

University of Southampton
Faculty of Engineering, Science and Mathematics
School of Engineering Sciences

A higher-order potential flow method
for thick bodies, thin surfaces and wakes

Daniel Joseph Bernasconi

Thesis for the degree of Doctor of Philosophy

September 2007

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS
SCHOOL OF ENGINEERING SCIENCES

Doctor of Philosophy

A HIGHER-ORDER POTENTIAL FLOW METHOD
FOR THICK BODIES, THIN SURFACES AND WAKES

by Daniel Joseph Bernasconi

A higher-order method is developed that models continuous source and doublet singularity distributions over three-dimensional curved surfaces. The singular on-surface influence coefficients are treated by a robust desingularisation algorithm, whereas off-surface coefficients are calculated by means of an efficient subdivision and variable cubature scheme. Whilst higher-order methods have previously been developed for thick bodies and Dirichlet boundary conditions, this method is also capable of modelling continuous geometry and singularity surfaces over thin bodies and wakes that require Neumann boundary conditions.

The Continuous Surface Method (CSM) has a number of advantages over conventional constant panel methods (CPMs). Firstly, as curved geometries are represented exactly, changing the order of the solution does not modify the physical shape of the configuration. Furthermore, as singularity solutions are continuous, the significant grid-dependency of CPMs does not arise. Finally, the continuous singularity distributions allow velocities to be evaluated accurately across the entire surface without interpolation: this enables the calculation of continuous pressure distributions and the construction of streamlines and wakes flowing very close to surfaces, without any problems of divergence.

Numerical results comparing the CSM to a CPM have shown that for equal run times, the CSM obtains greater accuracy in pressure distributions than a CPM, and produces much smoother velocity fields. However the CSM was not able to improve upon the efficiency of the CPM in determining total aerodynamic forces.

A wake relaxation scheme in which wakes are modelled as curved B-spline patches is developed, and is convergent for simple geometries. For a more complex example of wakes shed from two closely overlapping sails, the wake relaxation converges to within around 0.5% of total aerodynamic load, but the low panel resolutions employed in the CSM are insufficient to model the detail of the wake roll-up effectively. Three alternative schemes to address this problem are evaluated.

Contents

Nomenclature	viii
1 Introduction	1
1.1 Background	1
1.2 Motivation	2
1.3 Structure of the thesis	3
1.4 Preliminary work on the dynamics of yacht sails	4
2 Higher-order potential flow methods	6
2.1 Doublet – Vortex line equivalence	7
2.2 Review of existing methods	9
2.2.1 The Hess I code	9
2.2.2 The Hess II code	10
2.2.3 The Woodward II code	12
2.2.4 Johnson’s code	13
2.2.5 The PAN AIR code	14
2.2.6 Bi-linear vortex triangle codes	16
2.2.7 Horstmann’s code	17
2.2.8 Maniar’s code	19
2.2.9 Summary of methods reviewed	20
2.3 Evaluation study	20
2.3.1 Johnson’s bi-quadratic doublet panels	21
2.3.2 Mracek’s bi-linear vortex triangles	26
2.3.3 Comparison of computation times	32
2.4 Conclusions	33
3 The Continuous Surface Method	35
3.1 Problem formulation	35
3.1.1 Surface boundary conditions	37
3.1.2 Collocation point formulation	41
3.1.3 Galerkin formulation	42
3.1.4 Comparison of Collocation and Galerkin formulations	42
3.1.5 Leading edge boundary conditions	43
3.1.6 Trailing edges and wakes	48

3.1.7	Boundary conditions on other edges	51
3.1.8	The concept of a patch	52
3.1.9	Singularity distributions	54
3.2	Influence functions	56
3.2.1	Velocity influences due to point doublets and sources	56
3.2.2	Kernel function singularities	57
3.2.3	Integration of three dimensional hypersingular kernel functions	59
3.2.4	Integration of two dimensional hypersingular kernel functions	64
3.2.5	Less than hypersingular integrals	68
3.2.6	Free-terms	69
3.2.7	Unknown singularity distributions	73
3.2.8	Alternative integration of the doublet velocity kernel	73
3.3	Singularity distributions represented as B-Spline surfaces	75
3.3.1	Bi-cubic B-spline basis functions	75
3.3.2	Modified basis functions	77
3.3.3	Panels and super-panels	78
3.3.4	Off-surface influence evaluation	79
3.3.5	On-surface influence evaluation	84
3.3.6	Influence coefficient validation	88
3.3.7	Re-use of panel integration points	88
3.3.8	Problem configuration for bodies and wakes	90
3.4	Implementation details	93
3.4.1	Matrix assembly and solution	93
3.4.2	Main algorithm	94
3.4.3	Efficiency improvements	94
3.5	Post-processing of velocities, pressures and forces	99
3.5.1	Surface doublet gradients	99
3.5.2	Field point velocities	100
3.5.3	Surface velocities	101
3.5.4	Surface pressures and forces	103
3.5.5	Leading edge suction forces	103
3.5.6	Forces calculated in the Trefftz plane	104
4	Results and analysis	106
4.1	Parameter studies	106
4.1.1	Edge suction forces via a square-root fit	109
4.1.2	Panel spacing	112
4.1.3	Panel aspect ratio	115
4.1.4	Collocation point positions	116
4.1.5	Basis Type	119
4.2	Optimal configuration	126
4.3	Computation time	126

4.4	Validation of the CSM	128
4.5	Comparison of CSM and Vortex Lattice results	133
4.5.1	Basis for comparison	134
4.5.2	Run time and condition number	134
4.5.3	Lift and drag coefficients	135
4.5.4	Surface pressures	139
4.5.5	On-surface velocities	139
4.5.6	Off-surface velocities	141
4.5.7	Summary of VLM and CSM comparison	143
5	Wake relaxation	145
5.1	The requirement for relaxing the wake	145
5.2	Wake relaxation principles for the CSM	146
5.2.1	Wake sample points	146
5.2.2	Staged relaxation process	146
5.2.3	First column of wake panels	147
5.2.4	Far-field wake extension	149
5.2.5	Geometry constraints	149
5.2.6	Velocity sampling	150
5.2.7	Surface fitting	151
5.2.8	Singularity distributions during wake relaxation	151
5.2.9	Summary	152
5.3	Example of wake relaxation for a flat inclined plate	152
5.4	Example of wake relaxation for two overlapping sails	154
5.4.1	Pressure and force variation with wake relaxation	156
5.5	Alternative wake relaxation schemes	157
5.5.1	Use of wake cross-sections	159
5.5.2	Knot merging	159
5.5.3	Truncation and redistribution	159
5.5.4	Framework shape functions	160
5.5.5	Literature on vortex modelling	161
6	Conclusions	162
6.1	Results with fixed geometries	163
6.2	Wake relaxation	164
A	Power series expansions for three dimensional influence coefficients	165
A.1	Series expansions for basic functions	165
A.1.1	Expansion of powers of r	165
A.1.2	Expansion of the local Jacobian vector	166
A.1.3	Expansion of the singularity distribution	167
A.2	Kernel function expansions	168
A.2.1	Potential influence of a doublet distribution	168
A.2.2	Potential influence of a source distribution	168

A.2.3	Velocity influence of a doublet distribution	169
A.2.4	Velocity influence of a source distribution	170
A.3	Expansion of integration limit α	170
B	Power series expansions for two dimensional influence coefficients	172
B.1	Series expansions for basic functions	172
B.1.1	Expansion of powers of r	172
B.1.2	Expansion of the local Jacobian vector	173
B.1.3	Expansion of the singularity distribution	174
B.2	Kernel function expansions	174
B.2.1	Potential influence of a doublet distribution	174
B.2.2	Potential influence of a source distribution	175
B.2.3	Velocity influence of a doublet distribution	176
B.2.4	Velocity influence of a source distribution	176
B.3	Expansion of integration limit α	177
C	Vortex lattice parameter study	179
C.1	Effect of configuration on edge suction forces	181
C.2	Effect of configuration on sample point pressures	182
C.3	Effect of configuration on lift and drag coefficients	183
C.4	Variation of computation time with number of panels	185
C.5	Conclusions	185
D	Alternative wake relaxation schemes	188
D.1	Studies on a two dimensional cross-section	188
D.1.1	Basic two dimensional wake relaxation scheme	188
D.1.2	Limitations of a two dimensional model	189
D.1.3	Example of basic two dimensional wake relaxation	190
D.2	Knot merging	192
D.2.1	The two dimensional scheme	192
D.2.2	Results in two dimensions	194
D.2.3	Extension to three dimensions	196
D.3	Truncation and redistribution	196
D.3.1	The two dimensional scheme	197
D.3.2	Results in two dimensions	198
D.3.3	Extension to three dimensions	199
D.4	Framework shape functions	201
D.4.1	Implementation in two dimensions	202
D.4.2	Shape functions to model roll-up	203
D.4.3	Results in two dimensions	203
D.4.4	Extension to three dimensions	204
	Bibliography	205

Nomenclature

Greek symbols

$\alpha(\epsilon, \theta)$	Value of polar coordinate ρ on boundary of s_ϵ in 3D
$\alpha_1(\epsilon), \alpha_2(\epsilon)$	Values of ρ on left and right boundaries of s_ϵ in 2D
$\beta(\theta), \beta_m(\eta)$	Constant expansion coefficient of α in 3D; in 2D
$\gamma(\theta), \gamma_m(\eta)$	Linear expansion coefficient of α in 3D; in 2D
ϵ	Radius of spherical (3D) or circular (2D) area s_ϵ
ζ_1, ζ_2	Left and right limits of integration in 2D intrinsic coordinates
$\boldsymbol{\eta} = (\eta_1, \eta_2), \eta$	Image of \mathbf{y} in intrinsic coordinates for 3D; for 2D
θ	Intrinsic polar coordinate angle
κ	Power of ξ in B-spline basis function
λ	Curvilinear distance along wake cross section
$\mu(\mathbf{x}), \mu_0, \mu_1(\theta)$	Doublet strength; constant and linear terms in series expansion
$\boldsymbol{\xi} = (\xi_1, \xi_2), \xi$	Image of \mathbf{x} in intrinsic coordinates for 3D; for 2D
ρ	Intrinsic polar coordinate magnitude
$\bar{\rho}(\theta)$	Value of ρ at edge of area of integration
$\sigma(\mathbf{x})$	Source strength
$\tau(\mathbf{x})$	Singularity strength (either source or doublet)
$\phi^T(\mathbf{y}), \phi^\infty, \phi^P$	Total, free-stream and perturbation velocity potentials
ψ	Polar coordinate defining Archimedean spiral

Subscripts

b	Order of B-spline basis function
$(f, g), h$	Collocation or wake point index (either as 2D array or 1D index)
$(i, j), k$	Super-panel indices (referenced either as 2D array or 1D index)
$(p, q), r$	Panel indices (referenced either as 2D array or 1D index)
(m, n)	Quadrature point index

Roman symbols

$a(\mathbf{y})$	Free-term coefficient for hypersingular kernel functions
A, A	First order term in polar expansion of \mathbf{r} ; and its magnitude
$b(\mathbf{y})$	Free-term coefficient (always cancels in final integration)
B, B	Second order term in polar expansion of \mathbf{r} ; and its magnitude
B_{ij}	Weight (singularity strength) for super-panel ij
$c(\mathbf{y})$	Free-term coefficient for all kernel functions
C	Leading edge suction force coefficient
$e(\epsilon)$	Region around collocation point removed from integration
\mathbf{F}, F	Force vector; its magnitude
$\mathbf{F}(\rho, \theta)$	Product of kernel, singularity distribution and Jacobian
$\mathbf{F}_{-1}(\theta), \mathbf{F}_{-2}(\theta)$	Strongly singular/hypersingular components of $\mathbf{F}(\rho, \theta)$
$\mathbf{F}^1(\eta, \xi), \mathbf{F}^2(\eta, \xi)$	Value of $\mathbf{F}(\eta, \xi)$ either side of singularity in 2D analyses
\mathbf{I}_0	Weakly- and non-singular components of influence integral
$\mathbf{I}_{-1}, \mathbf{I}_{-2}$	Strongly/hypersingular components of influence integral
$\mathbf{J}(\boldsymbol{\xi}), \mathbf{J}(\boldsymbol{\xi})$	Jacobian determinant; multiplied by surface normal $\mathbf{n}(\boldsymbol{\xi})$
$\mathbf{J}_0, \mathbf{J}_1$	Constant and linear polar expansion coefficients of $\mathbf{J}(\boldsymbol{\eta})$
$\mathbf{n}(\mathbf{x})$	Surface normal unit vector (points into thick bodies)
N	Quantity of items: e.g. N_r number of panels
$p(\mathbf{y}_{\pm})$	Pressure on either side of surface
p_0	Static pressure
P	Panel
$\mathbf{q}(\mathbf{y})$	Potential or velocity influence
Q^{∞}	Free-stream velocity
R	Surface in intrinsic coordinate space
$\mathbf{r}(\mathbf{y}, \mathbf{x}), r(\mathbf{y}, \mathbf{x})$	Vector from \mathbf{y} to \mathbf{x} ; magnitude of \mathbf{r}
S, S_B	Super-panel (or general surface); surface enclosing a body
s	General scalar distance between two points
s_{ϵ}	Small circular or spherical region around collocation point
t	Time for 2D wake relaxations
$\mathbf{T}(\mathbf{y}, \mathbf{x})$	Kernel function
$U(\xi)$	B-spline basis function
\mathbf{v}	Wake sample point
w	Quadrature point weight
$\mathbf{x}(\xi)$	Point on singularity surface (usually a variable of integration)
$\mathbf{y}(\xi)$	Collocation, Galerkin integration, or velocity sample point
\mathbf{Z}, Z_i	Knot vector; element i of knot vector

Chapter 1

Introduction

1.1 Background

Boundary element potential flow methods have been used extensively within the aeronautical and marine industries since the pioneering work of Hess & Smith (1962). Commonly known as Panel Methods, their use remains widespread despite the introduction of more sophisticated techniques such as Reynolds-Averaged Navier Stokes (RANS) solvers. Their longevity must partly be a result of their simplicity: for the many classes of physical problems that approximate potential flow, a constant panel method (CPM) can produce a remarkably accurate solution (James, 1972). They are also extremely efficient: the flow-field around a complex geometry, modelled by hundreds of panels, can be solved by an average PC in a few seconds.

It is expected that as computing power relentlessly increases, potential flow methods will ultimately give way to more sophisticated codes capable of modelling greater realism with turbulence and boundary layers. However, particularly in situations where run-time is paramount, it will be many years before panel methods are abandoned. Examples of run-time constraints include real-time applications, and design optimisations which may require tens of thousands of iterations. In these cases, it is simply not feasible to use RANS simulations.

Despite their advantages of simplicity and efficiency, constant panel methods are, however, not without their problems. They exhibit a number of undesirable features that result from the discretisation of both their geometry and their singularity distributions. Most significantly, these are grid layout dependence¹,

¹The distinction between grid layout and grid density is important: a good scheme will provide a solution that converges as the grid density increases, and that will converge to the same solution irrespective of the particular panel arrangement chosen.

singularities in surface velocity fields, and the requirement of interpolation to obtain pressure distributions and loads. In situations with fixed geometries these characteristics are not a major concern, and for the analysis of rigid bodies with known wake shapes it is hard to improve on the efficiency of a CPM. However, analysis in more demanding fluid-structure interaction problems can be limited by the issues described. In particular, accurate relaxed wake shapes around closely interacting bodies can be difficult to obtain with conventional CPMs.

1.2 Motivation

The aim of this work is to develop a potential flow method that eliminates many of the problems associated with constant panel methods. In particular, a method that reduces grid dependence and singularities in velocity sampling is required. As will be seen in the following literature review, this is by no means the first time that this goal has been sought. One of the first higher-order methods (Woodward, 1973) appears in the literature only a decade after the introduction of the CPM, and since then a significant number of higher-order panel methods have been devised. These have mainly been evolutionary steps from constant panel methods, increasing the order of geometry or singularity distributions by one or two degrees whilst maintaining the concept of discrete panels. A few of these (e.g. Epton & Magnus, 1990) have enjoyed considerable patronage within industry. However, none of the early methods completely eradicated the problems of CPMs; they merely reduced them to a lower order.

More recently though, some work at MIT introduced the concept of using continuous surfaces to model geometry and singularity distributions. In particular, Maniar (1995)² developed a solution based on B-splines which provided much of the inspiration for the method described in this thesis. However, Maniar's method could only deal with thick bodies, and did not account for wakes. A subsequent MIT method by Lee & Kerwin (2003) also employed B-

²Despite the novelty and elegance of this method, it appears to be unpublished except for this PhD thesis.

splines, this time in two dimensions, but again did not allow for the analysis of thin bodies or wakes.

It may be that thin bodies and wakes have not been the subject of the continuous surface treatment because of the difficulty in forming their influence coefficients. Thick body potential flow problems are usually solved by means of Dirichlet boundary conditions, which require potential influences to be evaluated at collocation or Galerkin integration points. These influences are singular when the collocation point lies on the singularity surface, but the singularity is ‘weak’ (of order $1/r$ in three dimensions), and can be treated by a transformation to polar coordinates. In contrast, the *velocity* influence on a thin surface containing a doublet distribution, required by a Neumann formulation, is ‘hypersingular’ (of order $1/r^3$ in three-dimensions). Hypersingular functions are, in general, only integrable in the sense of Hadamard finite parts (Hadamard, 1952). Analytical boundary element formulations exist for some special cases, such as a constant doublet quadrilateral panel, but until recently no general method of solution was available. However, between 1987 and 1995, Guiggiani and his colleagues at Universita di Pisa published a series of papers in the solid mechanics literature (Guiggiani *et al.*, 1990; Guiggiani, 1995, 1998) that provided a semi-analytical method for treating the hypersingularity. Within the field of solid mechanics this method has become well established, and is used in commercial boundary element codes. Surprisingly though, despite the direct applicability of this technique to potential flow, it seems to have gone unmentioned in the fluid mechanics literature. It is the hypersingular integral treatment of Guiggiani combined with some of the modelling concepts from Maniar’s B-spline thick body method that have provided the basis of the new continuous surface method (CSM) described in this thesis.

1.3 Structure of the thesis

To set the new Continuous Surface Method in context, a comprehensive review of existing higher-order schemes is presented in chapter 2, and a number of preliminary studies into alternatives are described.

Chapter 3 provides a complete description of the new Continuous Surface Method, and shows how influence coefficients are determined. Some details of the influence coefficient calculation are contained within appendices A and B, for three and two dimensional situations respectively.

Chapter 4 begins with a detailed parameter study to define an optimum configuration for the CSM, investigating factors such as panel layouts and collocation point positions. Results from the CSM are then compared with those from a conventional vortex lattice method, both in terms of accuracy and efficiency. In order to provide a fair comparison, the configuration of the vortex lattice method is also subject to an optimisation process, which is detailed in appendix C.

Chapter 5 details a wake relaxation strategy for the CSM, and provides examples of wake relaxations for the simple case of a flat inclined plate, and the more complex configuration of two overlapping sails. Some problems relating to wake relaxations are discussed, and three alternative schemes to combat these issues are outlined, with further details provided in appendix D.

Chapter 6 summarises the development of the CSM and discusses the benefits of using this method as an alternative to conventional constant panel methods.

1.4 Preliminary work on the dynamics of yacht sails

The original aim of this research project, which was undertaken in association with the America's Cup team Alinghi, was to construct a system to investigate the dynamic behaviour of yacht sails. This was to be based primarily on conventional panel methods, and towards this goal, a dynamic vortex lattice code (DVLC) was created. The development and validation of the DVLC is described in a separate technical report (Bernasconi, 2007), which also contains a series of parameter studies and preliminary results.

Some way into the project, the author began to experience convergence issues whilst modelling wake relaxations on sails with the dynamic code; at around the same time, Alinghi were having similar problems with their steady-state panel code, partly as a result of sail plans becoming ever more closely overlapped.

With these issues becoming more pressing than the original desire to investigate dynamic sail behaviour, it was agreed to change direction to research higher-order methods, and ultimately to develop the CSM. However, many aspects of the DVLC are relevant to the CSM's development, and the vortex lattice code is used as a benchmark by which to compare the efficiency and accuracy of the CSM.

Chapter 2

Higher-order potential flow methods

The discretisation of surfaces into constant panels or concentrated vortex lines leads to irregular induced velocity fields close to vortex ring surfaces. This in turn can cause irregularities in the development of the wakes as they propagate downstream, as highlighted in figure 2.1. Whilst some strategies to overcome these problems are investigated and discussed by Bernasconi (2007, chapter 4), it is shown that it is difficult to eradicate them without sacrificing accuracy. Rather than managing the difficulties created by the concentration of vorticity into vortex lines, an alternative solution is to use a different fundamental panel type.

Constant panel and vortex ring formulations also have short-comings in terms of solution accuracy when used for the modelling of curved surfaces. Because the models comprise discrete flat or skewed panels over a pre-defined structured grid, the flow solution can be strongly influenced by the shape and alignment of the grid. Therefore, when comparing solutions for two different surface shapes, care must be taken to ensure that differences in the solution reflect genuine aerodynamic differences of the surfaces, and are not artefacts relating to the choice of grid shape or alignment. Further, since the constant panels or vortex rings have straight edges, they cannot accurately represent curved surfaces. This means that as the number of elements used to model a surface increases, the solution convergence is much slower than would be seen when modelling a flat plate (Bernasconi, 2007, chapter 5), or when using curved elements to model a curved surface.

In this chapter, alternative potential flow discretisations are reviewed with particular emphasis on their abilities to solve the above problems.

