show that for spectral/hp discretizations can be evaluated by

$$CFL = \max_{e=1}^{N_{el}} 0.2\Delta t |\mathbf{V}^{st}| P^2$$
(B.2)

where N_{el} is the number of elements in the mesh, $|\mathbf{V}^{st}|$ is the maximum local velocity calculated in a standard element, and P is the polynomial order. In all simulations of part III the choice of time-step led to a CFL which is approximately between 0.4 and 0.6.

Considering this restriction on the time-step, it is expected that the value $\Delta t = 0.0005$ used to make the simulation converge should be low enough to guarantee a good temporal resolution. To check if this is indeed the case, another simulation was performed reducing the time-step by half, with first-order time integration being used in both cases. As tables B.7 and B.8 show, reducing Δt has little influence on the results, and therefore whenever possible a time-step $\Delta t = 0.0005$ was used, with the time-step being reduced if this was necessary for a particular simulation to converge.

Δt	$\overline{C_d}$	$\overline{C_l}$	$ u - u_{\Delta t = 0,00025} _{\infty}$	$ v - v_{\Delta t = 0,00025} _{\infty}$
0.0005	0.1276	0.2438	1.79E-3	1.05E-3
0.00025	0.1276	0.2437		

Table B.7: Convergence tests for time-step for simulations with $\alpha = 5^{\circ}$.

Table B.8: Convergence tests for the time-step for simulations with $\alpha = 18^{\circ}$.

Δt	$\overline{C_d}$	$\overline{C_l}$	$C_l^\prime(r.m.s)$
0.0005	0.3861	0.8599	0.1614
0.00025	0.3857	0.8592	0.1603

B.1.5 Number of Fourier modes in the span direction

In the three-dimensional simulation of an infinite wing (straight or wavy), the methodology proposed by Karniadakis (1990) was employed, which consists in using a spectral/hp discretization in the plane of the cross-section of the body, and discretizing the span direction using a Fourier series. In order to determine the number of Fourier modes (N_z) that should be used, simulations were performed for the case of a wing with sinusoidal waviness of wavelength of 50% of the chord and amplitude of 5% of the chord. Here, N_z refers to the number of real modes (or the total number of degrees of freedom) in the Fourier direction. This is twice as much as the number of complex Fourier modes, which is sometimes used in this context. The periodic length in the span direction chosen for these simulations was equal to the waviness wavelength.

Tables B.9 and B.10 show the results for these tests. From these results, it was decided that using 8 modes for this domain with spanwise length equal to 0.5 leads to an acceptable resolution. Since the simulations in chapter 6 consider a periodic length equal to the chord, 16 modes were used for those simulations, in order to maintain the same spatial resolution.

B.2 Convergence tests for $\mathbf{Re} = \mathbf{10}, \mathbf{000}$

Since the simulations with Re = 10,000 are computationally more expensive, the convergence tests in this case could not be so extensive as for the lower Reynolds number.

Table B.9: Convergence tests for number of Fourier modes in the span direction. Simulations with $\alpha = 5^{\circ}$.

N_z	$\overline{C_d}$	$\overline{C_l}$	$ u - u_{Nz=16} _{\infty}$	$ v - v_{Nz=16} _{\infty}$
4	0.1324	0.2128	1.85-03	1.28E-03
8	0.1324	0.2129	2.02E-05	2.12E-05
16	0.1324	0.2129		

 $\overline{C_d}$ $\overline{C_l}$ N_z $C_l'(r.m.s)$ 0.7352 4 0.3482 0.0972 8 0.3152 0.6625 0.0437 0.3211 0.6797 0.0514 16

Table B.10: Convergence tests for number of Fourier modes in the span direction. Simulations with $\alpha = 18^{\circ}$.

Figure B.3: Detail of the mesh used for simulations with Re = 10,000.

Therefore, the tests focused in a first moment on determining the appropriate simulation time and averaging procedure to obtain representative values of the forces, and then on estimating the errors associated with the spatial resolution. All simulations considered a wing without any waviness and an angle of attack $\alpha = 12^{\circ}$, with a spanwise length of 1.0. The size of the computational domain was the same as in the previous case, but a new mesh with more elements close to the wing was used. This mesh consists of 1041 quadrilateral elements, and the region close to the airfoil is presented in figure B.3.