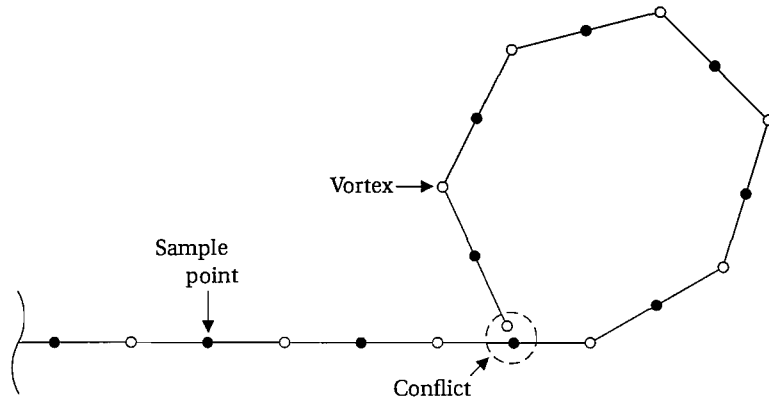


Figure 2.1: A cross-section of a typical wake roll-up within a vortex-lattice method. The wake is developed iteratively by evaluating induced velocities at discrete sample points, and relaxing the geometry accordingly. As the wake develops, vortex elements at the edges of the wake may be relaxed such that they come close other sections of the geometry. In the situation shown, the singular velocity influence around the edge vortex will exert an extremely strong and physically incorrect influence on the adjacent sample point velocity, often causing irrecoverable divergence of the process.

2.1 Doublet – Vortex line equivalence

Within the family of discrete element potential flow models, constant panel and vortex ring schemes are practically the simplest, or lowest-order. The alternative schemes considered will therefore by necessity be higher-order, where in this context the order is that of the polynomial approximation by which continuous vortex, source or doublet distributions are approximated.

Throughout this chapter, reference will be made to the equivalence between a doublet strength distribution and a distribution of vortex lines, particularly for the purposes of comparing doublet schemes with vortex line schemes. This equivalence relates a general surface doublet distribution to a corresponding surface vortex distribution

$$\mathbf{q} = -\frac{1}{4\pi} \int_S (\mathbf{n} \times \nabla \mu) \times \frac{\mathbf{r}}{r^3} dS + \frac{1}{4\pi} \int_C \mu \frac{d\mathbf{l} \times \mathbf{r}}{r^3}$$

whose order is one less than the order of the doublet distribution plus a vortex ring whose strength is equal to the edge value of the doublet distribution (Hess, 1972). In practice, the edge value of a doublet distribution over a wing is zero along edges to which no wake is attached, and is continuous where a wake is

attached, so the second term will generally be zero. However, some higher-order schemes contain doublet distributions that are discontinuous between panels, and this discontinuity is equivalent to a concentration of vorticity.

Figure 2.2 gives examples of the doublet strength-vorticity equivalence, showing some of the schemes referred to in Section 2.2. In these plots, a (2×2) grid of panels is shown, and the height of the surfaces above the base plane (i.e. the z -axis) represents the doublet strength. Red and green lines indicate an equivalent vortex line set: red lines indicate the strengths of concentrated vortex lines which lie along panel boundaries; green lines indicate sheets of infinitesimal vortex lines running over panels.

Case (a) shows the constant doublet strength equivalence of the standard vortex ring construction; here, the steps in doublet strength between the panels form the concentrated vortex rings. In cases (b) and (c), the linear and quadratic doublet strengths along one axis are equivalent to sheets of parallel vortex lines, but concentrated vortex lines segment the other axis. Case (d) shows a continuous bi-linear doublet strength distribution created by triangular panel elements.

These diagrams, and in particular case (d), show that for a continuous distribution, vortex lines form closed ‘contours’ of doublet strength. The greater the doublet strength gradient, the greater the intensity of the vortex lines (or the closer they become in the discrete representation of the diagrams). A step-change in doublet strength concentrates many infinitesimal vortex lines into a single line of finite vorticity.

The equivalence between doublet and vorticity distributions is referred to surprisingly little in the literature. Panel influence coefficients derived from the integration of vortex lines appear to be simpler to compute than those derived from the integration of doublet singularities: for example, compare Epton & Magnus (1990) with Mracek *et al.* (1992). In contrast, the construction of the potential flow problem to satisfy the boundary conditions is more naturally stated in terms doublet strength distributions. This might suggest that vorticity derived panels would commonly be used within equivalent doublet based solutions, but this does not appear to be the case: nearly all of the reviewed literature uses

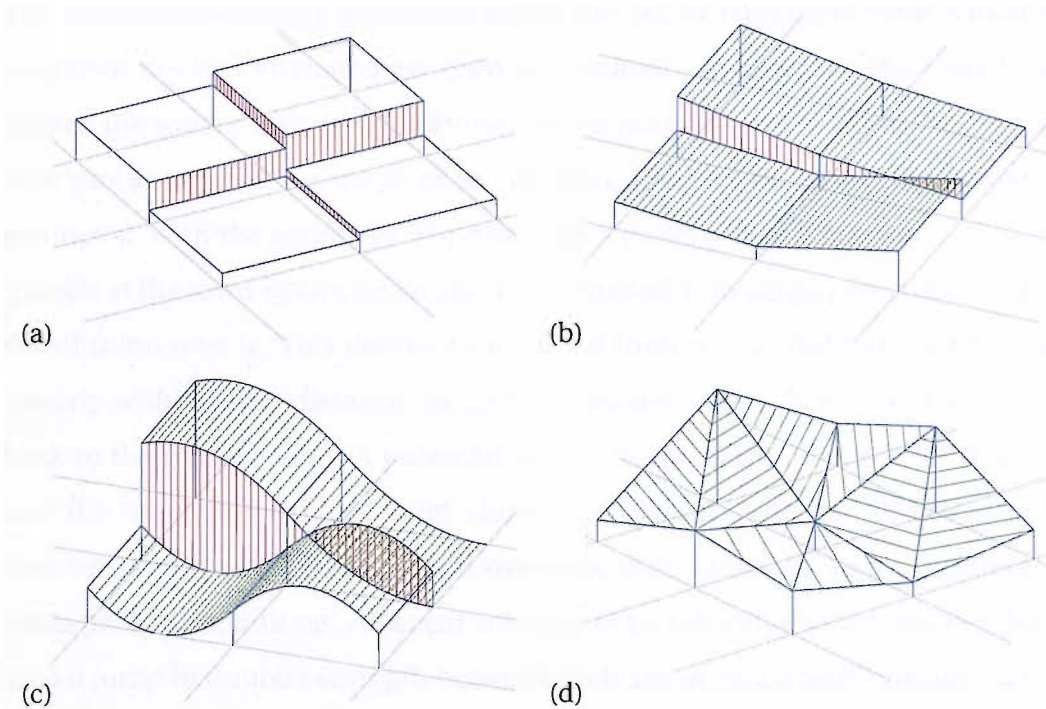


Figure 2.2: Examples of doublet strength-vorticity equivalence: (a) Constant-Constant doublet strength (vortex lattice); (b) Constant-Linear; (c) Constant-Quadratic; (d) Linear-Linear

the same singularity type for both the influence coefficient derivation and the solution mechanism.

2.2 Review of existing methods

2.2.1 The Hess I code

The first widely used three-dimensional panel code was developed by Hess & Smith (1962) of what was then the Douglas-McDonnell Aircraft Company. Known as the ‘Douglas-Neumann’ code, this was a non-lifting method with only constant source elements on flat panels. A lifting capability was later added by the introduction of doublet panels (Hess, 1972), and this program became an industry standard known as the Hess I code.

The majority of thick-body source/doublet panel codes in use today treat the doublet distribution as the fundamental unknown, and the source strength is set by the relationship between the local geometry and the free-stream flow.

The Dirichlet boundary condition forms the set of equations from which the unknown doublet strengths are then determined. In contrast, the Hess I code treated the source strength distribution as the principal unknown of the potential flow problem, with the shape of the doublet distribution being determined by geometry. With the geometry discretised by a panel grid, each column of panels (panels at the same spanwise location) was treated as having a prescribed doublet distribution over it. This distribution started from zero at the trailing edge, rose linearly with surface distance towards the leading edge, then continued to rise back to the trailing edge. A potential jump was thus formed at the trailing edge and the magnitude of this jump alone controlled the doublet distribution for that set of panels. These magnitudes were then found by use of an iterative Kutta pressure condition. Adjacent columns of panels either acted independently with a jump in doublet strength between each set, or had a quadratically varying doublet strength in the span direction with the variation being controlled by a quadratic 'stencil' acting over adjacent potential jumps at the trailing edge.

The source distribution was prescribed to be constant over each panel, and the velocity influences of the unknown source and doublet panels was determined at panel collocation points. With the doublet distribution controlled by the Kutta condition, a Neumann boundary condition applied at each collocation point provided the system of equations from which the unknown panel strengths were found. In summary then, the Hess I code employed flat quadrilateral panels with constant source strengths and either linear-constant or linear-quadratic doublet strengths, but with a very prescribed doublet strength variation.

2.2.2 The Hess II code

A higher-order version of the Hess I code was later developed, and became known as the Hess II code. This is described in a report by Hess (1979), with further details of implementation given by Hess & Friedman (1981).

At the time the Hess II code was introduced, potential flow problems in industry were being solved with over 3,000 panels, and the manual effort

in creating the panel grids was becoming very significant. This prompted the development of a new automatic panel generation code, introduced in McDonnell Douglas at around the same time, which became a significant development in its own right. As such, speed of computation of the panel codes remained a significant factor.

The aim in developing the new code was an “increase in speed of computation with no decrease in accuracy, or better still, with an increase in accuracy”, and this was achieved by augmenting the order of both the geometry and the source distribution. The Hess II code maintained the concept of the Hess I code, with source distributions forming the principal unknowns, but now provided paraboloidal panels (with small normal deviations from a flat panel) containing linear source and linear vortex distributions.

Hess would have preferred to employ a quadratic doublet strength, rather than a linear vortex strength, but he concluded that there was no analytic expression for the velocity influence of a quadratic doublet variation over a paraboloidal panel. These distributions are in fact equivalent to linear vortex distributions, except for a concentrated vortex filament around the edge of the equivalent vortex-based panel, and it is these filaments that prevent analytic integration. For a network of panels forming continuous geometric and singularity surfaces this would not present a problem, as the vortex filaments on neighbouring panels would cancel out; however, Hess’ formulation created small gaps between panels, such that these terms did not in fact disappear. Because they could not be integrated, the influence of these vortex filament terms was discarded, and Hess argued that the error induced by this simplification was no greater than the errors in approximating the solution with paraboloidal panels and linearly varying source strengths. However, the inconsistency in the approach implies that the condition of zero divergence of vorticity is not strictly met, and the resulting velocity fields would be non-physical.

Nevertheless, Hess reported that the new code did result in a significant increase in accuracy compared to the Hess I code, with no increase in computing time for the same number of panels. Given the significant extra complexity of the influence coefficients compared to the flat panels, the statement that run-

time was not increased is somewhat surprising, and it is likely that a number of efficiency gains were made in the Hess II code that could have been, but were not implemented in the Hess I code. It is also probable that on 1970's computers with slow memory access, the overall run-time was more significantly influenced by the matrix solution than by influence coefficient computation.

2.2.3 The Woodward II code

Alongside the Hess II code, the Woodward II code (Woodward, 1973) is one of the earliest published examples of a higher-order method. As a successor to the Woodward I code (Woodward, 1968), it was developed for NASA to analyse aircraft wings and bodies in combination. Constant strength source panels model non-lifting bodies, and linear vortex panels provide two options for modelling lifting surfaces:

- Planar boundary conditions: linearly varying source and vortex panels located on the mean plane of the wing are used to model the wing's thickness, camber, twist and incidence.
- Non-planar boundary conditions: linear vortex panels are placed on the upper and lower surfaces of the wing to simulate both lift and thickness effects, with no source panels. Collocation points are generally located at panel centroids, and Neumann boundary conditions are applied.

Woodward's vortex panels are required to be flat and trapezoidal in shape, with two edges parallel to the streamwise direction. Each panel comprises a sheet of bound spanwise vortex lines that vary linearly in strength line-to-line. For rectangular panels, the only streamwise vortex lines are those shed at the panel edges, and because the bound vortex lines vary linearly, these panel-edge streamwise lines have a quadratic strength. For tapered panels, the bound vortex lines must also shed streamwise vortex lines across the panel to avoid a concentration of vorticity as the panel narrows.

The vorticity distribution on rectangular panels is equivalent to a doublet distribution that varies quadratically in strength in the streamwise direction

and has constant strength in the spanwise direction (see Fig 2.2, case 'c'). For tapered panels, this constant-quadratic doublet surface may be tilted to provide a constant gradient perpendicular to the panel leading and trailing edges.

Induced velocities are derived analytically by first considering triangular panels, then subtracting the effect of one triangular panel from another to obtain a flat trapezoidal panel.

In constructing the system of equations to model the potential flow problem, the leading edge of each panel is assigned a (constant) unknown vortex strength. The vortex strength then varies linearly to the panel's trailing edge, whose vorticity matches that of the following panel's leading edge. The last streamwise row of panels has a trailing edge vorticity set to zero to satisfy the Kutta condition. The resulting expressions are relatively simple, and the method is likely to be quick in comparison to more complex methods described below.

The code developed by Woodward was released as *USSAERO* (Unified Subsonic and Supersonic Aerodynamic Analysis code), a more recent revision of which is described by Weise (1986).

2.2.4 Johnson's code

Johnson (1980) describes a higher-order method in which curved quadrilateral panels carry linearly varying source or quadratically varying doublet distributions. Both Neumann and Dirichlet boundary conditions are catered for, and a number of different schemes for placing collocation points are provided.

This code was developed as a successor to the original Boeing A-230 code (Rubbert & Saaris, 1972), which whilst extremely versatile for its time, required significant manual effort and expertise in order to generate a suitable panel layout. By introducing curved panels and higher-order singularity distributions, Johnson's code provided solutions which were much less sensitive to the choice of grid, and as such, it could allow automatic panelling.

In Johnson's code, each curved panel is formed by fitting a paraboloid to known corner points in an immediate neighbourhood by the method of least squares. Unlike straight edged panels between common corner points, there will

generally be small gaps between paraboloid panels. To determine the potential flow solution, discrete values of singularity strength are assigned as unknown singularity parameters, typically at panel centres. The method of least squares is again used to fit linear (source) or quadratic (doublet) distributions to the panels, creating *shape functions* to distribute the unknown centre-point values. Like the geometric gaps between panels, this approach produces small jumps in potential strength along panel borders. Johnson reports that insensitivity of results to the choice of panel layout is a major benefit of his scheme, and this is demonstrated by producing accurate results from surfaces that have been ‘randomly’ and irregularly panelled.

The derivation of the panels’ induced velocities is extremely complex, and although it does not resort to direct numerical integration, it does rely on linear recursion relations to evaluate integral expansions. Whilst the computation steps needed are clearly presented, the large number of intermediate calculations required are computationally quite expensive.

2.2.5 The PAN AIR code

The PAN AIR (Panel Aerodynamics) code (Epton & Magnus, 1990) was developed by Boeing for NASA in around 1980, and whilst capable of solving subsonic flow, was written with the harder problem of supersonic flow in mind. Whilst small geometric gaps, and small jumps in singularity strength between panels (in a method such as Johnson’s) can be insignificant for subsonic flow, they can present a much greater problem for supersonic flow. This is because disturbances caused by doublet discontinuities do not diminish with distance in supersonic flow, which is governed by the wave equation, in contrast to subsonic flow governed by Laplace’s equation. The elements that distinguish PAN AIR from previous higher-order codes are thus continuity both in geometry and singularity strength between panels.

Whilst the aerodynamics of interest in this thesis lie well within the subsonic domain, the features of some codes that present problems for supersonic flow are the same features that can cause problems in wake propagation: namely,

doublet strength discontinuities. It is possible therefore that a solution to the wake problem may be found by looking for solutions to the supersonic flow problem.

Within PAN AIR, continuous geometry is achieved by implementing piecewise flat quadrilateral panels, where each panel comprises a sub-quadrilateral whose vertices are the mid-points of the master panel's edges, and four triangles in the remainder. The central sub-quadrilateral is planar, but for the purposes of assigning singularity distributions is itself sub-divided into four triangles with a common central vertex.

Over these piecewise flat panels, PAN AIR implements a well-conceived scheme to ensure that source and doublet distributions are continuous across panel edges (in addition to a scheme to ensure the distributions are continuous across panel network boundaries). Like Johnson's code, PAN AIR treats as its fundamental unknowns the values of source or doublet strength at panel centres, such that there are approximately the same number of unknowns as panels. Johnson's method then defines the linear or quadratic variation in source or doublet strength on a given panel as a surface-fit to the centre point values of the current panel and its neighbours. Because two adjacent panels will share some neighbours, but also have some different ones, the source or doublet distribution along the edge joining the two panels will generally not match exactly. PAN AIR overcomes this problem by introducing an intermediate step, in which the values of singularity strength at panel vertices and mid-edge points is first determined by surface-fitting to values at neighbouring centre-points. The distribution of source or doublet strength along panel edges is then determined uniquely from the values at the vertices and mid-edge points. Because the same vertex and mid-edge values are shared between neighbouring panels, this guarantees continuity of singularity strength across panel boundaries. Strength distribution within the panel is then a function of the edge distribution and the remaining centre-point strength.

This formulation is possible in PAN AIR because of the way quadrilateral panels are sub-divided into eight triangles, providing additional degrees of freedom. Johnson acknowledges that the method he describes does not

guarantee continuity across panel edges, but notes that “it is virtually impossible to construct a quadratically accurate doublet spline with exact continuity across all panel edges of a network when the distribution of a panel has only six degrees of freedom”.

Like Johnson’s method, computation of the panel influence coefficients is a lengthy procedure (see Epton & Magnus, 1990, Appendix J, which spans 200 pages), but one that does produce an exact solution. Whilst the computation shares some features of Johnson’s procedure, it is a different formulation which should, for a given sub-element, produce the same results. Post-processing routines then determine surface velocities, pressures and load distributions in the usual way.

Versions of PAN AIR have been released into the public domain, the most recent release being known as the ‘ht2 version’, dated 1993. The Boeing Company has since continued its development for commercial use.

2.2.6 Bi-linear vortex triangle codes

Kandil *et al.* (1984) develop a scheme comprising planar quadrilateral panels which carry bi-linear vorticity distributions. The scheme is used for both the steady and the unsteady solutions of thin wing problems, with triangular elements replacing the quadrilaterals in deforming wakes because of their ability to model non-planar and twisted surfaces more accurately. Equations are set up to ensure that continuity of vorticity and the Kutta condition (where applicable) are satisfied at panel vertices; the no-penetration condition is enforced at panel centre points. Neither the derivation nor the resulting equations for the velocity influence coefficients are provided; the paper simply states that a closed-form solution was obtained that had acceptable computational efficiency.

The resulting system is over-constrained with respect to the unknowns of local panel vorticity distributions, and is solved together with the wake geometry relaxation by a nonlinear least squares method. Accordingly, the scheme is termed a *Nonlinear Hybrid Vortex* method by the authors, where *Hybrid* refers to

the use of a standard vortex-lattice approximation for far-field induced velocity calculations.

Mracek *et al.* (1992) present a very similar scheme to that of Kandil *et al.* (1984), but uses triangular elements throughout rather than planar quadrilateral elements for the lifting surfaces (note that planar quadrilateral elements can always be constructed from triangular elements, so there is no fundamental difference in complexity between the two). The problem formulation and boundary conditions appear similar for both schemes, and again, the over-constrained problem is solved by means of a linear least-squares procedure.

The principal difference between the two schemes is in their application. Whereas Kandil *et al.* (1984) present their method as a solution to the thin lifting surface flow problem, the Mracek *et al.* (1992) scheme is proposed as an alternative to conventional source-panel methods for the modelling of closed-body flow problems.

In the case of Mracek *et al.* (1992), the set of equations representing the velocity influence coefficients are provided, and appear relatively straightforward. The derivations are not provided, although the paper refers to theses by two of the authors for further details.

A further scheme using linear vorticity triangles is presented by Plotkin & Yeh (1986), in this case to model the wake roll-up behind a large aspect ratio wing. In common with the previous two methods described, an iterative solution is found, in this case with the relaxation of the wake geometry forming part of the iteration cycle. Closed-form influence coefficients are used, for details of which the paper refers to a thesis by one of the authors.

2.2.7 Horstmann's code

Horstmann (1987) presents an extension of the conventional vortex ring method in which the spanwise vortex lines are 'upgraded' from constant strength to quadratic strength. A standard horseshoe vortex element, which forms the basis of vortex ring panels, has constant vortex strength in its spanwise section, and can be regarded as shedding single concentrated vortex lines at its start and end

where there is a jump in spanwise vorticity. If however the spanwise element has a linearly varying vortex strength, it will shed a sheet of infinitesimal vortex lines of equal strength between the two outer lines; similarly, a spanwise vortex line having quadratic strength will shed a sheet of vortex lines which increase linearly in strength line-to-line. In the same way that a vortex ring is assembled from two equal and opposite standard horseshoe lifting lines located at its leading and trailing edges, Horstmann's elements are formed by two equal and opposite quadratic vortex lines, joined by a sheet of vorticity parallel to the stream.

A network of panels is constructed such that the spanwise vortex lines are continuous in value at element junctions, which ensures that there are no streamwise concentrations of vorticity in the trailing vortex sheet. However, there are clearly spanwise vorticity concentrations on the quadratic vortex lines. This scheme can be visualised more clearly in terms of its equivalent doublet distribution, which is quadratic continuous spanwise, and piecewise constant streamwise, as shown in figure 2.2, case (c).

Horstmann's elements are restricted to being planar and trapezoidal, with two edges parallel to the stream and aligned with the vortex sheet. Induced velocities are calculated in closed form via a relatively simple construction, and as such, the scheme is likely to be reasonably efficient.

A FORTRAN implementation of this method is available in the public domain, and is supported by the *Deutsches Zentrum für Luft und Raumfahrt (German Aerospace Centre)*. It has been widely used in Germany for the design of glider aerofoils, and has formed the basis for other more complex schemes. In particular, Bramesfeld & Maughmer (2004) extends Horstmann's method to include a free-wake via a time-stepping method. In this case, the absence of streamwise vortex line singularities makes the method particularly suitable to modelling wake roll-up, and is reported to alleviate some of the problems described in Bernasconi (2007, chapter 3). In Friedl (2002), a hybrid method derived from the schemes of Horstmann and Mracek *et al.* (1992) is linked to a nonlinear finite element membrane code to solve the coupled problem of membrane wings in unsteady potential flow. The method is applied to the analysis of a hang glider in unsteady flight.

It is interesting to compare Horstmann's scheme to that of Woodward (1973). Both schemes, derived from vortex lines, are equivalent to constant-quadratic doublet distributions, and both comprise trapezoidal elements in which two parallel edges are aligned with the streamwise direction. However, where Woodward's quadratic doublet variation is in the streamwise direction, Horstmann's quadratic doublet variation is spanwise. Since the pressure distribution across a wing or sail is likely to vary more rapidly in the streamwise direction, Woodward's choice of alignment appears better suited to the modelling task; however, a benefit of Horstmann's alignment is in the elimination of vortex line singularities in a (steady-state) wake.

2.2.8 Maniar's code

Maniar (1995) presents a method in which the surface geometry and surface singularity strength are both represented by (typically bi-cubic) B-spline tensors. The unknowns are the source or doublet strength spline coefficients, and a well-constrained (square) set of equations is set-up via a Galerkin procedure to satisfy the standard boundary conditions. Although the solution is constructed in the sense of minimising errors over the continuous B-spline surfaces, boundary conditions are ultimately evaluated at discrete integration points which are chosen by using a Gauss-quadrature rule over the intervals between spline knots.

Different techniques are used to compute the influence coefficients relating to the B-spline source or doublet strength terms, depending on the position of the field point relative to the panel. For self-influence coefficients, where the field point lies on the panel, the (four-sided) panel is divided into four triangles, where the apex of each triangle is the field point. An algebraic series expansion is then undertaken on each of the triangles, and their effects summed. For near-field coefficients, Maniar provides two alternative methods: one is a series expansion similar to that presented by Johnson (1980), and the other is an adaptive sub-panelling method. In this case, the panel is recursively sub-divided into smaller panels, refining the mesh in the region of the field point until all of the sub-panels

can be treated as far-field panels. For far-field coefficients, a Gauss-quadrature rule is employed to evaluate the surface integrals numerically.

Maniar presents a number of examples of his method applied to flow around solid bodies, and compares results with those computed by a conventional low-order panel method. For solutions calculated to the same accuracy (1.0% – 0.1%), Maniar reports that the higher-order method is 10 – 200 times faster than the conventional method. Although the computational cost per panel is obviously much higher, the solution is obtained with far fewer unknowns.

Elements of Maniar's method are incorporated into a scheme by Pyo & Kinnas (1997), who use curved panels with bi-quadratic doublet distributions to model propeller wake roll-up. The panel influence coefficients are taken from Maniar, and appear to produce very well formed and stable wake vortices without the use of any special wake management routines.

2.2.9 Summary of methods reviewed

Table 2.1 summarises the main higher-order codes that have been reviewed. For further information, similar tables can be found in Katz & Plotkin (2001, page 351), and in Epton & Magnus (1990, page 57); however, these two tables are less focussed on higher-order methods, and omit some more recent codes.

2.3 Evaluation study

All of the methods described in the previous section have the potential to provide greater accuracy and less grid dependency, and most are also likely to alleviate the problems associated with wake roll-up. However, in most cases, this will be at the cost of computational efficiency.

At the high-end of the accuracy/computation-time scale, Johnson's method is appealing because of its apparent extreme insensitivity to panel layouts. However, the complexity of the influence coefficient calculations may make the scheme impractical. At the lower end of the scale, bi-linear vortex triangles may offer many of the advantages of other higher-order schemes with a much smaller run-time penalty. In order to evaluate the possible benefits and costs of these

Table 2.1: Summary of higher-order methods

Method	Panel geometry	Singularity distribution	Boundary conditions	Remarks
Douglas-Neumann (Hess & Smith, 1962)	Flat quad	Constant source	Neumann	Non-lifting
Hess I (Hess, 1972)	Flat quad	Constant source Mixed doublet	Neumann	
Woodward II (USSAERO) (Woodward, 1973)	Flat trapezium	Linear source Linear vortex	Neumann	
Hess II (Hess, 1979)	Paraboloid quad	Linear source Linear vortex	Neumann	
Johnson (1980)	Paraboloid quad	Linear source Quadratic doublet	Neumann & Dirichlet	
PAN AIR (Epton & Magnus, 1990)	Piecewise flat super-panel	Linear source Quadratic doublet	Neumann & Dirichlet	
Kandil <i>et al.</i> (1984)	Flat quad	Linear vortex	Neumann	Thin bodies
Horstmann (1987)	Flat trapezium	Linear vortex	Neumann	Thin bodies
Mracek <i>et al.</i> (1992)	Triangle	Linear vortex	Neumann	Thick bodies
Maniar (1995)	B-spline patches	Any order source Any order doublet	Dirichlet (Galerkin)	Thick bodies
The new CSM method described herein	B-spline patches	Any order source Any order doublet	Neumann & Dirichlet	

two strategies, the following sections present preliminary evaluation studies of Mracek's and Johnson's methods.

2.3.1 Johnson's bi-quadratic doublet panels

Panel formulation

For a curved panel whose flat projection lies in the xy -plane, Johnson's formulation generates vector coefficients \mathbf{J}_{ij} , where the induced velocity at a field point (x, y, z) is given by

$$\mathbf{v} = \mu(x, y)\mathbf{J}_{11} + \mu_x(x, y)\mathbf{J}_{21} + \mu_y(x, y)\mathbf{J}_{12} \\ + \frac{1}{2}\mu_{xx}(x, y)\mathbf{J}_{31} + \mu_{xy}(x, y)\mathbf{J}_{22} + \frac{1}{2}\mu_{yy}(x, y)\mathbf{J}_{13}$$

where

$$\begin{aligned}
\mu(x, y) &= \mu_0 + \mu_\xi x + \mu_\eta y + \frac{1}{2}\mu_{\xi\xi}x^2 + \mu_{\xi\eta}xy + \frac{1}{2}\mu_{\eta\eta}y^2 \\
\mu_x(x, y) &= \mu_\xi + \mu_{\xi\xi}x + \mu_{\xi\eta}y \\
\mu_y(x, y) &= \mu_\eta + \mu_{\xi\eta}x + \mu_{\eta\eta}y \\
\mu_{xx}(x, y) &= \mu_{\xi\xi} \\
\mu_{xy}(x, y) &= \mu_{\xi\eta} \\
\mu_{yy}(x, y) &= \mu_{\eta\eta}
\end{aligned}$$

and the panel's doublet distribution is

$$\mu = \mu_0 + \mu_\xi x + \mu_\eta y + \frac{1}{2}\mu_{\xi\xi}x^2 + \mu_{\xi\eta}xy + \frac{1}{2}\mu_{\eta\eta}y^2 \quad (2.1)$$

For incorporation into the problem structure detailed below, it is necessary to define the influence coefficients in terms of doublet strength and doublet strength gradients at panel-fixed reference points, rather than at the field points. Thus, a function was created to perform the necessary global-to-local coordinate transformations, and to redefine the influence coefficients \mathbf{q} by

$$\mathbf{v} = \mu_0\mathbf{q}_{11} + \mu_\xi\mathbf{q}_{21} + \mu_\eta\mathbf{q}_{12} + \frac{1}{2}\mu_{\xi\xi}\mathbf{q}_{31} + \mu_{\xi\eta}\mathbf{q}_{22} + \frac{1}{2}\mu_{\eta\eta}\mathbf{q}_{13} \quad (2.2)$$

where by substitution,

$$\begin{aligned}
\mathbf{q}_{11} &= \mathbf{J}_{11} \\
\mathbf{q}_{21} &= x\mathbf{J}_{11} + \mathbf{J}_{21} \\
\mathbf{q}_{12} &= y\mathbf{J}_{11} + \mathbf{J}_{12} \\
\mathbf{q}_{31} &= x^2\mathbf{J}_{11} + 2x\mathbf{J}_{21} + \mathbf{J}_{31} \\
\mathbf{q}_{22} &= xy\mathbf{J}_{11} + y\mathbf{J}_{21} + x\mathbf{J}_{12} + \mathbf{J}_{22} \\
\mathbf{q}_{13} &= y^2\mathbf{J}_{11} + 2y\mathbf{J}_{12} + \mathbf{J}_{13}.
\end{aligned}$$

The influence function has the form

$$[q] = \text{INFLUENCEJ}(rF, rV, rC, nXiEtZe)$$

where rF is the (3×1) field point, rV is the (3×4) set of panel vertices, rC is the (3×1) panel centre (origin of the local coordinate system), $nXiEtZe$ is the (3×3) set of unit vectors defining the local system, and q is the (3×6) set of influence coefficients.

The influence coefficients produced by Johnson's method were successfully verified against a numerical integration of a sheet of finite strength vortex lines whose influences were found with a standard vortex line function. Results of timing tests are given at the end of this section.

Problem structure

In order to evaluate the elements in the context of a potential flow problem, a solution structure similar in concept to (but not taken from) that presented by Johnson was implemented. The doublet strengths μ_0 at panel centroids were treated as the unknowns of the problem, and for each panel, the doublet strengths of neighbouring panels were used to estimate local doublet strength gradients. For a panel P , with neighbours A, B, C , etc., an estimate of the gradients defined in equation 2.1 is provided by the least-squares solution of

$$S\mu_P = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & x_{PA} & y_{PA} & x_{PA}^2 & x_{PA}y_{PA} & y_{PA}^2 \\ 1 & x_{PB} & y_{PB} & x_{PB}^2 & x_{PB}y_{PB} & y_{PB}^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix} \begin{pmatrix} \mu_{0,P} \\ \mu_{\xi,P} \\ \mu_{\eta,P} \\ \frac{1}{2}\mu_{\xi\xi,P} \\ \mu_{\xi\eta,P} \\ \frac{1}{2}\mu_{\eta\eta,P} \end{pmatrix} \approx \begin{pmatrix} \mu_{0,P} \\ \mu_{0,A} \\ \mu_{0,B} \\ \vdots \end{pmatrix}$$

where (x_{PA}, y_{PA}, z_{PA}) are the coordinates of panel A 's centre within the local coordinate system centred on panel P , and $\mu_{0,P}$ is the doublet strength at panel

P 's centre, etc. The influences

$$\mathbf{q}_{0,P} = \begin{pmatrix} q_P \\ q_A \\ q_B \\ \vdots \end{pmatrix}$$

of each unknown μ on the velocity induced by panel P are then given by

$$\mathbf{q}_{0,P} = \mathbf{q}_{\xi\eta,P}(\mathbf{S}'\mathbf{W}\mathbf{S})^{-1}(\mathbf{S}\mathbf{W})$$

where

$$\mathbf{q}_{\xi\eta,P} = \begin{pmatrix} \mathbf{q}_{11} & \mathbf{q}_{21} & \mathbf{q}_{12} & \mathbf{q}_{31} & \mathbf{q}_{22} & \mathbf{q}_{13} \end{pmatrix}$$

are the influences defined by equation 2.2, and \mathbf{W} is a diagonal matrix of weights. The first element of \mathbf{W} is set to 10^6 , and the remaining diagonal elements to 1, to ensure that the 'estimate' of doublet strength of panel P is almost exactly the actual assigned doublet strength of that panel. The term $(\mathbf{S}'\mathbf{W}\mathbf{S})^{-1}(\mathbf{S}\mathbf{W})$, which is the (weighted) pseudo-inverse of \mathbf{S} , is effectively a shape function for panel P .

This construction provides the influence of each unknown doublet strength on the velocity induced at the field point by panel P . In practice, the number of neighbouring panels used to define the local doublet strength gradients at P is limited to the eight surrounding panels, such that the influences of most panels will be zero. This process is carried out for each panel, and the influences of each unknown on each panel's contribution are summed to determine the total influence of each panel on the field point. The function developed to achieve this has the form

$$[\mathbf{q}] = \text{INFLUENCEJG}(\mathbf{rF}, \mathbf{P})$$

where \mathbf{rF} is the (3×1) field point, \mathbf{P} is a structure containing descriptions of the panels, and \mathbf{q} is the $(3 \times N)$ set of influence coefficients where N is the number of panels.




Figure 2.3: Doublet strength distribution over a rectangular wing with a (5×5) grid of bi-quadratic panels.

Application to a rectangular wing

The elements and problem structure outlined above were tested using the standard planar rectangular wing with aspect ratio equal to 2 and a fixed semi-infinite wake, subject to a steady inclined flow. Timing results for this test are given at the end of this section, and figure 2.3 shows the calculated doublet strength distribution.

It can be seen that there are small gaps between the doublet panels, corresponding to step changes in the doublet strength. This is a result of the doublet strength gradients being determined separately for each panel without enforcing the continuity at boundaries as the PAN AIR method does. Also note that using the basic scheme outlined above, the doublet strength does not fall to zero at the leading-edge and sides of the wing as it should; also, the trailing-edge has a falling doublet-strength gradient, where a zero streamwise gradient would be expected. To correct these errors, additional edge conditions need to be added to the problem structure; currently, edge panels are treated in the same way as

internal panels, except that two inner rows of neighbouring panels are used for defining the shape functions, instead of one row either side.

2.3.2 Mracek's bi-linear vortex triangles

Panel formulation

Mracek's panels comprise a bi-linear vortex line distribution of the form:

$$\boldsymbol{\gamma} = (\gamma_x + \gamma_{xx}x + \gamma_{xy}y)\mathbf{i} + (\gamma_y + \gamma_{yx}x + \gamma_{yy}y)\mathbf{j}$$

where to satisfy the Helmholtz divergence theorem,

$$\nabla \cdot \boldsymbol{\gamma} = 0$$

so that

$$\gamma_{xy} = -\gamma_{yx}$$

reducing the number of vortex strength parameters to five for each panel.

A function was developed that computes the influences due to the five terms for a given field point. However, for this preliminary study, only the two constant vortex terms were used, corresponding to a linear doublet strength distribution. For this case, the influences can be presented as functions of the doublet strengths at each of the triangle's vertices, providing a function of the form:

$$[\mathbf{q}] = \text{INFLUENCEM}(\mathbf{rF}, \mathbf{rV})$$

where \mathbf{rF} is the (3×1) field point, \mathbf{rV} is the (3×3) set of panel vertices, and \mathbf{q} is the (3×3) set of velocities induced by a unit doublet strength at each vertex, assuming a linear doublet distribution over the panel.

Influence coefficients produced by Mracek's method were successfully verified against a numerical integration of a sheet of finite strength vortex lines whose influences were found with a standard vortex line function. Results of timing tests are given at the end of this section.

Super-panel formulation

In this chapter, the term super-panel is used to refer to a structured group of panels linked by geometry and singularity distribution. The super-panel can be treated as a single panel by higher level functions, but its influence coefficients are determined by summing the contributions of its individual sub-panels.

To allow alternative panel schemes to be investigated (see Section 2.3.2), a quadrilateral super-panel type was created comprising four bi-linear doublet strength triangles with a common central vertex. The quadrilateral panel, which is not necessarily planar, is defined by the coordinates and doublet strengths of each of its vertices. The central vertex is assumed to have coordinates and a doublet strength equal to the mean of the four corner values, and a linear doublet strength distribution over each of the four triangles is assumed. The super-panel's influence function has the form:

$$[q] = \text{INFLUENCEM4}(rF, rV)$$

where rF is the (3×1) field point, rV is the (3×4) set of vertex coordinates and q is the (3×4) set of influences due to each corner vertex.

Super-panel influence coefficients were successfully verified as above.

Distinction between vortex line and doublet panels

Whilst the influences of the triangular panels (and quad super-panels) are described in terms of the doublet strengths of their vertices, they are not directly equivalent to panels containing linear doublet strengths. These latter panels, when described in terms of a vortex line equivalence, would contain the same internal vortex line distribution but would also comprise concentrated vortex lines of linearly varying strength around their perimeters. In the present vortex line based panels, it is only the difference between doublet strengths at vertices that generates a velocity influence, whereas in a true doublet panel the absolute doublet strength is also significant.

When panels are aggregated into a mesh, this distinction becomes less important because the solution is configured such that panels always meet

exactly along edges holding matching doublet strength distributions. The missing vortex lines around each panel thereby cancel out in the overall solution, meaning that when combined in this way the panels can indeed be considered in terms of doublet strength.

However, there remains a disadvantage with the concept of using these vortex line panels in a doublet formulation, which is that the panels' self-influences are generally low. If the panels were true doublet panels with concentrated edge vortices, the influence of a doublet gradient remote from the panel itself would effectively be shifted into the self-influence of the local panel, forming a more strongly diagonal solution matrix and allowing far-field influence coefficients to be used. In contrast, vortex line panels have zero self-influence if the local doublet strength is constant, and rely on contributions from the entire mesh.

Application to a rectangular wing

Three solution structures utilising linear doublet strength triangular panels were investigated (no use was made of the quadratic components in this preliminary study). The studies were based on the standard planar rectangular wing with aspect ratio equal to 2, with a fixed semi-infinite wake, subject to a steady inclined flow. The three studies were:

- The rectangular wing and wake were divided into regular square regions, and each square sub-divided into four triangular panels (the panels were treated individually in this case, not as super-panels). The doublet strengths at nodes on the wing's edges were constrained to zero, whilst the doublet strengths at the remaining nodes formed the problem's unknowns. Collocation points were located at each wing panel's centre (defined by the average of the nodes), and Neumann boundary conditions were imposed. This formed an over-constrained problem which provided a linear least-squares solution.
- The wing and wake were divided into regular square super-panels (such that the sub-panel grid was the same as above). Unknowns were the doublet strengths at internal super-panel nodes, and collocation points were located

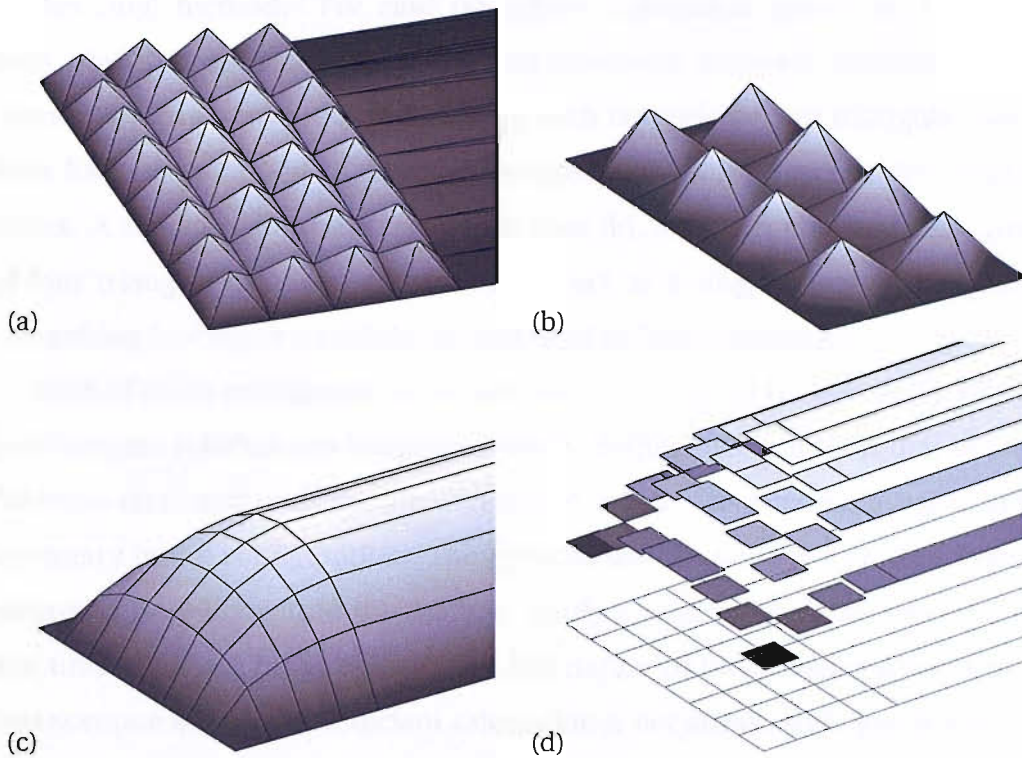


Figure 2.4: (a) Constant vorticity triangles; (b) Super-panels from triangles, collocation points at panel centres; (c) Super-panels from triangles, collocation points at vertices; (d) Constant doublet strength quadrilaterals (vortex ring)

at super-panel centres. Note that the influence of a rectangular panel with constant doublet strength is zero at its centre, so this contribution can be ignored – thus avoiding the difficulty described in the third case below.

- The super-panels were distributed as above, but in this case the collocation points were located at super-panel nodes instead of at their centres. Because the super-panel centres coincided with the sub-panel vertices, at which the influence coefficients are logarithmically unbounded, this was achieved in practice by averaging the influence at field points just above and just below the panel centres.¹

Figure 2.4 shows the resulting distribution in doublet strength for the three cases, in addition to the distribution for an equivalent constant doublet strength

¹This averaging process is sufficient to demonstrate the principle of the scheme, but is not suitable for producing accurate results, since the exact distance of the collocation points above and below the panels significantly affects the results. The scheme described by Maniar (1995) also requires calculation of influences of this type, and he notes that whilst the surface integrals are singular, they are integrable. A series expansion is then used to evaluate the integrals.

(vortex ring) method. For case (a), where collocation points are located at each triangle's centre, a surprising and obviously incorrect doublet strength distribution has been produced. Groups each comprising four triangular panels have formed 'pyramids' of doublet strength, each with a peak at the common vertex. A similar pattern has resulted in case (b), although now that each group of four triangles has been constrained to act as a single super-panel, groups comprising four super-panels have combined to form 'pyramids'.

Both of these configurations formed over-constrained problems, for which a least-squares solution was sought. However, despite this, solutions that satisfied the boundary conditions exactly were found, which was possible because of the symmetry in the configuration. The pyramid distributions have arisen from an inherent instability within the method, partly a result of the collocation points and unknowns not being coincident. The nature of these results demonstrates that accurate influence coefficient calculation is not always sufficient to ensure a stable solution.

Case (c), in which unknowns are collocation points coincident with super-panel nodes, appears to have been more successful. The doublet distribution resembles that expected of a flat plate, with a steep gradient (high pressure coefficient) at the leading edge, and almost zero gradient towards the trailing edge.