The first step consisted in determining an averaging procedure to obtain the forces around the wing. To this extent, a simulation was performed for 250 time units, using 10^{th} order polynomials in the xy discretisation and 128 planes in the z direction. Then, moving averages of the instantaneous lift coefficient were calculated, considering different widths for the average. Figure B.4 shows the results of these averages, based on the final time and the average width. This can be interpreted as the result that would be obtained from a simulation until the final time represented in the x axis, with the last T_{ma} time units considered to obtain the result. The dotted lines in the plot represent 2% deviations from the value 0.76, which was considered as an acceptable error for the approximation of using a finite time. Based on this result, T_{ma} was chosen as 100 time units, since shorter series do not guarantee that the result remains within the error bounds. Also, the total simulation



Figure B.4: Moving average of the lift coefficient for different time-averaging intervals, described by the width of the interval T_{ma} and by the final time. The dotted lines represent 2% deviations from the value 0.76.

time was taken as 140, so that the first 40 time units, which might be affected by the initial condition, are not considered.

Although figure B.4 suggests that the results obtained through this averaging process are still very sensitive to the simulation time, this impression is a result of the scale employed. To clarify this point, figure B.5 compares the moving average with $T_{ma} = 100$ which was chosen to calculate the results, with the moving average with $T_{ma} = 0.01$, which is close to the instantaneous lift value. From this figure, it becomes clear that the oscillations in the average are very low when compared with the fluctuations of the lift coefficient. Also, although the discussion considered only the lift coefficient, the drag coefficient presents similar behaviour.

In order to determine the effect of the spatial resolution on the results, the polynomial order and the number of Fourier modes were reduced, independently. Table B.11 shows the effect on the forces of reducing the number of degrees of freedom in the z direction to 64, while table B.12 presents the results obtained by reducing the polynomial order to 9. It is clear that each of these changes lead to a significant change of the result, and therefore the original resolution was employed in the simulations. Although these results indicate that the simulations are not fully resolved, the level of error implied was considered acceptable for the purpose of evaluating the modifications in the flow arising from the use of waviness in the wing.

The further illustrate the resolution power of our simulations, figure B.6 shows the energy



Figure B.5: Comparison of the lift coefficient obtained using the chosen averaging interval to the result with a short interval, which is close to the instantaneous lift coefficient.

Table B.11: Convergence tests for number of Fourier modes in the span direction. Simulations with Re = 10,000.

N_z	$\overline{C_d}$	$\overline{C_l}$	$C_l^\prime(r.m.s)$
64	0.1990	0.6954	0.1255
128	0.2109	0.7569	0.1281

Table B.12: Convergence tests for basis functions polynomial order. Simulations with Re = 10,000.

_

Order	$\overline{C_d}$	$\overline{C_l}$	$C_l^\prime(r.m.s)$
9 10	0.2073	0.7327	0.1638



Figure B.6: Modal energy for simulation of baseline wing with Re = 10,000 and $\alpha = 18^{\circ}$.

contained in each Fourier mode of the spanwise expansion for a simulation of the baseline wing with $\alpha = 18^{\circ}$. We note that when the artificial dissipation of the spectral-vanishing viscosity becomes significant, around mode 40, the energy contained in each mode is already 4 orders of magnitude lower than the energy in the first mode. Note that mode zero, corresponding to the mean spanwise flow, is not shown in this figure. This result provides an idea of how well the Fourier expansion chosen to discretize the span direction resolves the flow is this case.

B.3 Tests for Re = 50,000

Although several tests were performed in order to obtain an appropriate setup for the simulations with Re = 50,000, these did not follow a systematic approach like the previous sections. This was in part due to the extremely high computational cost of these simulations. Also, the experimental data available in this case helped validating the numerical results without resorting to extensive convergence tests. In any case, some of the process followed for determining the simulation parameters in this case will now be briefly presented.

The simulations used a mesh with 4674 quadrilateral elements, with figure B.7 displaying the region close to the wing for this mesh. The mesh follows a similar structure to the one used for Re = 10,000, but is much more refined close to the wing. As explained in chapter 8, the averaging interval for this Reynolds number was chosen independently for each simulation, and therefore results for the convergence of the average forces will not be presented here.

In terms of spatial resolution, the polynomial order was varied using the adaptive procedure from section 5.3, and therefore the relevant parameter for determining the spatial resolution in xy is the maximum polynomial order P_{max} . All simulations used $P_{max} = 11$,



Figure B.7: Detail of the mesh used for simulations with Re = 50,000.