Figure 2.5 shows the normal induced velocity component along a spanwise traverse over the wing, starting beyond the wing's edge and moving towards its centre. It compares the case of a standard vortex ring solution (constant doublet strength) with that of the most successful linear doublet strength construction. It is clear that whilst the linear case has some discontinuities in velocity gradient, these are less severe than those associated with the vortex ring method.

Inspection of the curves in this figure can help to explain the poor solutions obtained for two of the formulations described above. For the linear doublet velocity traverse shown, collocation points are aligned with the peaks seen at span coordinates -0.50 and -0.75 , although at different chordwise coordinates to those at which the traverse is taken. Whilst the solution ensures the normal velocity boundary condition is satisfied at collocation points, the sharp peaks




Figure 2.5: A traverse across the span of a wing, just above its surface, showing normal velocities computed by constant and linear doublet strength solutions. In a perfect solution, all velocities within the span should be zero, whereas in a numerical solution this condition is generally only satisfied at collocation points. The traverse shown lies between rows of collocation points for both cases.

in the velocity curve demonstrate that it is in fact only satisfied very locally. The inflow velocity had a normal component magnitude of around 0.1, so the overall velocity profile is seen to be quite poor. In the alternative panel formulations, where collocation points (between the velocity peaks) did not coincide with the unknowns of the problem (which remained at the velocity peaks), the resulting solutions were understandably less predictable. By analogy, the solution of a constant doublet strength formulation would become unstable if the collocation points were moved towards the vortex lines represented by the velocity singularities in the figure's blue curve.

The scheme using Johnson's quadratic panels has not yet been developed to the stage where inclusion on this graph would give a comparable result. However, it is anticipated that the bi-quadratic solution would provide a more accurate velocity curve that would be smooth except in regions immediately surrounding

panel boundaries, where slight discontinuities in doublet strength would produce very local velocity singularities.

2.3.3 Comparison of computation times

The table below shows the comparative computation times for the methods described. The times are shown relative to a vortex ring method, and are based on the solution of a rectangular wing with a (5×5) grid plus semi-infinite wake. For the vortex-ring and quadrilateral doublet methods, these grids comprise 30 panels and 25 collocation points, requiring a total of 750 influence calculations. For the triangular panel methods, costs are shown for the same number of elements, but also for the case where each quadrilateral super-panel was subdivided into 4 triangles.

Method	Cost
Constant doublet strength (standard vortex ring)	1.0
Linear doublet strength triangles (Mracek <i>et al.</i> , 1992) ^a	3.9
Linear doublet strength triangles (Mracek <i>et al.</i> , 1992) ^b	26.0
Quadratic doublet strength triangles (Mracek <i>et al.</i> , 1992) ^a	4.7
Quadratic doublet strength triangles (Mracek <i>et al.</i> , 1992) ^b	31.3
Quadratic doublet strength quadrilaterals (Johnson, 1980)	60.0

^a 30 panels, 25 collocation points, to match other cases.

^b 120 panels, 36 collocation points, required for super-panel scheme.

The quadratic doublet strength triangles use the same influence calculations as the linear doublet strength triangles, but in the linear case some of the code can be omitted resulting in a small time-saving. In practice, the times shown for all of the higher-order methods are likely to be over-estimates, because cheaper far-field approximations of influence coefficients could be made for a large proportion of the collocation point – panel pairs. Also, the higher-order codes have not been optimised to the extent that the vortex ring code has been.

2.4 Conclusions

The literature review and evaluation studies have shown that there are a number of alternatives to simple constant panel and vortex ring methods. Most of these higher-order schemes are likely to produce more accurate results which are less sensitive to panel layout, and which would improve the stability of the wake development. These performance improvements are naturally offset by an increase in computational run-time, with the simplest higher-order method being four times slower than a vortex ring scheme for the same number of panels and collocation points.

However, the relevant comparison is not the computational cost for a given number of panels, but the cost to achieve a given accuracy. This comparison has been made before, and examples such as Maskew (1982) have generally found in favour of low-order methods. For the aims of this project though, basic accuracy is not the only consideration, and the potential benefits to wake roll-up will certainly shift the balance more in favour of higher-order methods. Furthermore, previous studies have mostly compared ‘heavy-weight’ alternatives such as PAN AIR, whereas simpler higher-order schemes may offer a reasonable compromise.

It is thought that neither the constant-quadratic scheme of Woodward (1973), nor the quadratic-constant scheme of Horstmann (1987) are ideal candidates to solve the main issues: the former because the concentrated streamwise vortex lines will not alleviate the wake roll-up problem, and the latter because the streamwise doublet strength is the same order as the vortex ring method. If a piecewise constant distribution in this direction is maintained, it is unlikely that significant improvements in terms of grid insensitivity will be made.

Johnson’s method, whilst apparently highly accurate and insensitive to panel layout, does not produce a completely smooth doublet distribution. Moreover, it is very computationally expensive, and as such would be a last resort. The PAN AIR method, whilst producing a continuous solution, has an element formulation that is expected to be as costly as Johnson’s.

In the scheme presented by Mracek *et al.* (1992), the panels’ influence coefficients are relatively cheap to compute, but the overall scheme appears

unnecessarily complex. The vorticity vectors at each node are treated as unknowns (six parameters) and constraints are introduced into the solution to ensure zero divergence both within the elements and where elements meet. However, as detailed in section 2.1, the zero divergence criterion is automatically satisfied by any doublet distribution. As such, by regarding the panels as bi-quadratic doublet elements, and deriving each panel's vorticity distribution from the local doublet strength gradient, the additional constraints could have been dispensed with.

Finally, the significant performance gains reported for Maniar's method are impressive. The accuracy and computation time benefits, combined with a smooth doublet distribution to simplify wake roll-up, would appear to make this method very well suited to the requirements. However, it should be noted that the quoted computation-time gains are in comparison to a specific low-order panel code, which may not itself be particularly efficient.

Maniar's scheme is significantly different to the others discussed, in that it is not panel-based, but models curved geometric surfaces directly. However, despite its higher order, the formulation of the influence coefficients is relatively simple, particularly in comparison to the hugely complex algorithms such as those of PAN AIR or Johnson. What the method lacks, though, is the ability to model thin surfaces and wakes, and because of the higher orders of singularity involved in forming velocity coefficients, this omission is far from trivial. Whilst the Continuous Surface Method described in this thesis differs in many ways to Maniar's scheme, it is influenced by his representation of geometry and singularity distribution, which were considered to be the most promising amongst those reviewed in this chapter.

Chapter 3

The Continuous Surface Method

3.1 Problem formulation

Our goal is to determine the three dimensional velocity field of an inviscid, incompressible and irrotational fluid subject to boundary conditions imposed by one or more bodies within the region of interest. We take as our starting point Laplace's equation, which is the continuity equation for an incompressible irrotational fluid

$$\nabla^2 \phi^T(\mathbf{y}) = 0 \quad (3.1)$$

where $\phi^T(\mathbf{y})$ is the total velocity potential of the fluid, and \mathbf{y} is a general point in the fluid region. In common with traditional panel methods, we seek a distribution of singularity elements over the boundaries of the fluid region to produce a flow field satisfying the boundary conditions. The two types of singularity elements we consider are doublets μ and sources σ , which when located at a point \mathbf{x} on a fluid boundary produce the perturbation potentials at $\mathbf{y} \neq \mathbf{x}$

$$\phi_\mu^P(\mathbf{y}) = -\frac{\mathbf{r}(\mathbf{y}, \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})\mu(\mathbf{x})}{4\pi r^3(\mathbf{y}, \mathbf{x})} \quad (3.2)$$

and

$$\phi_\sigma^P(\mathbf{y}) = -\frac{\sigma(\mathbf{x})}{4\pi r(\mathbf{y}, \mathbf{x})} \quad (3.3)$$

where $\mathbf{r}(\mathbf{y}, \mathbf{x}) = \mathbf{x} - \mathbf{y}$, $r(\mathbf{y}, \mathbf{x}) = |\mathbf{r}(\mathbf{y}, \mathbf{x})|$ and $\mathbf{n}(\mathbf{x})$ is a unit vector orienting the doublet. Both of these singularity elements satisfy equation 3.1, and have the property that the perturbation potentials they induce decay to zero at infinity.

By integrating the influence of source and doublet distributions over the external boundaries S_B and S_W of a body and its wake, the total potential $\phi^T(\mathbf{y})$

can be written as the sum of the free-stream potential $\phi^\infty(\mathbf{y})$ and the perturbation potential $\phi^P(\mathbf{y})$. Wakes will be modelled with zero thickness, and it is assumed that $\partial\phi/\partial n$ is continuous across them, which constrains their surface source strength distributions to be zero. This means that no fluid-dynamic loads will be supported by the wake,

$$\phi^T(\mathbf{y}) = \phi^P(\mathbf{y}) + \phi^\infty(\mathbf{y}) \quad (3.4)$$

$$= \frac{1}{4\pi} \int_{S_B+S_W} \mu \mathbf{n} \cdot \nabla \left(\frac{1}{r} \right) dS - \frac{1}{4\pi} \int_{S_B} \sigma \left(\frac{1}{r} \right) dS + \phi^\infty \quad (3.5)$$

where the functional dependencies of $\sigma(\mathbf{x})$, $\mu(\mathbf{x})$, $\mathbf{n}(\mathbf{x})$, $r(\mathbf{x}, \mathbf{y})$, $dS(\mathbf{x})$ and $\phi^\infty(\mathbf{y})$ have been omitted for brevity.

Equation 3.5 is a form of Green's third identity (Kellogg, 1954)

$$\phi^P(\mathbf{y}) = \frac{1}{4\pi} \int \frac{1}{r} \nabla^2 \phi^P dV + \frac{1}{4\pi} \int \phi^P \frac{\partial}{\partial \xi_3} \frac{1}{r} dS - \frac{1}{4\pi} \int \frac{1}{r} \frac{\partial \phi^P}{\partial \xi_3} dS \quad (3.6)$$

where ξ_3 is an intrinsic coordinate aligned with the surface normal. Comparing equation 3.6 with equation 3.5, the first term on the right is zero by equation 3.1, the surface potential $\phi^P = \phi^P(\mathbf{x})$ is represented by the doublet strength $\mu(\mathbf{x})$, and the potential's normal derivative is equivalent to the source strength $\sigma(\mathbf{x})$.

The potential and velocity influences of the singularity distributions are unbounded as $r \rightarrow 0$, or $\mathbf{x} \rightarrow \mathbf{y}$ (hence the term 'singularity'). Therefore, when determining the potential or velocity on a boundary, a small region of S_B or S_W must be excluded from the integration. The limiting value of that region's influence is then determined as its size tends to zero, and this principal value is added to the result of the integration. This matter is covered in detail in section 3.2.

Alternatively, we can consider the potential induced by a singularity distribution as we approach the boundary along its surface normal, as shown in figure 3.1. If $\mathbf{n}(\mathbf{x})$ is a unit vector normal to the boundary, we denote by \mathbf{x}_- and




Figure 3.1: Singularity distribution on a surface

\mathbf{x}_+ the limiting points on either side of the boundary such that as $\varepsilon \rightarrow 0$

$$\mathbf{x}_- = \mathbf{x} - \varepsilon \mathbf{n} \quad (3.7)$$

$$\mathbf{x}_+ = \mathbf{x} + \varepsilon \mathbf{n} \quad (3.8)$$

By application of the divergence theorem to an arbitrary fluid region (Katz & Plotkin, 2001), it can be shown that a distribution of source elements over the boundary of the region produces a difference between the normal derivatives of the potentials either side of the boundary,

$$-\sigma(\mathbf{x}) = \left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{x}_-} - \left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{x}_+} \quad (3.9)$$

It can similarly be shown that a distribution of doublet elements produces a difference between the potentials across the boundary,

$$-\mu(\mathbf{x}) = \phi^T(\mathbf{x}_-) - \phi^T(\mathbf{x}_+) \quad (3.10)$$

These relationships will be used in forming the Dirichlet boundary condition below.

3.1.1 Surface boundary conditions

The method described herein allows either Neumann and Dirichlet boundary conditions to be specified, or both to be mixed, for example using a Dirichlet formulation on a thick mast and a Neumann formulation on a sail modelled with

zero thickness. We only consider non-porous surfaces over which the flow must have zero normal component, except that resulting from motion of the surface. Thus for the static case, the boundary conditions constrain the flow adjacent to a body to be tangential to its surface. Whilst the Neumann condition states this physical constraint directly in terms of surface velocities, the Dirichlet condition is applied in terms of potentials.

Neumann boundary condition

The Neumann (or direct) boundary condition on a solid boundary S_B is stated as

$$\mathbf{q}(\mathbf{y}_-) \cdot \mathbf{n}(\mathbf{y}) = 0 \quad (3.11)$$

where \mathbf{y}_- (or \mathbf{y}_+) is a limiting point on one side of the boundary or the other, and $\mathbf{q}(\mathbf{y}_-)$ is the total velocity comprising free-stream and perturbation components, measured in a body-fixed reference frame. The additional boundary condition, that flow perturbations far from the body tend to zero, is automatically satisfied by any source or doublet distribution over the surface of the body. To apply the Neumann boundary condition we therefore simply specify the normal velocity over all body surfaces. Writing the velocity as the gradient of the potential, employing equation 3.5, we obtain

$$\left\{ \frac{1}{4\pi} \int_{S_B+S_W} \mu \nabla \left[\mathbf{n} \cdot \nabla \left(\frac{1}{r} \right) \right] dS - \frac{1}{4\pi} \int_{S_B} \sigma \nabla \left(\frac{1}{r} \right) dS + \nabla \phi^\infty \right\} \cdot \mathbf{n} = 0 \quad (3.12)$$

Dirichlet boundary condition

The Dirichlet (or indirect) boundary condition can be applied to boundaries that form a closed surface, with the region of interest lying outside the closed boundary. We choose to orient the surface normal $\mathbf{n}(\mathbf{y})$ such that it points into the closed region, such that \mathbf{y}_- lies in the fluid region of interest, and \mathbf{y}_+ lies just inside the closed boundary as shown in figure 3.2.




Figure 3.2: A thick surface

The Dirichlet boundary condition can be derived by considering the Neumann condition in potential form

$$\mathbf{q}(\mathbf{y}_-) \cdot \mathbf{n}(\mathbf{y}) = \nabla \phi^T(\mathbf{y}_-) \cdot \mathbf{n}(\mathbf{y}) = \left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{y}_-} = 0 \quad (3.13)$$

This Neumann condition implies that no streamlines enter or leave the region bounded by the surface S_B . Furthermore, as streamlines cannot begin or end within a region, or form closed curves lying wholly within it, there can be no motion within the region. The conclusion can be drawn that

$$\phi^T(\mathbf{y}_i) = \text{const.} \quad (3.14)$$

where \mathbf{y}_i is any point within the boundary: an 'internal' point. This theorem can otherwise be stated as follows: no continuous irrotational motion of a fluid can take place in a simply-connected region bounded entirely by fixed rigid walls (Lamb, 1932, Art.40 β).

Within a numerical method, it is not sufficient to require a constant unknown internal potential, and a value for this potential must be chosen. From a physical point of view the choice is somewhat arbitrary and should not affect the predicted external flow field; numerically, it is convenient to set the internal potential to be equal to the free-stream potential such that the perturbation potential

$$\phi^P(\mathbf{y}_i) = \phi^T(\mathbf{y}_i) - \phi^\infty(\mathbf{y}_i) = 0 \quad (3.15)$$

This step has been taken with the assumption that the free-stream potential $\phi^\infty(\mathbf{y}_i)$ is constant, implying that the fluid is at rest, or that the reference frame in which the potential is measured is attached to fluid moving with uniform velocity. Although body-fixed reference frames will generally be used in this thesis, this ‘fluid-fixed frame’ is a valid concept for the purpose of applying the Dirichlet boundary conditions. This requires the ‘simply-connected region bounded entirely by fixed rigid walls’ described by Lamb above to be extended, correctly, to a ‘region bounded entirely by rigid walls moving with a constant uniform linear velocity’. Constraining the internal flow to match the free-stream flow in this way can produce solutions with greater numerical stability than those in which the internal potential is set to zero (see for example Maskew, 1982).

Since the perturbation potential ϕ^P has been set to zero throughout the internal region V_i , it follows also that

$$\frac{\partial \phi^P}{\partial n} = \frac{\partial \phi^T}{\partial n} = 0 \quad (3.16)$$

within V_i , and particularly on the internal side of the boundary S_B . Equation 3.9 can then be simplified

$$-\sigma(\mathbf{x}) = \left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{x}_-} - \left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{x}_+} = \left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{x}_-} \quad (3.17)$$

and by application of the Neumann boundary condition

$$\left. \frac{\partial \phi^T}{\partial n} \right|_{\mathbf{x}_-} = -\mathbf{n}(\mathbf{x}) \cdot \mathbf{Q}^\infty \quad (3.18)$$

the source distribution around S_B is seen to be equal to the component of the free-stream normal to the surface

$$\sigma(\mathbf{x}) = \mathbf{n}(\mathbf{x}) \cdot \mathbf{Q}^\infty \quad (3.19)$$

With the perturbation potential equal to zero on the inner surface of the boundary S_B , equation 3.5 can be written

$$\frac{1}{4\pi} \int_{S_B+S_W} \mu \mathbf{n} \cdot \nabla \left(\frac{1}{r} \right) dS - \frac{1}{4\pi} \int_{S_B} \sigma \left(\frac{1}{r} \right) dS = 0 \quad (3.20)$$

As the source distribution is provided by equation 3.19, the problem is reduced to finding the body and wake doublet distributions that satisfy

$$\frac{1}{4\pi} \int_{S_B+S_W} \mu \mathbf{n} \cdot \nabla \left(\frac{1}{r} \right) dS - \frac{1}{4\pi} \int_{S_B} \mathbf{n}(\mathbf{x}) \cdot \mathbf{Q}^\infty \left(\frac{1}{r} \right) dS = 0 \quad (3.21)$$

everywhere on the internal surface S_B . In practice, the unknown doublet distribution on the wake is usually related to the unknown doublet distribution on the body by the Kutta condition (see section 3.1.6).

3.1.2 Collocation point formulation

The Dirichlet and Neumann boundary conditions specify respectively the potential and velocity required over the surface. When applied to a numerical method, it is generally not possible to ensure these boundary conditions are met everywhere, but they can be approximated by applying them at discrete points on the surface referred to as collocation points. It is then reasonably assumed that the boundary conditions will be approximately satisfied in the regions between collocation points.

To construct a set of equations to represent the problem, we must determine the potential or velocity induced at each collocation point in terms of the unknown source or doublet distributions. In order to produce an exactly- or over-determined set of equations, and in the absence of any other constraints, we therefore require at least as many surface collocation points as there are unknowns in the function that describes the singularity distributions. The question of how best to distribute the collocation points over the surface has been the subject of much research (e.g. Lan, 1974; Bernasconi, 2007, chapter 5), and is discussed in more detail in chapter 4.

3.1.3 Galerkin formulation

An alternative to the collocation point formulation is the Galerkin method. In this approach, instead of applying the boundary conditions at discrete surface points, they are applied in a distributed sense over orthogonal basis functions that span the surface region. For example, in the case of a Neumann boundary condition on a surface, a series of orthogonally weighted integrals of the normal velocity over the surface is evaluated in terms of the unknown singularity distributions. Then by choosing appropriate singularity distributions, these weighted integrals are set to zero.

3.1.4 Comparison of Collocation and Galerkin formulations

For a given number of unknowns to model the singularity distributions, the Galerkin method is computationally much more expensive than the collocation point method to produce a square system of equations. This is because for each equation in the system, velocity influences must be evaluated at a distribution of points, from which the weighted integral is estimated, rather than a single collocation point. Despite some reported advantages in the accuracy of the Galerkin method, preliminary two dimensional studies did not produce significantly greater accuracy in comparison to the much quicker collocation point method, and so it has not been pursued to date. However, this may be one area worth returning to in further work.

Conventional constant panel methods employ the collocation method, usually with a square system of equations with one collocation point on each constant panel. It would be inappropriate to employ a Galerkin scheme within such a constant panel method as it would be mixing higher-order constraint equations with a low-order singularity distribution. Of the higher-order methods, the majority (e.g. Hess & Friedman, 1981; Johnson, 1980; Epton & Magnus, 1990) use a collocation formulation, often with an over-determined system comprising more collocation points than unknowns. However, the higher-order B-spline method by Maniar (1995) is based on a Galerkin approach.

3.1.5 Leading edge boundary conditions

In constant panel methods, collocation points are located within panel interiors, usually at the mid-point or the three-quarter chord point of each panel. On thin surfaces, for which the leading edge is represented by a row of panels supporting a doublet distribution, there are consequently no collocation points on the leading edge itself. In the case of the present higher-order method however, it is necessary to impose additional constraints towards the leading edge to produce an exactly or over-determined system of equations (see section 3.3.8). Whilst these additional constraints could be applied by additional collocation points close to the leading edge, it is more efficient and numerically stable to impose a direct condition of zero doublet strength along the leading edge itself. This constraint, whilst trivial to impose, requires some work to justify its application.

To demonstrate the condition of zero doublet strength at the leading edge of a thin surface, we consider the simple case of flat two dimensional plate of chord length c in steady uniform flow. The plate is aligned with the x_1 -axis in a cartesian reference frame (x_1, x_2) , and the onset flow of velocity Q^∞ makes an angle-of-attack α with the plate, as shown in figure 3.3. We seek a vortex distribution $\gamma(x_1)$ over the plate which will produce a velocity field satisfying the Neumann boundary condition, which here implies zero flow over the plate parallel to the x_2 -axis.

The velocity induced at $(y, 0)$ by a vortex element of strength Γ located at $(x, 0)$ is parallel to the x_2 -axis and of magnitude

$$q_2 = -\frac{\Gamma}{2\pi(y-x)} \quad (y \neq x) \quad (3.22)$$

The normal velocity induced by the distribution $\gamma(x)$ over the plate is therefore

$$q_2(y, 0) = \int_0^c -\frac{\gamma(x)}{2\pi(y-x)} dx \quad (3.23)$$




Figure 3.3: Glauert representation of inclined flow on a flat plate

By adding the component of the inflow velocity normal to the plate, the Neumann boundary condition can be written

$$\int_0^c -\frac{\gamma(x)}{2\pi(y-x)} dx = -Q^\infty \sin \alpha \quad (3.24)$$

Following the method of Glauert (1959), it is convenient to perform the coordinate transformation

$$x = \frac{c}{2}(1 - \cos \theta) \quad (3.25)$$

with

$$dx = \frac{c}{2} \sin \theta d\theta \quad (3.26)$$

The integral equation 3.24 is now transformed to

$$\int_0^\pi -\frac{\gamma(\theta) \sin \theta}{2\pi(\cos \theta - \cos \theta_y)} d\theta = -Q^\infty \sin \alpha \quad (3.27)$$

The solution to this equation is

$$\gamma(\theta) = 2Q^\infty \sin \alpha \left(\frac{1 + \cos \theta}{\sin \theta} \right) \quad (3.28)$$

which will be proved by writing 1 as $\cos 0\theta$ and substituting into the left-hand side of equation 3.27, producing

$$LHS = -\frac{Q^\infty \sin \alpha}{\pi} \int_0^\pi \left(\frac{\cos 0\theta + \cos \theta}{\sin \theta} \right) \left(\frac{\sin \theta}{\cos \theta - \cos \theta_y} \right) d\theta \quad (3.29)$$

The integration can now be carried out with the assistance of the *Glauert integral* (Glauert, 1959), which states

$$\int_0^\pi \frac{\cos n\theta}{\cos \theta - \cos \theta_y} d\theta = \frac{\pi \sin n\theta_y}{\sin \theta_y}, \quad n = 0, 1, 2, \dots \quad (3.30)$$

so that equation 3.29 becomes

$$LHS = -\frac{Q^\infty \sin \alpha}{\pi} \left[\frac{\pi \sin 0\theta_y + \pi \sin \theta_y}{\sin \theta_y} \right] = -Q^\infty \sin \alpha \quad (3.31)$$

which satisfies equation 3.27, thus verifying that the distribution $\gamma(\theta)$ of equation 3.28 is indeed correct.

From this vortex distribution we can obtain an equivalent doublet distribution, by recalling the equivalence relationship described in section 2.1. For this two dimensional case we can simply integrate the vortex distribution with respect its length coordinate x from the aerofoil's leading edge

$$\mu(\theta_y) = \int_0^{\theta_y} 2Q^\infty \sin \alpha \left(\frac{1 + \cos \theta}{\sin \theta} \right) dx \quad (3.32)$$

which, with the aid of equation 3.26, becomes

$$\mu(\theta_y) = \int_0^{\theta_y} 2Q^\infty \sin \alpha \left(\frac{1 + \cos \theta}{\sin \theta} \right) \frac{c}{2} \sin \theta d\theta \quad (3.33)$$

This can be integrated directly to obtain

$$\mu(\theta_y) = cQ^\infty(\theta_y + \sin \theta_y) \sin \alpha \quad (3.34)$$

In the vicinity of the leading edge $\theta_y \approx \sin \theta_y$ producing the approximation

$$\mu(\theta_y) \approx 2cQ^\infty \sin \alpha \sin \theta_y \quad \text{for small } \theta_y \quad (3.35)$$

which can be transformed back into an expression in terms of y to produce

$$\begin{aligned}\mu(y) &\approx 4cQ^\infty \sin \alpha \sqrt{\frac{y}{c} \left(1 - \frac{y}{c}\right)} \\ &\approx 4cQ^\infty \sin \alpha \sqrt{\frac{y}{c}} \quad \text{for small } y\end{aligned}\tag{3.36}$$

The equivalent vortex distribution is found by taking the derivative of the doublet distribution with respect to y resulting in

$$\gamma(y) \approx 2Q^\infty \sin \alpha \left(\frac{y}{c}\right)^{-\frac{1}{2}} \quad \text{for small } y\tag{3.37}$$

It is seen that whilst vortex strength is infinite at the leading edge, producing an infinite gradient (with x) in the doublet distribution, the doublet strength is zero. Note that despite the doublet distribution being the result of an integration of vorticity from the leading edge, it would not be zero valued here for all conceivable vortex distributions: for instance, a concentrated point vortex at the leading edge would correspond to a non-zero leading edge doublet strength. However in this case, we see that the singularity in the vortex distribution is less strongly singular than a concentrated vortex (although the resulting velocity field at the leading edge is still singular).

It can be shown (Glauert, 1959) that small variations in the geometry of the flat two dimensional aerofoil only produce terms in the vortex distribution that are zero at the leading-edge, so the boundary condition can be extended to any such two dimensional foil. We then make the reasonable assumption that the two dimensional boundary condition can be extended to the leading edge of a thin, finite-span aerofoil in three dimensions such that it can be applied to our higher-order method.

Thin aerofoil as the limit of a shrinking thick aerofoil

It is interesting to consider an alternative approach to obtaining the above result, in which we model a thin two or three dimensional aerofoil as the limit of a thick, smooth aerofoil as its thickness tends to zero. As the aerofoil's thickness reduces,

we denote the distance from the leading edge to be x , with surface normals $\mathbf{n}_u(x)$ on the upper surface and $\mathbf{n}_l(x)$ on the lower surface.

We imagine the thick aerofoil to be supporting a doublet distribution over its surface, $\mu_u(x)$ and $\mu_l(x)$. As the aerofoil's geometry is smooth, the velocity field and doublet distribution will be continuous, and in particular at the rounded leading edge

$$\mu_u(0) = \mu_l(0) \quad \text{and} \quad \left. \frac{d\mu_u}{dx} \right|_{x=0} = - \left. \frac{d\mu_l}{dx} \right|_{x=0} \quad (3.38)$$

where the negative sign arises because the sign of dx changes at the leading edge as we trace a point from the upper surface to the lower surface.

Now, as the thickness t tends to zero, $\mathbf{n}_l \rightarrow -\mathbf{n}_u$, and the doublet distributions on the two surfaces can be combined to form

$$\mu(x)\mathbf{n}_u = \mu_u(x)\mathbf{n}_u(x) - \mu_l(x)\mathbf{n}_u(x) \quad (3.39)$$

If we assume that in the limit the continuity equations 3.38 still holds, we can deduce

$$\mu(x) = 0 \quad (3.40)$$

and

$$\left. \frac{d\mu}{dx} \right|_{x=0} = \left. \frac{d\mu_u}{dx} \right|_{x=0} - \left. \frac{d\mu_l}{dx} \right|_{x=0} = 2 \left. \frac{d\mu_u}{dx} \right|_{x=0} \quad (3.41)$$

This should not be regarded as a rigorous demonstration of zero doublet strength on the leading edge of an arbitrary aerofoil, as we have rather casually applied the assumptions of a smooth geometry to the limiting case of a leading edge that tends to infinite curvature. However, by comparing the two dimensional case with that described by the Glauert representation, it is interesting to see that equation 3.38 does indeed hold, although as the thickness decreases the doublet distribution gradients tend to infinite gradient around the infinite curvature of the leading edge. Both approaches also show that whilst the strong singularity of a non-zero doublet strength leading edge is avoided, the weaker singularity (in terms of the resulting velocity field) of a non-zero doublet gradient exists.

3.1.6 Trailing edges and wakes

Wake geometry

Wakes shed from aerofoils are modelled in a similar way to the aerofoils themselves. A separate patch is defined comprising both geometric and singularity surface distributions. Since the wake represents a shear layer across which there should be no change in normal velocity, the singularity distribution is limited to a doublet surface, and can include no sources.

The wake should ideally be aligned with the velocity flow-field, and schemes to relax a wake to achieve this are discussed in detail in chapter 5. However, a fixed pre-defined wake shape can also be modelled, which in many cases will produce an acceptable solution.

It is convenient to define the wake such that the stream direction (or assumed stream direction for an un-relaxed wake) is aligned with one of the two intrinsic patch coordinates ξ_2 . We can then assume that wake lines of constant ξ_1 are streamlines, or for a relaxing wake, adjust the wake shape until this condition is met. As all vorticity shed from the trailing-edge of the wing is carried downstream in the wake parallel to the streamlines, this alignment of ξ_2 implies that, for a steady-state solution, the doublet strength is constant in this direction. This allows the wake doublet surface to be modelled as a function of ξ_1 only. (For a dynamic solution, in which doublet strength solutions at the trailing edge are propagated downstream, lines of constant ξ_2 would represent cross-sections of the wake that were shed at the same point in time.)

Kutta condition

The well-established Kutta condition (for example see Katz & Plotkin, 2001) states that flow leaves the sharp trailing edge of an aerofoil smoothly, and for a thin wing or cusped trailing edge of a thick wing this implies that the geometry must be continuous in value and gradient over the junction between the wing and wake. For a thick wing with a finite-angled trailing edge, we interpret the condition to mean that the wake must be aligned with its angle bisector. The Kutta condition also states that at the trailing edge of a body, surface velocities should be finite.

This has different modelling implications for thin and thick bodies as described in the following two sections.

In the equations that follow, note that the edge of a wake attached to the trailing edge of a body is given by

$$\xi^w = (\xi_1^w, 0) \quad (3.42)$$

It is then useful to write the equation describing the trailing edge of the body as

$$\xi_{\text{TE}}^b = \xi_{\text{TE}}^b(\xi_1^w) \quad (3.43)$$

such that the line joining the body and wake is

$$\mathbf{y}^b(\xi_{\text{TE}}^b(\xi_1^w)) = \mathbf{y}^w(\xi_1^w, 0) \quad (3.44)$$

Trailing edge conditions for a thin body

The Kutta condition for a thin body implies that the doublet distribution should be continuous in both value and gradient across the boundary on to the wake, since a step change in either produces a singular velocity distribution. As the wake's streamwise doublet distribution is constant in the steady-state, and therefore has zero streamwise gradient, the doublet strength gradient at the wing's trailing edge should also be zero in a direction parallel to the wake streamlines,

$$\mu^w(\xi_1^w, \xi_2^w) = \mu^b(\xi_{\text{TE}}^b(\xi_1^w)) \quad (3.45)$$

and

$$\left. \frac{\partial \mu^b}{\partial \xi_2^w} \right|_{\xi^b = \xi_{\text{TE}}^b} = 0 \quad (3.46)$$

When wake relaxation is modelled this condition can only be applied as part of the wake relaxation iteration cycle, as the direction in which the gradient is measured will change. However in practice, the intrinsic coordinate grid on the body can sometimes be constructed such that ξ_2^b is aligned with the

anticipated flow direction at the trailing edge, and the gradient condition can be approximated to

$$\left. \frac{\partial \mu^b}{\partial \xi_2^b} \right|_{\xi^b = \xi_{TE}^b} = 0 \quad (3.47)$$

Trailing edge conditions for a thick aerofoil

In this section we replace the subscripts ‘TE’ with ‘UTE’ and ‘LTE’ for the upper and lower sides of a body’s trailing edge.

Continuous doublet strength In the absence of a wake at the trailing edge, the difference between the doublet strength on the upper and lower surface would form a vortex line of strength

$$\gamma(\xi_1^w) = \mu^b(\xi_{UTE}^b) - \mu^b(\xi_{LTE}^b) \quad (3.48)$$

In order to maintain a finite velocity field, this vortex must be cancelled by the doublet wake sheet of strength

$$\mu^w(\xi_1^w, \xi_2^w) = \mu^b(\xi_{UTE}^b) - \mu^b(\xi_{LTE}^b) \quad (3.49)$$

This is a ‘standard’ Kutta condition generally applied in constant panel methods, and in the present higher order method.

Zero pressure jump Pressures must be equal on the upper and lower surfaces at the trailing edge, otherwise an infinite pressure gradient would exist. Thus, the upper and lower surface velocities must have equal magnitude at the trailing edge. Since velocity magnitude is not a linear function of singularity strength (see equation 3.142), this condition would need to be applied iteratively, which would be extremely inefficient for a solution without wake relaxation that would otherwise be linear. However, for a two dimensional case, or a three dimensional case where the trailing edge flow direction can be approximately stated beforehand, the condition can be simplified. For instance, if the flow at

the trailing edge is parallel to lines of constant ξ_1^b , it can be written as

$$\mathbf{Q}^\infty \cdot \mathbf{e}_2^b(\xi_{\text{UTE}}^b) + \left. \frac{\partial \mu^b}{\partial \xi_2^b} \right|_{\xi^b = \xi_{\text{UTE}}^b} = -\mathbf{Q}^\infty \cdot \mathbf{e}_2^b(\xi_{\text{LTE}}^b) - \left. \frac{\partial \mu^b}{\partial \xi_2^b} \right|_{\xi^b = \xi_{\text{LTE}}^b} \quad (3.50)$$

where \mathbf{e}_2^b is a unit vector in the direction of ξ_2^b , and the minus signs result from ξ_2^b pointing forwards on the upper trailing edge and backwards on the lower edge, or vice versa.

This zero pressure-jump condition is generally not applied in constant panel methods, and it has been found that in this higher-order method, valid solutions can be obtained without it being explicitly imposed. However, it provides a valid constraint equation that can replace a more computationally expensive equation (such as an extra collocation point condition) with no loss in accuracy, and is therefore employed in this method.

Stagnation point For the case of a thick foil with a finite trailing edge angle, the trailing edge must be a stagnation point for velocities to be continuous. This requirement is more stringent than the zero pressure jump condition above, forcing both velocities to be zero, rather than equal. Whilst physically valid, this boundary condition is so dominant within a discretised problem formulation that it can be detrimental to overall accuracy. It is therefore unhelpful to impose it within the potential flow solution.

3.1.7 Boundary conditions on other edges

Conditions for the leading and trailing edges of bodies have been discussed in detail, but as yet no consideration has been given to other body edges that are perhaps aligned with the flow, rather than perpendicular to it. In fact, we will treat all edges as either leading or trailing edges, classifying each edge according to whether flow is expected to be towards it or away from it. In principle, for cases where the flow is approximately parallel to the edge, either classification can be chosen without hugely affecting the results; however, the choice would normally be made based on considerations of required accuracy, stability and run-time.

For example, consider the case of an inclined rectangular flat plate, aligned with the free-stream. A wake will clearly be shed from the trailing-edge, but flow is expected to be approximately parallel to each of the two side edges. If a side was treated as a trailing-edge, a wake would be shed along its length, propagating the plate's doublet distribution downstream. In particular, at the front corner of the plate, the leading edge condition of zero doublet strength would propagate along one edge of the side's shed wake. Now as the side wake relaxed, it would compress into a very thin region approximately aligned with the side of the plate, and the whole side of the body would be connected over a short (albeit potentially tangled) distance to a line of zero doublet strength.

Compare this with the situation that would occur if the side had been classified as a leading edge: in this case, a boundary condition of zero doublet strength would be applied along the side. In terms of the doublet solution on the wing these two situations are, in principle, similar, and similar results should be obtained. However, the very significant distortion and tight wrapping that will exist in a wake shed from an edge parallel to the flow direction is likely to produce numerical instabilities, and modelling the edge without a wake would normally be preferable in this situation.

3.1.8 The concept of a patch

In defining the geometry of the bounding surfaces, we refer to the basic building block of the method as a patch. A patch is a 'large' $C1$ -continuous surface (one with continuous slope) over which a source or doublet singularity distribution, or both, may exist. A patch is parameterised by the intrinsic coordinates $\xi = (\xi_1, \xi_2)$ which always span the rectangular parametric space

$$0 \leq \xi_1 \leq 1 \tag{3.51}$$

$$0 \leq \xi_2 \leq 1 \tag{3.52}$$

as shown in figure 3.4.

The geometric surface of the patch can be represented by any continuous function. For example, it could be a NURBS surface (Piegl & Tiller, 1997), or

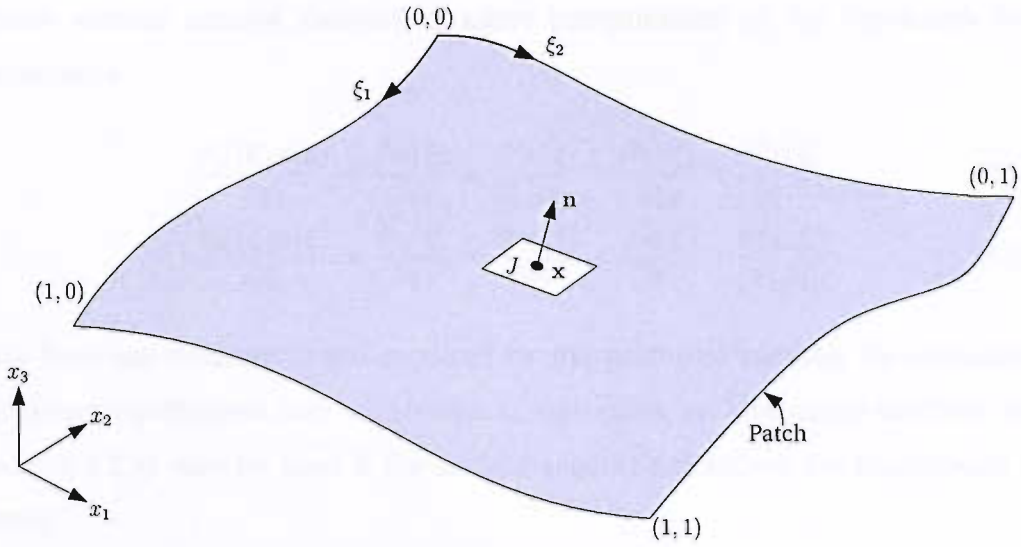


Figure 3.4: A geometric patch

a surface created by smooth interpolation into a discrete set of coordinates. The only requirements of the surface function are that the following surface parameters can be evaluated for any point ξ

$$\mathbf{x}(\xi) \quad \frac{\partial \mathbf{x}(\xi)}{\partial \xi_1} \quad \frac{\partial \mathbf{x}(\xi)}{\partial \xi_2} \quad (3.53)$$

where $\mathbf{x}(\xi)$ are the physical coordinates of the surface. From the tangents, the two other fundamental surface parameters required by the method can be readily obtained, these being the surface unit normal $\mathbf{n}(\xi)$ and the Jacobian determinant $J(\xi)$ relating the physical coordinates to the intrinsic coordinates.

$$J(\xi)\mathbf{n}(\xi) = \frac{\partial \mathbf{x}(\xi)}{\partial \xi_1} \times \frac{\partial \mathbf{x}(\xi)}{\partial \xi_2} \quad (3.54)$$

Ideally, the following additional surface derivatives could also be obtained from the surface function

$$\frac{\partial^2 \mathbf{x}(\xi)}{\partial \xi_1^2} \quad \frac{\partial^2 \mathbf{x}(\xi)}{\partial \xi_1 \xi_2} \quad \frac{\partial^2 \mathbf{x}(\xi)}{\partial \xi_2^2} \quad (3.55)$$

These surface second derivatives allow computation of the Jacobian's first derivatives

$$\frac{\partial J(\boldsymbol{\xi})\mathbf{n}(\boldsymbol{\xi})}{\partial \xi_1} = \frac{\partial \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_1} \times \frac{\partial^2 \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_1 \partial \xi_2} - \frac{\partial \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_2} \times \frac{\partial^2 \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_1^2} \quad (3.56)$$

$$\frac{\partial J(\boldsymbol{\xi})\mathbf{n}(\boldsymbol{\xi})}{\partial \xi_2} = \frac{\partial \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_1} \times \frac{\partial^2 \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_2^2} - \frac{\partial \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_2} \times \frac{\partial^2 \mathbf{x}(\boldsymbol{\xi})}{\partial \xi_1 \partial \xi_2} \quad (3.57)$$

The Jacobian derivatives are required by the preferred method for evaluating influence coefficients (see section 3.2.3). However, an alternative method (see section 3.2.8) may be used if the surface second derivatives are unavailable or 'noisy'.

A patch will generally represent an entire body, and may be thin (like a sail) or thick (like an aerofoil). Two or more patches can be connected to produce more complex shapes, and wakes are modelled by one or more separate patches attached to a body patch. Although the underlying intrinsic coordinates $\boldsymbol{\xi}$ lie on a rectangular grid, it is permissible for the patch to taper to zero width at a vertex to form a three sided surface. At such a vertex the surface normal must still be defined, and the Jacobian determinant will be zero.

With a body's surface modelled as one or more continuous patches, there is no requirement to define any further geometric grid or discretisation over the body. This is in contrast to conventional constant or higher-order panel methods, which require the surface to be approximated by several flat, twisted or low-order polynomial panels. The ability to model a truly curved surface, without discretisation, is one of the principal advantages of this method.

3.1.9 Singularity distributions

The strength of a source or doublet distribution over a patch is modelled by an additional surface of the same intrinsic coordinates $\boldsymbol{\xi}$. Figure 3.5 illustrates the correspondence between the geometric and singularity surfaces, in this case showing how the three dimensions of the geometry can be viewed as separate functions of $\boldsymbol{\xi}$.




Figure 3.5: Geometric and singularity surfaces are smooth functions of the intrinsic coordinates ξ .

As the singularity distributions form the unknowns of the problem to which we seek a solution, the surfaces that represent them must be parameterised in a way sufficiently flexible to allow physically realistic distributions to be modelled. Beyond this requirement, the singularity surfaces can be of any form which is at least $C1$ -continuous. They could be generic, or specifically tailored to model a restricted family of solution surfaces, should prior knowledge of the solution shape be known.

For a given set of parameters \mathbf{B}^μ or \mathbf{B}^σ , the functional description of the surface should provide the value of the distribution at any point ξ

$$\mu = \mu(\xi, \mathbf{B}^\mu) \quad (3.58)$$

$$\sigma = \sigma(\xi, \mathbf{B}^\sigma) \quad (3.59)$$

The function must also be capable of providing the following derivatives

$$\frac{\partial \mu(\boldsymbol{\xi})}{\partial \xi_1} \quad \frac{\partial \mu(\boldsymbol{\xi})}{\partial \xi_2} \quad \frac{\partial \sigma(\boldsymbol{\xi})}{\partial \xi_1} \quad \frac{\partial \sigma(\boldsymbol{\xi})}{\partial \xi_2} \quad (3.60)$$

In section 3.2, the method of calculating influence coefficients is outlined with reference to a generic singularity surface. In later sections, a (bi-cubic) B-spline surface is chosen to demonstrate an implementation of the method.