Figure B.8: Modal energy for simulations of baseline wing with Re = 50,000.

since for $\alpha = 15^{\circ}$ using $P_{max} = 9$ led to a lift coefficient which differed significantly from the experimental results. The spanwise resolution was fixed at 256 degrees of freedom per chord length. This is twice as much as what was employed for Re = 10,000, and seems to be sufficient for the intents of the present work. Figure B.8 shows the modal energy of the Fourier expansion for the simulations of the baseline wing performed with these parameters. There is a significant reduction in the energy along the resolved scales, with the energy dropping by 4 orders of magnitude for $\alpha = 15^{\circ}$.

For the largest simulation performed, which used $L_z = 2.0c$, a simulation with a constant polynomial order equal to P_{max} would have approximately 340 million degrees of freedom. Due to the adaptive polynomial order, the number of degrees of freedom in this simulation is actually around 175 million. This is still a very large amount, illustrating the challenges posed by this simulation.

C. Simulations of the Taylor-Green Vortex

In order to illustrate the effects of the dealiasing techniques and of the spectral-vanishing viscosity (SVV) on under and marginally resolved simulations, some results of simulations for the Taylor-Green Vortex will be briefly presented here. For dealiasing, the method of Kirby and Sherwin (2006a) was used in the xy plane with the 3/2 padding rule being used for the Fourier expansion, while SVV was described in section 5.1.

The Taylor-Green Vortex is a vortex breaking down into turbulent structures until it is eventually dissipated. The simulations presented here consist of solving the incompressible Navier-Stokes equations in a periodic box of side 2π with initial conditions

$$u = sin(x)cos(y)cos(z),$$

$$v = -cos(x)sin(y)sin(z),$$
 (C.1)

$$w = 0$$

using a Reynolds number $Re = \frac{1}{\nu} = 1600$.

All simulations used a Fourier expansion in the *z*-direction and a regular quadrilateral mesh with polynomial order P = 8 in the *xy* plane. Also, a 2^{nd} order time integration scheme was used with a time step 0.005. In all simulations using SVV the parameters were $P_{cut} = 0.7P$ and $\epsilon = 0.1$. Several simulations were performed for 64^3 degrees of freedom, considering several combinations of SVV and using or not dealiasing. An additional simulation was performed for 128^3 degrees of freedom using dealiasing without SVV, in order to serve as a reference of a better resolved result.

Figure C.1 presents the results of the enstrophy for all simulations considered here. The first thing we notice is that without dealiasing or SVV the simulation with 64³ degrees of freedom is unstable, justifying the need for a stabilization technique. Also, dealiasing does not degrade the accuracy of the solution. In fact, since it is a consistent integration of the non-linear terms, it is likely to improve the accuracy. This is a clear contrast to SVV, which degrades the accuracy by introducing an artificial dissipation into the solution. However, although in this particular case dealiasing alone was capable of leading to a stable solution, this is not always the case. Therefore, in general SVV is still necessary in marginally resolved simulations.



Figure C.1: Enstrophy for the Taylor-Green Vortex simulations with different configurations.

To illustrate the previous results, figure C.2 presents iso-surfaces of Q-criterion for some of the simulations. Although the simulation with dealiasing and without SVV is clearly less accurate than the higher resolution simulation, it is able to resolve more of the small scale structures when compared to the simulation using SVV, where these structures are dissipated by the artificial viscosity.

Considering the previous discussion, we can conclude that although the dealiasing techniques alone are not always capable of stabilizing the solution, they are important for two reasons. First, by consistently integrating the non-linear terms, they potentially improve the accuracy of the solution. Second, they allow us to reduce the amount of SVV used in the simulation. This is advantageous, since SVV can have a significant impact on the accuracy of under resolved simulations, and therefore should be used with moderation.



(a) $128^{3},$ with dealiasing, without $\ensuremath{\mathsf{SVV}}$



(b) 64^3 , with dealiasing, without SVV



(c) 64^3 , without dealiasing, with SVV

Figure C.2: Iso-surfaces of the second invariant of the velocity gradient tensor Q = 10 coloured by u velocity for simulations of the Taylor-Green Vortex at t = 10. Only one eighth of the domain is displayed.