3.2 Influence functions

Whatever solution method is used, the potential or velocity influence due to the singularity distributions ultimately needs to be calculated at a number of discrete points. For a collocation solution, these are the collocation points, and for a Galerkin solution, they are the integration points used to approximate the Galerkin integral via a quadrature scheme. We will generally refer to collocation points, but this should be understood to mean either collocation or Galerkin integration points.

In this section we first set out the influences at a point resulting from point sources or doublets, then consider how the effect of source or doublet sheets can be integrated. It will be seen that this integration requires special attention when the collocation point lies on the surface containing the singularity distribution.

3.2.1 Velocity influences due to point doublets and sources

Recalling equation 3.2, the velocity potential $\phi(\mathbf{y})$ induced at \mathbf{y} by a doublet element $\mu(\mathbf{x})$ is

$$\phi(\mathbf{y}) = -\frac{1}{4\pi} \frac{\mathbf{r}(\mathbf{y}, \mathbf{x}) \cdot \mathbf{n}(\mathbf{x}) \mu(\mathbf{x})}{r^3(\mathbf{y}, \mathbf{x})} \quad (3.61)$$

where $\mathbf{n}(\mathbf{x})$ is the orientation of the doublet, or the normal of the surface with which the doublet is associated. The velocity influence of the doublet distribution

is obtained by differentiating the velocity potential

$$\begin{aligned}
\mathbf{q} = \nabla\phi &= -\frac{\mu}{4\pi} \nabla \frac{\mathbf{r} \cdot \mathbf{n}}{r^3} \\
&= -\frac{\mu}{4\pi} \left[\frac{\nabla(\mathbf{r} \cdot \mathbf{n})}{r^3} + \mathbf{r} \cdot \mathbf{n} \nabla \left(\frac{1}{r^3} \right) \right] \\
&= \frac{\mu}{4\pi} \left[\frac{3\mathbf{n} \cdot \mathbf{r}}{r^5} \mathbf{r} - \frac{1}{r^3} \mathbf{n} \right]
\end{aligned} \tag{3.62}$$

where the functional dependencies have been omitted for clarity.

Similarly, from equation 3.3, the velocity potential induced by a source element $\sigma(\mathbf{x})$ is

$$\phi(\mathbf{y}) = -\frac{\sigma(\mathbf{x})}{4\pi r(\mathbf{y}, \mathbf{x})} \tag{3.63}$$

The associated velocity influence is therefore

$$\begin{aligned}
\mathbf{q} = \nabla\phi &= -\frac{\sigma}{4\pi} \nabla \frac{1}{r} \\
&= \frac{\sigma}{4\pi} \frac{\mathbf{r}}{r^3}
\end{aligned} \tag{3.64}$$

3.2.2 Kernel function singularities

In the subsequent analysis it is convenient to represent a generic singularity distribution by $\tau(\mathbf{x})$, where

$$\tau(\mathbf{x}) = \mu(\mathbf{x}) \quad \text{or} \quad \tau(\mathbf{x}) = \sigma(\mathbf{x}) \tag{3.65}$$

We also separate the singularity distribution strength from the ‘kernel function’ describing the influence of a point singularity of unit strength, and represent this function by $T(\mathbf{y}, \mathbf{x})$ or $\mathbf{T}(\mathbf{y}, \mathbf{x})$. A summary of the kernel functions for source and doublet, potential and velocity influences is shown in table 3.1. Finally, we use $\mathbf{q}(\mathbf{y})$ to represent both potential and velocity influences due to source and doublet distributions.

Armed with our new terminology, we can consider the influence at a point \mathbf{y} of a singularity distribution over a surface S , which is simply the integration of the

influence of point singularities given in the previous section

$$\mathbf{q}(\mathbf{y}) = \int_S \tau(\mathbf{x}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) \quad (3.66)$$

For a general field point not lying on the surface S , this is a regular integral and can be evaluated directly (see section 3.3.4 for details of an appropriate numerical method). However, a problem arises when \mathbf{y} lies on the surface S , in which case the influence expression becomes singular for all kernel functions at the point $\mathbf{y} = \mathbf{x}$ ($r = 0$). This is highly significant for our potential flow method, which relies on evaluation of kernel function influences at discrete points on the singularity surface itself.

The source potential kernel function is of the form $\frac{1}{r}$ in three dimensions, or $\ln r$ in two dimensions, and is said to be weakly singular. The three dimensional doublet potential kernel is of the form $\frac{1}{r^2}$ and appears to be strongly singular, but in fact the numerator contains the term $\mathbf{r} \cdot \mathbf{n}$ which is zero as $r \rightarrow 0$, and this lowers the order to weakly singular. The source velocity kernel is also of the form $\frac{1}{r^2}$, but is a vector quantity without a vanishing numerator, and is therefore strongly singular. The doublet velocity influence is one order higher, with a kernel function of the form $\frac{1}{r^3}$, and is said to be hypersingular. The corresponding two dimensional kernel functions shown in table 3.1, whilst being of different order of r , are always of the same order of singularity as their three dimensional counterparts. The classification of order of singularity is relevant to the method of integration, as will become apparent in section 3.2.3.

In conventional panel methods, these functions have only been integrated for special cases such as flat or skewed quads with constant or low-order singularity distributions. However, the method that follows in section 3.2.3 shows how the integration of all kernel functions can be transformed into a sum of regular integrals, each of which can be evaluated by standard numerical techniques. This method is applicable to any continuous singularity distribution over any continuous geometric surface.

3.2.3 Integration of three dimensional hypersingular kernel functions

The method described here applies to the integration of all the relevant singular kernel functions, although some simplifications can be made for kernel functions that are less than hypersingular, as described in section 3.2.5. The scheme can be employed to solve both two and three dimensional kernel integrals, although solution in three dimensions is naturally a little more complicated. As such, the three dimensional case is detailed in full here, and the differences required for a two dimensional analysis are outlined in section 3.2.4. The method described is that developed by Guiggiani, and further details can be found in the series of papers Guiggiani & Casalini (1987); Guiggiani & Gigante (1990); Guiggiani *et al.* (1990); Guiggiani (1995). A good overview of the scheme can be found in Guiggiani (1998).

To begin, we recall the potential or velocity influence at a point \mathbf{y} due to a singularity distribution over a surface $S(\mathbf{x})$

$$\mathbf{q}(\mathbf{y}) = \int_S \tau(\mathbf{x})\mathbf{T}(\mathbf{y}, \mathbf{x})dS(\mathbf{x}) \quad (3.67)$$

where $\tau(\mathbf{x}) = \sigma(\mathbf{x})$ or $\tau(\mathbf{x}) = \mu(\mathbf{x})$ is the strength of the distribution. In order to proceed with the integration it is helpful to consider a surface which just excludes the singular point \mathbf{y} , rather than one which passes through the point. This is achieved by replacing a small region e_ϵ of the surface surrounding \mathbf{y} by the surface s_ϵ of a partial sphere of radius ϵ centred on \mathbf{y} , and considering the limiting behaviour of the integral as $\epsilon \rightarrow 0$. With reference to figure 3.6, the integral can now be written

$$\mathbf{q} = \lim_{\epsilon \rightarrow 0} \left\{ \int_{S-e_\epsilon} \tau(\mathbf{x})\mathbf{T}(\mathbf{y}, \mathbf{x})dS(\mathbf{x}) + \int_{s_\epsilon} \tau(\mathbf{x})\mathbf{T}(\mathbf{y}, \mathbf{x})dS(\mathbf{x}) \right\} \quad (3.68)$$

It is shown in section 3.2.6 that the integration over the small spherical region e_ϵ results in up to three 'free' terms

$$\mathbf{q}_{e_\epsilon} = \tau(\mathbf{y}) \left[\frac{\mathbf{b}(\mathbf{y})}{\epsilon} + \mathbf{a}(\mathbf{y}) \right] + \nabla\tau(\mathbf{y}) \cdot \mathbf{c}(\mathbf{y}) + O(\epsilon) \quad (3.69)$$





Figure 3.6: The physical and intrinsic representations of the surface, with a small region surrounding the singular point replaced by a spherical surface.

For the moment, the component $c(\mathbf{y})$ is neglected, and the remaining free terms can replace the integral over e_ϵ in equation 3.68 to produce

$$\mathbf{q} = \lim_{\epsilon \rightarrow 0} \left\{ \int_{S-e_\epsilon} \tau(\mathbf{x}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) + \frac{\mathbf{b}(\mathbf{y})}{\epsilon} \tau(\mathbf{y}) \right\} + \mathbf{a}(\mathbf{y}) \tau(\mathbf{y}) \quad (3.70)$$

We now consider the integral over $S - e_\epsilon$, and begin by transferring it into the intrinsic coordinate space of the surface, $\xi = (\xi_1, \xi_2)$. In this transformation, the plane region R is the image of the surface S , and σ_ϵ is the image of the circular region e_ϵ surrounding the singular point. Note that whilst both e_ϵ and σ_ϵ have the singular point at their centres and shrink to zero as $\epsilon \rightarrow 0$, σ_ϵ is generally not circular, but elliptical as shown in figure 3.6. With Jacobian determinant of the transformation $J(\xi)$, the integral becomes

$$\mathbf{q} = \lim_{\epsilon \rightarrow 0} \left\{ \int_{R-\sigma_\epsilon} \tau(\xi) \mathbf{T}(\mathbf{y}, \xi) J(\xi) d\xi_1 d\xi_2 + \frac{\mathbf{b}(\mathbf{y})}{\epsilon} \tau(\mathbf{y}) \right\} + \mathbf{a}(\mathbf{y}) \tau(\mathbf{y}) \quad (3.71)$$

We then define a set of polar coordinates within the intrinsic frame, centred on the singular point η (image of \mathbf{y}), with radius ρ and angle θ :

$$\xi_1 = \eta_1 + \rho \cos \theta \quad (3.72a)$$

$$\xi_2 = \eta_2 + \rho \sin \theta \quad (3.72b)$$

Re-writing the singular integral in terms of the polar coordinates we obtain

$$\mathbf{q} = \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \mathbf{F}(\rho, \theta) d\rho d\theta + \frac{\mathbf{b}(\eta)}{\epsilon} \tau(\eta) \right\} + \mathbf{a}(\mathbf{y}) \tau(\mathbf{y}) \quad (3.73)$$

where $\mathbf{F}(\rho, \theta) = \tau(\xi) \mathbf{T} J(\xi) \rho$ and $\rho d\rho d\theta = d\xi_1 d\xi_2$. The limits $\rho = \alpha(\epsilon, \theta)$ and $\bar{\rho}(\theta)$ are the equations of the boundaries of σ_ϵ and S respectively.

We now employ a Laurent series expansion of $\mathbf{F}(\rho, \theta)$ to separate terms that are singular as $\epsilon \rightarrow 0$ from the non-singular remainder:

$$\mathbf{F}(\rho, \theta) = \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} + \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \quad (3.74)$$

Explicit expressions for \mathbf{F}_{-1} and \mathbf{F}_{-2} are derived in appendix A. Substituting this expansion into equation 3.73 we can write

$$\begin{aligned} \mathbf{q} &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \left[\mathbf{F}(\rho, \theta) - \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} - \frac{\mathbf{F}_{-1}(\theta)}{\rho} \right] d\rho d\theta \right. \\ &\quad + \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \frac{\mathbf{F}_{-1}(\theta)}{\rho} d\rho d\theta \\ &\quad \left. + \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \left[\frac{\mathbf{F}_{-2}(\theta)}{\rho^2} d\rho d\theta + \frac{\mathbf{b}(\boldsymbol{\eta})}{\epsilon} \tau(\boldsymbol{\eta}) \right] \right\} + \mathbf{a}(\mathbf{y})\tau(\mathbf{y}) \\ &= \mathbf{I}_0 + \mathbf{I}_{-1} + \mathbf{I}_{-2} + \mathbf{a}(\mathbf{y})\tau(\mathbf{y}) \end{aligned} \quad (3.75)$$

We now consider each of the integrals \mathbf{I}_0 , \mathbf{I}_{-1} and \mathbf{I}_{-2} separately.

As shown by the Laurent expansion in equation 3.74, \mathbf{I}_0 does not contain any singularities and applying the limit is trivial:

$$\begin{aligned} \mathbf{I}_0 &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \left[\mathbf{F}(\rho, \theta) - \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} - \frac{\mathbf{F}_{-1}(\theta)}{\rho} \right] d\rho d\theta \right\} \\ &= \int_0^{2\pi} \int_0^{\bar{\rho}(\theta)} \left[\mathbf{F}(\rho, \theta) - \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} - \frac{\mathbf{F}_{-1}(\theta)}{\rho} \right] d\rho d\theta \end{aligned} \quad (3.76)$$

The resulting integral is regular, and can be integrated directly by means of a generic cubature formula (see section 3.3.5). Note however that it is crucial that the cubature is carried out across the variables ρ and θ (rather than $\boldsymbol{\xi}$). This effectively concentrates integration points towards the collocation point \mathbf{y} , correctly assigning more ‘weight’ to the region closer to the singularity.

In order to integrate \mathbf{I}_{-1} and \mathbf{I}_{-2} we first introduce a Taylor series expansion for $\alpha(\epsilon, \theta)$, the equation for ρ describing the boundary of σ_ϵ

$$\alpha(\epsilon, \theta) = \epsilon\beta(\theta) + \epsilon^2\gamma(\theta) + O(\epsilon^3) \quad (3.77)$$

Explicit expressions for $\beta(\theta)$ and $\gamma(\theta)$ are given in appendix A.

We are now able to perform the integration of \mathbf{I}_{-1} with respect to ρ analytically

$$\begin{aligned}\mathbf{I}_{-1} &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \frac{\mathbf{F}_{-1}(\theta)}{\rho} d\rho d\theta \right\} \\ &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \mathbf{F}_{-1}(\theta) [\ln \bar{\rho}(\theta) - \ln \alpha(\epsilon, \theta)] d\theta \right\}\end{aligned}\quad (3.78)$$

Applying the limit on ϵ to terms independent of ϵ , and employing the Taylor expansion of equation 3.77, we can reduce \mathbf{I}_{-1} to

$$\mathbf{I}_{-1} = \int_0^{2\pi} \mathbf{F}_{-1}(\theta) \ln \bar{\rho}(\theta) d\theta - \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \mathbf{F}_{-1}(\theta) \ln \epsilon \beta(\theta) d\theta \right\}\quad (3.79)$$

It can be shown from formulae in appendix A that $F_{-1}(\theta) = -F_{-1}(\theta + \pi)$. It follows that

$$\lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \mathbf{F}_{-1}(\theta) \ln \epsilon d\theta \right\} = 0\quad (3.80)$$

and hence,

$$\mathbf{I}_{-1} = \int_0^{2\pi} \mathbf{F}_{-1}(\theta) \ln \frac{\bar{\rho}(\theta)}{\beta(\theta)} d\theta\quad (3.81)$$

In order to treat \mathbf{I}_{-2} , we require a series expansion for $\alpha^{-1}(\epsilon, \theta)$ which we obtain by taking a Taylor series of the existing expansion given by equation 3.77

$$\begin{aligned}\frac{1}{\alpha(\epsilon, \theta)} &= \frac{1}{\epsilon\beta(\theta) + \epsilon^2\gamma(\theta) + O(\epsilon^3)} \\ &= \frac{1}{\epsilon\beta(\theta)} \times \frac{1}{1 + \epsilon\gamma(\theta)/\beta(\theta) + O(\epsilon^2)} \\ &= \frac{1}{\epsilon\beta(\theta)} \left[1 - \frac{\epsilon\gamma(\theta)}{\beta(\theta)} \right] + O(\epsilon)\end{aligned}\quad (3.82)$$

This enables us to treat \mathbf{I}_{-2} in a similar manner to \mathbf{I}_{-1}

$$\begin{aligned}\mathbf{I}_{-2} &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \int_{\alpha(\epsilon, \theta)}^{\bar{\rho}(\theta)} \left[\frac{\mathbf{F}_{-2}(\theta)}{\rho^2} d\rho d\theta + \frac{\mathbf{b}(\eta)}{\epsilon} \tau(\eta) \right] \right\} \\ &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \mathbf{F}_{-2}(\theta) \left[-\frac{1}{\bar{\rho}(\theta)} + \frac{1}{\alpha(\epsilon, \theta)} \right] d\theta + \frac{\mathbf{b}(\eta)}{\epsilon} \tau(\eta) \right\} \\ &= \lim_{\epsilon \rightarrow 0} \left\{ \int_0^{2\pi} \frac{\mathbf{F}_{-2}(\theta)}{\epsilon\beta(\theta)} d\theta + \frac{\mathbf{b}(\eta)}{\epsilon} \tau(\eta) \right\} - \int_0^{2\pi} \mathbf{F}_{-2}(\theta) \left[\frac{\gamma(\theta)}{\beta^2(\theta)} + \frac{1}{\bar{\rho}(\theta)} \right] d\theta\end{aligned}\quad (3.83)$$

It is reasoned by Guiggiani *et al.* (1990) that as the final integral must be finite, the two parts of the singular term in equation 3.83 must cancel out such that

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left\{ \int_0^{2\pi} \frac{\mathbf{F}_{-2}(\theta)}{\beta(\theta)} d\theta + \mathbf{b}(\boldsymbol{\eta})\tau(\boldsymbol{\eta}) \right\} = 0 \quad (3.84)$$

leaving the expression

$$\mathbf{I}_{-2} = - \int_0^{2\pi} \mathbf{F}_{-2}(\theta) \left[\frac{\gamma(\theta)}{\beta^2(\theta)} + \frac{1}{\bar{\rho}(\theta)} \right] d\theta \quad (3.85)$$

The expressions for \mathbf{I}_{-1} and \mathbf{I}_{-2} , given by equations 3.81 and 3.85, are both one-dimensional regular integrals, that can be evaluated numerically by standard quadrature formulae (see section 3.3.5).

Terms can now be collected to produce the final expression

$$\begin{aligned} \mathbf{q} = & \int_0^{2\pi} \int_0^{\bar{\rho}(\theta)} \left\{ \mathbf{F}(\rho, \theta) - \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} - \frac{\mathbf{F}_{-1}(\theta)}{\rho} \right\} d\rho d\theta \\ & + \int_0^{2\pi} \left\{ \mathbf{F}_{-1}(\theta) \ln \frac{\bar{\rho}(\theta)}{\beta(\theta)} - \mathbf{F}_{-2}(\theta) \left[\frac{\gamma(\theta)}{\beta^2(\theta)} + \frac{1}{\bar{\rho}(\theta)} \right] \right\} d\theta + \mathbf{a}(\mathbf{y})\tau(\mathbf{y}) \end{aligned} \quad (3.86)$$

where all the terms are regular and can be evaluated numerically.

3.2.4 Integration of two dimensional hypersingular kernel functions

The principals of the two dimensional method are equivalent to those of the three dimensional version, and once again we follow the method of Guiggiani *et al.* (1990) with modifications to model a continuous surface rather than discrete elements. The integration is along a curved line, rather than over a curved surface, and a shrinking circular region of radius ϵ around the singular point \mathbf{y} is excluded, as shown in figure 3.7. In intrinsic coordinates, the distance from the singular point $\boldsymbol{\eta}$ (image of \mathbf{y}) is measured by ρ , and the extent of the excluded region is given by the equation

$$\rho = \alpha_m(\epsilon) \quad (3.87)$$

There can be no transformation to polar coordinates within the one dimensional parameter plane, and in place of the integration with respect to θ , which swept




Figure 3.7: Integration of a two dimensional kernel function over a surface. Integration over the regions e_ϵ^m are replaced by an integration over the spherical region s_ϵ .

around the singular point, we simply add components from the two regions either side of y . The regions are identified by a suffix $m = 1$ or $m = 2$ and in general, where three dimensional quantities were continuous functions of θ , their two dimensional counterparts are only dependent on the value m . It is important to note that if the Jacobian determinant, tangent vector or curvature is discontinuous at y (e.g. where a wake flows from a finite trailing edge), the region s_ϵ will not be symmetric about η in the intrinsic plane, and $\alpha_1(\epsilon) \neq \alpha_2(\epsilon)$. However, if y is an interior point on a smooth curve, the analysis presented below can be simplified.

Following the same initial arguments of the three dimensional case, we arrive at the two dimensional equivalent of equation 3.73 which gives the total influence as the sum of limiting terms as $\epsilon \rightarrow 0$ and a free-term coefficient.

$$\mathbf{q} = \lim_{\epsilon \rightarrow 0} \left\{ \int_{\zeta_1}^{\eta - \alpha_1(\epsilon)} \mathbf{F}^1(\eta, \xi) d\xi + \int_{\eta + \alpha_1(\epsilon)}^{\zeta_2} \mathbf{F}^2(\eta, \xi) d\xi + \frac{\mathbf{b}(\eta)}{\epsilon} \tau(\eta) \right\} + \mathbf{a}(\mathbf{y}) \tau(\mathbf{y}) \quad (3.88)$$

where $\mathbf{F}^m(\eta, \xi) = \tau(\xi) \mathbf{T}_m(\mathbf{x}(\eta), \mathbf{x}(\xi)) J_m(\xi)$. The limits $\xi = \zeta_1$ and $\xi = \zeta_2$ are the two extents of the total region containing the singularity distribution, and it is assumed that $\zeta_1 < \eta < \zeta_2$. If the singular point y falls on the join between two patches, this will not necessarily hold, in which case some minor changes to signs and integration limits should be made to the procedure outlined below.

The Laurent series expansion of $\mathbf{F}^m(\eta, \xi)$ is given by

$$\mathbf{F}^m(\eta, \xi) = \frac{\mathbf{F}_{-2}^m(\eta, \xi)}{\rho^2} + \frac{\mathbf{F}_{-1}^m(\eta, \xi)}{\rho} + O(1) \quad (3.89)$$

where, in contrast to the three dimensional method, the term $\rho = \xi - \eta$ can be positive or negative. This requires particular care when deriving the series expansions for $r^n = |\mathbf{y} - \mathbf{x}|^n$ and the integration limit $\alpha_m(\epsilon)$, presented in appendix B. The explicit expansions for $\mathbf{F}_{-2}^m(\eta, \xi)$ and $\mathbf{F}_{-1}^m(\eta, \xi)$ are also derived in this appendix.

Substituting the expansions for $\mathbf{F}^m(\eta, \xi)$ given by equation 3.89 into equation 3.88, we obtain the influence as the sum of regular integrals

$$\mathbf{q} = \mathbf{I}_0 + \mathbf{I}_{-1} + \mathbf{I}_{-2} + \mathbf{a}(\mathbf{y})\tau(\mathbf{y}) \quad (3.90)$$

where

$$\mathbf{I}_0 = \int_{\zeta_1}^{\zeta_2} \left\{ \left[\mathbf{F}^1(\eta, \xi) - \frac{\mathbf{F}_{-2}^1}{\rho^2} - \frac{\mathbf{F}_{-1}^1}{\rho} \right] + \left[\mathbf{F}^2(\eta, \xi) - \frac{\mathbf{F}_{-2}^2}{\rho^2} - \frac{\mathbf{F}_{-1}^2}{\rho} \right] \right\} d\xi \quad (3.91)$$

This is a regular integral which can be evaluated by a numeric quadrature formula.

In order to integrate \mathbf{I}_{-1} and \mathbf{I}_{-2} we introduce Taylor series expansions for $\alpha_m(\epsilon)$ which mirror those of the three dimensional integration limit expansion $\alpha(\epsilon, \theta)$

$$\alpha_m(\epsilon) = \epsilon\beta_m(\eta) + \epsilon^2\gamma_m(\eta) + O(\epsilon^3) \quad (3.92)$$

As shown in appendix B, the expressions for $\gamma_m(\eta)$ contain a term $\rho/|\rho|$, which is another example of the translation of θ into the two dimensional case.

We are now able to integrate \mathbf{I}_{-1} analytically with respect to ξ (as $d\rho = d\xi$, this is equivalent to the three dimensional integration with respect to ρ , but the

integration variable here is chosen as ξ in order to simplify the limits).

$$\begin{aligned}
\mathbf{I}_{-1} &= \lim_{\epsilon \rightarrow 0} \left\{ \int_{\zeta_1}^{\eta - \alpha_1(\epsilon)} \frac{\mathbf{F}_{-1}^1}{\xi - \eta} d\xi + \int_{\eta + \alpha_2(\epsilon)}^{\zeta_2} \frac{\mathbf{F}_{-1}^2}{\xi - \eta} d\xi \right\} \\
&= \lim_{\epsilon \rightarrow 0} \left\{ \mathbf{F}_{-1}^1 \ln \frac{-\alpha_1(\epsilon)}{\zeta_1 - \eta} + \mathbf{F}_{-1}^2 \ln \frac{\zeta_2 - \eta}{\alpha_2(\epsilon)} \right\} \\
&= \lim_{\epsilon \rightarrow 0} \left\{ \mathbf{F}_{-1}^1 [\ln(\epsilon \beta_1) - \ln(\eta - \zeta_1)] + \mathbf{F}_{-1}^2 [\ln(\zeta_2 - \eta) - \ln(\epsilon \beta_2)] \right\} \\
&= \mathbf{F}_{-1}^1 \ln \frac{\beta_1}{\eta - \zeta_1} + \mathbf{F}_{-1}^2 \ln \frac{\zeta_2 - \eta}{\beta_2} + \lim_{\epsilon \rightarrow 0} \left\{ (\mathbf{F}_{-1}^1 - \mathbf{F}_{-1}^2) \ln \epsilon \right\} \tag{3.93}
\end{aligned}$$

It is seen in appendix B that the expansions of \mathbf{F}_{-1}^1 and \mathbf{F}_{-1}^2 are equal when the geometry is smooth at \mathbf{y} such that the remaining limit term will disappear to leave

$$\mathbf{I}_{-1} = \mathbf{F}_{-1}^1 \ln \frac{\beta_1}{\eta - \zeta_1} + \mathbf{F}_{-1}^2 \ln \frac{\zeta_2 - \eta}{\beta_2} \tag{3.94}$$

In order to evaluate \mathbf{I}_{-2} we require a series expansion for $\alpha_m^{-1}(\epsilon)$. The derivation for the three dimensional case is given by equation 3.82, and the same treatment in this case yields

$$\frac{1}{\alpha_m(\epsilon)} = \frac{1}{\epsilon \beta_m} \left[1 - \frac{\epsilon \gamma_m}{\beta_m} \right] + O(\epsilon) \tag{3.95}$$

This enables us to integrate \mathbf{I}_{-2} analytically with respect to ξ

$$\begin{aligned}
\mathbf{I}_{-2} &= \lim_{\epsilon \rightarrow 0} \left\{ \int_{\zeta_1}^{\eta - \alpha_1(\epsilon)} \frac{\mathbf{F}_{-2}^1}{(\xi - \eta)^2} d\xi + \int_{\eta + \alpha_2(\epsilon)}^{\zeta_2} \frac{\mathbf{F}_{-2}^2}{(\xi - \eta)^2} d\xi + \tau(\mathbf{y}) \frac{\mathbf{b}(\mathbf{y})}{\epsilon} \right\} \\
&= \lim_{\epsilon \rightarrow 0} \left\{ \mathbf{F}_{-2}^1 \left[\frac{1}{\alpha_1(\epsilon)} + \frac{1}{\zeta_1 - \eta} \right] + \mathbf{F}_{-2}^2 \left[\frac{1}{\alpha_2(\epsilon)} - \frac{1}{\zeta_2 - \eta} \right] + \tau(\mathbf{y}) \frac{\mathbf{b}(\mathbf{y})}{\epsilon} \right\} \\
&= \lim_{\epsilon \rightarrow 0} \left\{ \frac{\mathbf{F}_{-2}^1}{\epsilon \beta_1} \left[1 - \frac{\epsilon \gamma_1}{\beta_1} \right] + \frac{\mathbf{F}_{-2}^2}{\epsilon \beta_2} \left[1 - \frac{\epsilon \gamma_2}{\beta_2} \right] + \tau(\mathbf{y}) \frac{\mathbf{b}(\mathbf{y})}{\epsilon} \right\} + \left[\frac{\mathbf{F}_{-2}^1}{\zeta_1 - \eta} - \frac{\mathbf{F}_{-2}^2}{\zeta_2 - \eta} \right] \\
&= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left\{ \frac{\mathbf{F}_{-2}^1}{\beta_1} + \frac{\mathbf{F}_{-2}^2}{\beta_2} + \tau(\mathbf{y}) \mathbf{b}(\mathbf{y}) \right\} - \frac{\mathbf{F}_{-2}^1 \gamma_1}{\beta_1^2} - \frac{\mathbf{F}_{-2}^2 \gamma_2}{\beta_2^2} + \frac{\mathbf{F}_{-2}^1 \gamma_1}{\zeta_1 - \eta} - \frac{\mathbf{F}_{-2}^2 \gamma_2}{\zeta_2 - \eta} \tag{3.96}
\end{aligned}$$

As $\epsilon \rightarrow 0$ the limit terms in \mathbf{F}_{-2}^1 and \mathbf{F}_{-2}^2 must cancel (Guiggiani *et al.*, 1990) with the free-term $\tau(\mathbf{y})\mathbf{b}(\mathbf{y})$ to produce a finite integral, leaving

$$\mathbf{I}_{-2} = \mathbf{F}_{-2}^1 \left[\frac{1}{\zeta_1 - \eta} - \frac{\gamma_1}{\beta_1^2} \right] - \mathbf{F}_{-2}^2 \left[\frac{1}{\zeta_2 - \eta} - \frac{\gamma_2}{\beta_2^2} \right] \tag{3.97}$$

Terms can now be collected to produce the final expression

$$\begin{aligned} \mathbf{q} = \int_{\zeta_1}^{\zeta_2} \left\{ \left[\mathbf{F}^1(\eta, \xi) - \frac{\mathbf{F}_{-2}^1}{\rho^2} - \frac{\mathbf{F}_{-1}^1}{\rho} \right] + \left[\mathbf{F}^2(\eta, \xi) - \frac{\mathbf{F}_{-2}^2}{\rho^2} - \frac{\mathbf{F}_{-1}^2}{\rho} \right] \right\} d\xi \\ + \mathbf{F}_{-1}^1 \ln \frac{\beta_1}{\eta - \zeta_1} + \mathbf{F}_{-1}^2 \ln \frac{\zeta_2 - \eta}{\beta_2} \\ + \mathbf{F}_{-2}^1 \left[\frac{1}{\zeta_1 - \eta} - \frac{\gamma_1}{\beta_1^2} \right] - \mathbf{F}_{-2}^2 \left[\frac{1}{\zeta_2 - \eta} - \frac{\gamma_2}{\beta_2^2} \right] + \mathbf{a}(\mathbf{y})\tau(\mathbf{y}) \quad (3.98) \end{aligned}$$

where all the terms are regular and can be evaluated numerically.

3.2.5 Less than hypersingular integrals

The method described shows how the hypersingular surface integral of $\tau(\mathbf{x})\mathbf{T}(\mathbf{y}, \mathbf{x})$ can be transformed into the sum of two line integrals $\mathbf{F}_{-2}(\theta)$ and $\mathbf{F}_{-1}(\theta)$, and a regular surface integral $\mathbf{F}(\rho, \theta)$ which is evaluated over a polar coordinate frame. The theory allows integration of any order of singularity, with successive higher orders requiring more terms in the various power series expansions, and ultimately more \mathbf{F}_{-n} terms. However, as the greatest order of singularity we have is of the order $\frac{1}{r^3}$ (for a doublet velocity influence), we have restricted our analysis to cover only this and lower orders.

As the other kernel functions are of lower orders, as shown in table 3.1, we should expect the analysis required for them to be simpler. Indeed this is the case, and although the same degree of power series expansions was used for all kernel functions in appendix A, it is seen that term $\mathbf{F}_{-2}(\theta)$ is zero for the strongly singular source velocity kernels, whereas both $\mathbf{F}_{-2}(\theta)$ and $\mathbf{F}_{-1}(\theta)$ are zero for weakly singular kernel functions. In fact, simpler power series expansions could have been used for these cases, as no Jacobian or singularity distribution derivatives appear in their final expressions, and only the lowest order terms for the expansion of $\mathbf{r}(\rho, \theta)$ are present in the strongly singular kernels.

For the weakly singular potential influences, the single non-zero terms $F(\rho, \theta)$ represent exactly the original integrands. However, it is important to realise that this does not imply the integration can be carried out directly (such an integration would not be convergent with increasing numbers of integration points). The transformation of the integral from a rectangular coordinate frame (ξ_1, ξ_2) to

a polar coordinate frame (ρ, θ) is significant, and sufficient to deal with the singularity.

3.2.6 Free-terms

The method of integration described in sections 3.2.3 and 3.2.4 splits the surface or line integral into two regions: one over a small spherical or circular region e_ϵ around the singular point, and one over the remainder $S - e_\epsilon$. This is shown by equation 3.68 in the three dimensional case. We refer to the influence due to the region s_ϵ as

$$\mathbf{q}_{s_\epsilon} = \lim_{\epsilon \rightarrow 0} \left\{ \int_{s_\epsilon} \tau(\mathbf{x}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) \right\} \quad (3.99)$$

An equivalent expression results from the two dimensional analysis, except that here the integral is over the linear region $d\Gamma$.

In the neighbourhood of the small sphere s_ϵ , it is possible to write the singularity distribution as a Taylor series expansion about the singular point \mathbf{y} as follows

$$\tau(\mathbf{x}) = \tau(\mathbf{y}) + (\mathbf{x} - \mathbf{y}) \cdot \nabla \tau(\mathbf{y}) + O(|\mathbf{x} - \mathbf{y}|^2) \quad (3.100)$$

Substituting this expansion into equation 3.99 whilst observing that in the limit over s_ϵ terms $O(|\mathbf{x} - \mathbf{y}|^2)$ disappear, we obtain

$$\mathbf{q}_{s_\epsilon} = \lim_{\epsilon \rightarrow 0} \left\{ \int_{s_\epsilon} \tau(\mathbf{y}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) + \int_{s_\epsilon} (\mathbf{x} - \mathbf{y}) \cdot \nabla \tau(\mathbf{y}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) \right\} \quad (3.101)$$

As the singularity distribution is only evaluated at \mathbf{y} , the terms $\tau(\mathbf{y})$ and $\nabla \tau(\mathbf{y})$ can be brought outside the integrals, and the integration on the remaining kernel functions performed (Guiggiani *et al.*, 1990) to produce the following expression

$$\mathbf{q}_{s_\epsilon} = \tau(\mathbf{y}) \left[\frac{\mathbf{b}(\mathbf{y})}{\epsilon} + \mathbf{a}(\mathbf{y}) \right] + \nabla \tau(\mathbf{y}) \cdot \mathbf{c}(\mathbf{y}) + O(\epsilon) \quad (3.102)$$

For kernel functions with less than hypersingular integrals, the equivalent expression is simply

$$\mathbf{q}_{s_\epsilon} = \tau(\mathbf{y}) \mathbf{c}(\mathbf{y}) + O(\epsilon) \quad (3.103)$$

It was shown in sections 3.2.3 and 3.2.4 that the singular term $\mathbf{b}(\mathbf{y})/\epsilon$ must cancel with a singular part of the integral $\mathbf{q}_{S-\epsilon}$ in order to produce a finite influence function. This leaves the so-called ‘free-terms’ $\mathbf{a}(\mathbf{y})$ (for hypersingular integrals) and $\mathbf{c}(\mathbf{y})$. Methods for deriving these coefficients can be found in Guiggiani *et al.* (1990), and the results are summarised here for the kernel functions needed within the CSM.

Velocity influence of a doublet distribution

In the original paper that presented the method of solving the hypersingular kernel integral equation (Guiggiani *et al.*, 1990), the free-term $\mathbf{a}(\mathbf{y})$ was overlooked. This was corrected in a later paper (Guiggiani, 1995), but it was shown that the term is only non-zero when the singular point \mathbf{y} lies at a point with a geometric tangent or curvature discontinuity. Since the patches used in this CSM have curvature continuity throughout, the term can be neglected here.

At a smooth boundary point \mathbf{y} , the other free-term coefficient is simply

$$\mathbf{c}(\mathbf{y}) = \begin{pmatrix} 0.5 \\ 0.5 \\ 0.5 \end{pmatrix} \quad \text{in three dimensions, or} \quad \mathbf{c}(\mathbf{y}) = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \quad \text{in 2D} \quad (3.104)$$

This produces the familiar result that the tangential velocity jump across a surface is equal to the gradient of the doublet strength on that surface. As we are concerned generally with velocities normal to the surface (or mid-surface velocities for the case of sampling wake velocities) we disregard this term in the main analysis, but add it back on in post-processing when evaluating surface velocities (see section 3.5.1).

Velocity influence of a source distribution

Equation 3.9 shows that there is a discontinuity in the normal velocity across a surface containing a source distribution, and as such, care is required when determining the singular velocity influence at a point on the surface. In particular, the limiting behaviour cannot be determined correctly by placing the singular

point on the surface and treating the integral in two parts, $\mathbf{q}_{S-e_\epsilon}$ and \mathbf{q}_{s_ϵ} . This is because the limiting influence $\mathbf{q}_{S-e_\epsilon}$ for a singular point close to the surface is not zero, whereas a direct integration of the kernel function over this region for a point on the surface would produce zero. The correct influence is determined by the limiting velocity as the point approaches the surface S (treated as a whole without separating e_ϵ) from the relevant side.

In fact with the CSM (and most common panel methods), sources are distributed on a surface specifically to set the normal velocity component to zero, so there is no need to evaluate this component numerically.

Potential influence of a doublet distribution

Like the velocity jump across a source distribution, there is a potential jump across a doublet distribution, as shown by equation 3.10. However, unlike the case of the source velocity, the region immediately outside e_ϵ does not contribute to the doublet potential influence. This is because the term $\mathbf{n}(x) \cdot \mathbf{r}(\mathbf{y}, \mathbf{x})$ in the kernel function tends to zero just outside e_ϵ , assuming the geometry is smooth at \mathbf{y} .

The free-term coefficient $c(\mathbf{y})$ can therefore be calculated by simply integrating the kernel function over the surface of s_ϵ . This is most readily accomplished by a transformation to polar coordinates centred on \mathbf{y} , and on a smooth boundary the result is given by

$$c = 0.5 \quad (\text{in 2D and 3D}) \quad (3.105)$$

Potential influence of a source distribution

There are no pitfalls relating to the source potential influence free-term, and a similar integration on the weakly singular kernel shows that in this case

$$c = 0 \quad (\text{in 2D and 3D}) \quad (3.106)$$

Table 3.1: Summary of kernel functions and methods of integration

Influence type	Kernel function	Singularity order	Singular terms	Free terms
3D doublet potential	$-\frac{1}{4\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{r^3(\mathbf{y}, \mathbf{x})} \right\}$	Weak		$0.5\mu(\mathbf{y})$
3D source potential	$-\frac{1}{4\pi} \left\{ \frac{1}{r(\mathbf{y}, \mathbf{x})} \right\}$	Weak		
3D doublet velocity	$-\frac{1}{4\pi} \left\{ \frac{\mathbf{n}(\mathbf{x})}{r^3(\mathbf{x})} - \frac{3[\mathbf{r}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})]\mathbf{r}(\mathbf{x})}{r^5(\mathbf{x})} \right\}$	Hypersingular	$\mathbf{F}_{-1}(\theta), \mathbf{F}_{-2}(\theta)$	$\mathbf{a}(\mathbf{y})\mu(\mathbf{y}) = 0^*$
3D source velocity	$\frac{1}{4\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x})}{r^3(\mathbf{y}, \mathbf{x})} \right\}$	Strong	$\mathbf{F}_{-1}(\theta)$	
2D doublet potential	$-\frac{1}{2\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{r^2(\mathbf{y}, \mathbf{x})} \right\}$	Weak		$0.5\mu(\mathbf{y})$
2D source potential	$-\frac{1}{2\pi} \ln r(\mathbf{y}, \mathbf{x})$	Weak		
2D doublet velocity	$-\frac{1}{2\pi} \left\{ \frac{\mathbf{n}(\mathbf{x})}{r^2(\mathbf{x})} - \frac{2[\mathbf{r}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})]\mathbf{r}(\mathbf{x})}{r^4(\mathbf{x})} \right\}$	Hypersingular	$\mathbf{F}_{-1}, \mathbf{F}_{-2}$	$\mathbf{a}(\mathbf{y})\mu(\mathbf{y}) = 0^*$
2D source velocity	$\frac{1}{2\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x})}{r^2(\mathbf{y}, \mathbf{x})} \right\}$	Strong	\mathbf{F}_{-1}	

* $\mathbf{a}(\mathbf{y}) = 0$ when the geometry has continuous tangent and curvature at \mathbf{y}

3.2.7 Unknown singularity distributions

The method presented in sections 3.2.3 and 3.2.4 describes how the potential or velocity influence of a known singularity distribution can be found. This is useful for evaluating on-surface wake velocities (see chapter 5), but not for solving the potential flow problem. However, it is straightforward to apply the scheme to an unknown distribution, and in fact Guiggiani *et al.* (1990) present their method in terms of element shape functions instead of singularity strengths, since it was designed for use in higher-order boundary element methods. The CSM though does employ discrete elements as such, and the concept of a shape function is not directly applicable. It will be seen in section 3.3.5 that rather than evaluating the total integral in one go, the constituent parts of $F_{-2}(\theta)$, $F_{-1}(\theta)$ and $F(\rho, \theta)$ are stored for later construction into an influence matrix. In particular, the line integrals $F_{-2}(\theta)$ and $F_{-1}(\theta)$ are separated into three parts that are multiplied by $\tau(\xi_1, \xi_2)$, $\partial\tau/\partial\xi_1$ and $\partial\tau/\partial\xi_2$ respectively.

3.2.8 Alternative integration of the doublet velocity kernel

This section presents an alternative method for evaluating the hypersingular doublet velocity integral. This scheme was developed for use within the CSM before the author became aware of work on hypersingular kernels by Guiggiani, but although slightly more computationally expensive than the standard method already covered, can be used on surfaces where the gradient of the ‘Jacobian normal’ function $\mathbf{J}(\mathbf{x}) = \mathbf{n}(\mathbf{x})J(\mathbf{x})$ is unavailable, unreliable, or computationally expensive to determine.

The method is based on splitting the surface integral into two parts, one relating to the value of the doublet strength at the singular point \mathbf{y} , and one relating to the remaining higher order terms:

$$\begin{aligned} \mathbf{q}(\mathbf{y}) &= \int_S \mu(\mathbf{x}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) \\ &= \int_S \mu(\mathbf{y}) \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) + \int_S [\mu(\mathbf{x}) - \mu(\mathbf{y})] \mathbf{T}(\mathbf{y}, \mathbf{x}) dS(\mathbf{x}) \end{aligned} \quad (3.107)$$

Integration of the constant part

The first term in equation 3.107 represents the velocity induced by a doublet distribution of constant strength on a general curved surface. This is evaluated by employing a special case of the transformation described by Hess (1979) (see section 2.1). This transformation allows the constant doublet influence to be represented solely by curved vortex lines around the perimeter of the surface.

This transformation is commonly used on straight-edged panels in vortex lattice and constant panel codes, where constant doublet influences are usually calculated by summing the influence of vortex lines around the panels' perimeters. The applicability of this result to a general curved surface can be appreciated by first considering the curved surface to be approximated by an array of contiguous quadrilateral panels. Along the lines between flat panels, the vortex influence disappears as the two common edges represent equal and opposite vortex lines. This leaves only the vortex lines on the edge of the surface, which are all equal in strength to the constant doublet distribution. If the density of the panels is then increased, in the limit the vortex lines on the edge of the surface become smooth curves.

Unlike the hypersingular doublet velocity kernels, the influence due to the curved vortex lines around the edge of the surface can be evaluated directly (except on the lines themselves, where the velocity is undefined regardless of how it is calculated). In practice, the velocity influence is evaluated using a one dimensional numeric quadrature around the edge of each surface (or each 'super-panel' as defined in section 3.3.3).

Integration of the higher-order part

Integration of the second part of equation 3.107 is carried out using the standard scheme for singular integrals laid out in section 3.2.3. However, with the doublet distribution now $\mu(\mathbf{x}) - \mu(\mathbf{y})$ instead of $\mu(\mathbf{y})$, the Taylor expansion given by equation A.9a becomes

$$\mu(\boldsymbol{\xi}) = \mu_0 + \rho\mu_1(\theta) + O(\rho^2) \quad (3.108)$$

where now $\mu_0 = \mu(\boldsymbol{\eta}) = 0$. The expansion of $\mathbf{F}(\rho, \theta)$ as derived in appendix A is then simplified to

$$\begin{aligned}\mathbf{F}(\rho, \theta) &= \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} + \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \\ &= \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1)\end{aligned}\tag{3.109}$$

since

$$\mathbf{F}_{-2}(\theta) = \frac{1}{4\pi} \left\{ \frac{\mathbf{J}_0}{A^3} \right\} \mu_0 = 0\tag{3.110}$$

thus reducing the kernel function from hypersingular to strongly singular. The expression for $\mathbf{F}_{-1}(\theta)$ is also substantially simplified to

$$\begin{aligned}\mathbf{F}_{-1}(\theta) &= \frac{1}{4\pi} \left\{ \left[\frac{\mathbf{J}_1}{A^3} - 3 \frac{(\mathbf{A} \cdot \mathbf{B})\mathbf{J}_0 + (\mathbf{B} \cdot \mathbf{J}_0)\mathbf{A} + (\mathbf{A} \cdot \mathbf{J}_1)\mathbf{A}}{A^5} \right] \mu_0 + \frac{\mathbf{J}_0}{A^3} \mu_1 \right\} \\ &= \frac{1}{4\pi} \left\{ \frac{\mathbf{J}_0}{A^3} \mu_1 \right\}\end{aligned}\tag{3.111}$$

thus eliminating surface second derivatives and Jacobian derivative terms from the calculation.

3.3 Singularity distributions represented as B-Spline surfaces

3.3.1 Bi-cubic B-spline basis functions

Until this point, we have considered the singularity distributions to be generic continuous surfaces, and have had no need to prescribe their form. Indeed, the general method can be applied to any such surface, and as will be seen in section 3.3.2, certain physical situations may benefit from the use of specifically tailored surface definitions. However, for a practical implementation we must define a form of singularity surface whose shape is a function of discrete parameters \mathbf{B}^μ or \mathbf{B}^σ which themselves comprise the set of unknowns of the potential flow equations. Initially we will choose a B-spline tensor product surface with basis functions $U_{i,r}$ defined by the well-established Cox-de-Boor recursion formulae (Rogers & Adams, 1990). The i -th basis function of order r

(degree $r - 1$) is evaluated by

$$U_{i,1}(\xi) = \begin{cases} 1 & \text{if } Z_i \leq \xi < Z_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (3.112a)$$

and

$$U_{i,r}(\xi) = \frac{[\xi - Z_i]U_{i,r-1}(\xi)}{Z_{i+r-1} - Z_i} + \frac{[Z_{i+r} - \xi]U_{i+1,r-1}(\xi)}{Z_{i+r} - Z_{i+1}} \quad (3.112b)$$

where Z_i are the elements of a knot vector with the property $Z_i \leq Z_{i+1}$.

Employing B-spline curves across both ξ_1 and ξ_2 dimensions, the value of the singularity distribution $\tau(\xi, \mathbf{B}^T)$, at a point with intrinsic coordinates $\xi = (\xi_1, \xi_2)$ is the sum of a set of weighted basis function products

$$\tau(\xi, \mathbf{B}^T) = \sum_{i=1}^{N_i} \sum_{j=1}^{N_j} B_{ij}^T U_i^{T1}(\xi_1) U_j^{T2}(\xi_2) \quad (3.113)$$

where $U^{T1}(\xi_1)$ and $U^{T2}(\xi_2)$ are two independent sets containing N_i and N_j basis functions respectively.

Any degree of B-spline basis functions could be used to describe the singularity distribution surface, and the degree need not be the same for $U^{T1}(\xi_1)$ and $U^{T2}(\xi_2)$. In fact, if a first order (zero-degree) were chosen for both axes, with a corresponding piecewise-linear geometric surface, the scheme would reduce to a standard constant panel or vortex lattice method. In this work, we will generally employ bi-cubic basis functions with knot vectors of the form

$$\mathbf{Z}^{T1} = [0, 0, 0, 0, Z_4^{T1}, Z_5^{T1} \dots 1, 1, 1, 1] \quad (3.114)$$

$$\mathbf{Z}^{T2} = [0, 0, 0, 0, Z_4^{T2}, Z_5^{T2} \dots 1, 1, 1, 1] \quad (3.115)$$

These vectors, with four equal knots at each end, produce basis functions with continuous curvature and independent control over value, gradient and curvature at both ends of the resulting spline. A set of nine basis functions defined by one such knot vector of this form is shown in figure 3.8. A set of N_i basis functions of the above form requires a knot vector with $N_i + 4$ elements producing an intrinsic parameter space divided into $N_p = N_i - 3$ inter-knot regions. Within




Figure 3.8: The set of basis functions for a cubic B-spline with nine control points and an evenly spaced knot vector with repeated end knots.

the regions $0 < \xi_1 < 1$ and $0 < \xi_2 < 1$, the distribution of knots is a tuning parameter that will be discussed in chapter 4.

3.3.2 Modified basis functions

As B-spline curves are piecewise polynomials of finite degree, their gradients are also polynomial functions and as such are always finite (for finite control points). In general this presents no problems to modelling singularity distributions over smooth surfaces; however, as was shown in section 3.1.5, leading edges of thin bodies have a doublet gradient with distance s from the edge of the form

$$\frac{\partial \mu}{\partial s} = \frac{C}{2\sqrt{s}} \quad (3.116)$$

which is singular at the leading edge $s = 0$. As discussed in section 3.1.7, this characteristic also exists at other edges that do not shed wakes.

Despite the finite-gradient limitation of a B-spline curve, the doublet shape can still be approximated reasonably well except very close to the edge, particularly if knots are concentrated towards that edge. However, as we are free to choose any set of smooth basis functions to suit the physical problem, we can sometimes improve on the standard B-spline set by modifying the functions $U_1(\xi)$

closest to non-wake shedding edges by a square-root basis function as follows

$$U_1'(\xi) = \sqrt{\xi}U_1(\xi) \quad (3.117)$$

Diagrams showing this and other forms of alternate basis functions are shown in section 4.1.5, with a discussion of the results obtained by using them.

3.3.3 Panels and super-panels

When two sets of basis functions, either B-spline or otherwise, are combined according to equation 3.113 to form a surface, the grid formed by the knot lines divides the patch into regions we refer to as panels. Each panel P_{pq} covers a rectangular region of the intrinsic parameter space and is bounded by

$$Z_{p+2}^{\mu_1} < \xi_1 < Z_{p+3}^{\mu_1}, \quad Z_{q+2}^{\mu_2} < \xi_2 < Z_{q+3}^{\mu_2} \quad (3.118)$$

As can be seen from figure 3.8, a single B-spline basis function influences between one and four segments of a spline, with four segments being the standard region of influence away from the edges. Similarly for a bi-cubic spline surface, the product of two basis functions influences a rectangular grid of panels ranging in number from 1×1 to 4×4 . We refer to the set of panels influenced by given basis-function product $U_i^{\mu_1}(\xi_1)U_j^{\mu_2}(\xi_2)$ as *super-panel* S_{ij} . Note that a square-root basis function influences an entire row or column of panels within a patch, and a super-panel formed by the product of two square-root basis functions will in fact cover the entire patch.

Each super-panel's basis function product $U_i^{\mu_1}(\xi_1)U_j^{\mu_2}(\xi_2)$ has a weight B_{ij}^μ associated with it, and it is this array of weights that form the unknowns of the potential flow problem. Thus the task can be stated as finding the super-panel weights, which when multiplied by the pre-defined super-panel basis functions produce a singularity surface that satisfies the boundary conditions. For the collocation formulation this is expressed as

$$A_{hk}B_k = C_h \quad (3.119)$$




Figure 3.9: Panels, super-panels and basis-functions for a singularity distribution

where element A_{hk} represents the influence of super-panel S_k with unit weight on collocation point y_h . The vector C is determined by application of the boundary conditions. In this equation, for ease of matrix manipulation, the two dimensional array of super-panels has been reshaped into a vector

$$S_k \equiv S_{ij} \tag{3.120}$$

where if $i = \{1 \dots N_i\}$ and $j = \{1 \dots N_j\}$ then $k = \{1 \dots N_i N_j\}$.

A surface with $N_p \times N_q$ panels (and B-spline basis functions) will comprise $N_i \times N_j = (N_p + 3) \times (N_q + 3)$ super-panels, and will require an equal number of equations to provide a solution. Note however that whilst most of these equations will represent the influence of a singularity distribution at a collocation point, some may represent fixed boundary conditions as described in section 3.3.8.

3.3.4 Off-surface influence evaluation

In this section we describe how a term of the influence matrix A_{hk} is constructed for the general off-surface case. The term expresses the potential or velocity induced at a collocation point y_h by a basis-function product of unit. As detailed in the previous section, the basis function product is only non-zero within a

defined super-panel S_k , and here we consider the case where the collocation point does not lie on the patch containing the super-panel.

The panels that make up the super-panel provide a useful division of the surface from which we can base our numerical integration of the off-surface influence coefficients. Within each panel P_{pq} the doublet distribution is described by a distinct polynomial (typically bi-cubic) of the intrinsic coordinates, and depending on how the surface geometry is represented (eg. with bi-cubic splines using the same knot grid), this may also be known to vary in a relatively simple way over the extent of a single panel. There are also efficiency gains from dividing super-panels into panels, as described in section 3.3.7. We can therefore write the super-panel influence as the sum of influences due to each panel P_r within the super-panel S_k

$$A_{hk} = \sum_i \sum_j q_{hij} \quad (3.121)$$

where the influence of panel P_{pq} is

$$q_{hijr} = \int_S \mu(\boldsymbol{\xi}, B_{ij}^\mu) \mathbf{T}(\mathbf{y}_h, \mathbf{x}) dS(\mathbf{x}) \quad (3.122)$$

$$= \int \int U_i^{\mu_1}(\xi_1) U_j^{\mu_2}(\xi_2) \mathbf{T}(\mathbf{y}_h, \boldsymbol{\xi}) J(\boldsymbol{\xi}) d\xi_1 d\xi_2 \quad (3.123)$$

the integration being over the limits of panel P_r defined by equation 3.118. To reduce unnecessary suffixes, panels here are referred to by a one dimensional vector index r instead of their previous two dimensional array index pq , but the two are equivalent such that $P_{pq} = P_r$.

The continuous integral above is easily transformed into a discrete sum by the application of a standard cubature formula of order $(N_m \times N_n)$ to produce

$$q_{hijr} = \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} w_m w_n U_i^{\mu_1}(\xi_{1m}) U_j^{\mu_2}(\xi_{2n}) \mathbf{T}(\mathbf{y}_h, \boldsymbol{\xi}_{mn}) J(\boldsymbol{\xi}_{mn}) \quad (3.124)$$

This sum can then be computed directly, evaluating the singularity basis function product and the surface geometry quantities at each integration point over the panel. For a bi-cubic singularity distribution, a (4×4) Legendre-Gauss cubature

rule² (Engels, 1980) would typically be used, although due to higher orders of ξ in the kernel function and possibly the surface geometry, this will produce a good approximation rather than an exact solution.

Note that whilst the integration proceeds, for convenience, by dividing the surface into regions we refer to as panels, this neither implies any discretisation of the geometry, nor any of doublet distribution beyond that imposed by the type of parametric surface chosen to represent it. The accuracy to which the geometry and doublet distributions are represented is limited only by the order of the cubature formulae chosen for the final integration, which itself is limited only by computation time. This is in contrast to conventional constant and higher-order panel methods, in which the geometry is constrained to the same level of discretisation as the singularity distributions.

Variable cubature orders and panel subdivision

Conventional panel methods often employ different panel formulations for ‘near field’ and ‘far field’ panels; within the CSM, efficiency can be gained by adjusting the order of the cubature according to the anticipated significance of the panel’s influence. In estimating a panel’s significance, and hence the order of cubature it is worth ‘spending’ on the panel, a criteria should be chosen which is accurate, but quick enough to evaluate that its computational cost is much less than the cost of simply always using a high order of cubature.

The closest distance of the panel to the collocation point is a reasonable measure of this anticipated significance, although more elaborate criteria including the panel’s aspect ratio and its swept angle as viewed from the collocation point could be considered. Results presented in this thesis employ a closest distance approximation, whereby the distance between the collocation point and five points on the panel are determined, these being the four corners and the centre-point (in the sense of the intrinsic coordinates) as shown in figure 3.10. The minimum of these five distances is then compared with the

²Following the terminology of Engels (1980), ‘cubature’ refers to quadrature over more than one dimension of integration

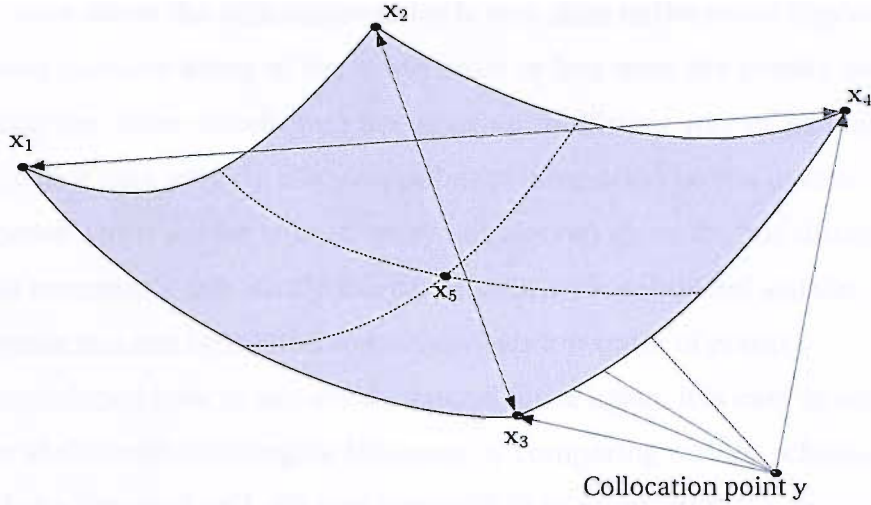


Figure 3.10: Points used to calculate the ‘nearness’ of a panel to a collocation point




Figure 3.11: Variable cubature orders on panels within a super-panel close to the collocation point. Integration points are shown as black dots, with the white circle above the super-panel plane being the collocation point.

maximum diagonal of the panel to create a normalised distance

$$s_n = \frac{\min(|y - \mathbf{x}_{1..5}|)}{\min(|\mathbf{x}_1 - \mathbf{x}_4|, |\mathbf{x}_2 - \mathbf{x}_3|)} \quad (3.125)$$

The distance s_n can then be used as the index into a pre-defined look-up table to determine an appropriate order of cubature. Figure 3.11 shows how the order of cubature might vary over a super-panel in close proximity to a collocation point.

For cases where the collocation point is very close to the panel (typically with the closest distance being of the same order or less than the panel's diagonal), increasing the order of cubature becomes an inefficient way of increasing the accuracy, as it may result in a large number of integration points in some regions of the panel which are far enough away not to need them. In this situation, it is better to recursively sub-divide the panel until each sub-panel satisfies the far-field criteria and can be treated with a relatively low order of points.

In considering how to sub-divide a panel, once again, it is easy to conceive a number of alternative strategies. However, by comparing various schemes, it was found that a practical and efficient approach is to simply divide a panel into two halves, splitting along a centreline (in the sense of its intrinsic coordinates), as shown in figure 3.12. The choice of centreline (top-to-bottom or left-to-right) is made by comparing the sum of lengths of the top and bottom sides with the sum of the lengths of the left and right sides, and then splitting the sides that produce the greater sum (for speed of computation, side length here is estimated as the straight-line distance between corners). In addition to the simplicity benefits of such a scheme, the fact that a given panel is always sub-divided in the same way means it can often be re-used, bringing efficiency gains (see section 3.3.7). This would not be possible if the method of sub-division was dependent on the relative position of the collocation point. Once a panel has been divided, the value of s_n is evaluated for each sub-panel which can then be integrated by a cubature rule or further sub-divided as necessary.

The values of s_n that form the steps in cubature order and subdivision criteria have a strong influence over both the accuracy and computational cost of the overall solution. For the sake of efficiency it is therefore important that the set of criteria is 'balanced', such that each step in the look-up table represents a common trade-off between accuracy and computation time. An approximately balanced table can be constructed by using very high accuracy settings as a baseline (high cubature orders and widespread subdivision), and then investigating the effect of changing each s_n step in terms of computation time and deviation from the baseline solution. Table 3.2 shows a typical set of cubature order and sub-division criteria that were drawn up in this way.

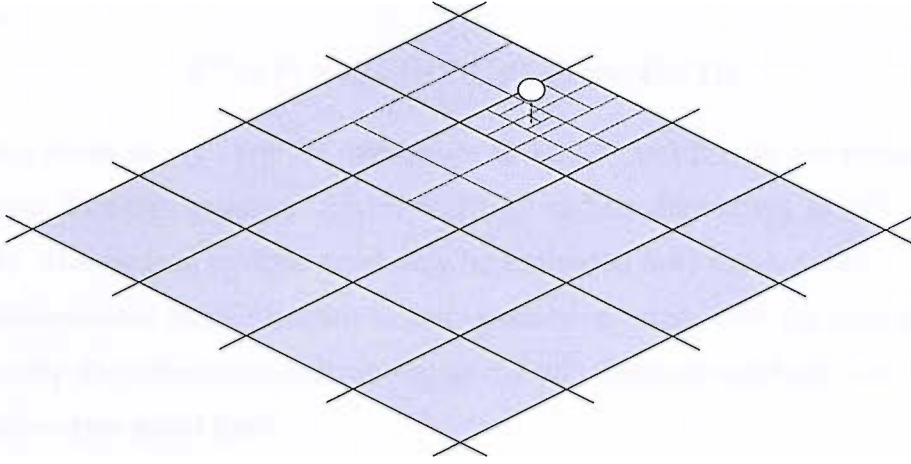


Figure 3.12: Recursive subdivision within panels. The panel closest to the collocation point (the white circle) has six levels of subdivision, forming a total of 17 smaller panels. Integration points within each panel or sub-panel are not shown.

Table 3.2: Typical sub-division and cubature order criteria

s_n	≤ 0.8	$> 0.8, \leq 1.6$	$> 1.6, \leq 3.2$	$> 3.2, \leq 6.4$	> 6.4
Grid	Sub-divide	4×4	3×3	2×2	1×1

3.3.5 On-surface influence evaluation

For the case in which the collocation point y_h lies on the influencing surface, it was seen in section 3.2 that the integrand over the surface is singular. Equation 3.86 shows that the integral can be evaluated in two parts: a surface integral of the desingularised integrand over the entire patch, and a line integral around the singular point. Both integrals are evaluated over a polar representation of the intrinsic coordinate space.

Whereas equation 3.86 provides the velocity or potential influence of a surface with known singularity distribution $\tau(\mathbf{x})$, we require the separate influences \mathbf{q}_{hk} due to each super-panel S_k

$$\mathbf{q}_{hk} = \int_0^{2\pi} \int_0^{\rho_{\bar{h}k}(\theta)} \left\{ \mathbf{F}^{hk}(\rho, \theta) - \frac{\mathbf{F}_{-2}^{hk}(\theta)}{\rho^2} - \frac{\mathbf{F}_{-1}^{hk}(\theta)}{\rho} \right\} d\rho d\theta + \int_0^{2\pi} \left\{ \mathbf{F}_{-1}^{hk}(\theta) \ln \frac{\rho_{\bar{h}k}(\theta)}{\beta_h(\theta)} - \mathbf{F}_{-2}^{hk}(\theta) \left[\frac{\gamma_h(\theta)}{\beta_h^2(\theta)} + \frac{1}{\rho_{\bar{h}k}(\theta)} \right] \right\} d\theta + \mathbf{a}(y_h)\tau(y_h) \quad (3.126)$$

where

$$\mathbf{F}^{hk}(\rho, \theta) = U_{i(k)}^{\tau_1}(\xi_1)U_{j(k)}^{\tau_2}(\xi_2)\mathbf{T}(\boldsymbol{\eta}_h, \boldsymbol{\xi})J(\boldsymbol{\xi})\rho \quad (3.127)$$

and the terms of $\tau(\boldsymbol{\eta})$ and its derivatives in $\mathbf{F}_{-2}(\theta)$ and $\mathbf{F}_{-1}(\theta)$ are replaced by the basis function product $U_{i(k)}^{\tau_1}(\eta_1)U_{j(k)}^{\tau_2}(\eta_2)$ and its derivatives in $\mathbf{F}_{-2}^{hk}(\theta)$ and $\mathbf{F}_{-1}^{hk}(\theta)$. The surface integral need only be evaluated over super-panel S_k as the two dimensional basis function is zero elsewhere. Note that the line integral takes only the value (and derivatives) of the two dimensional basis function at the collocation point itself.

It is straightforward to convert both the continuous surface and the continuous line integrals into discrete forms suitable for numeric integration. Employing a standard cubature formula of order N_m for θ and N_n for ρ , we obtain

$$\begin{aligned} \mathbf{q}_{hk} = & \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} w_m w_n(m) \left\{ \mathbf{F}^{hk}(\rho_n, \theta_m) - \frac{\mathbf{F}_{-2}^{hk}(\theta_m)}{\rho_n^2} - \frac{\mathbf{F}_{-1}^{hk}(\theta_m)}{\rho_n} \right\} \\ & + \sum_{m=1}^{N_m} w_m \left\{ \mathbf{F}_{-1}^{hk}(\theta_m) \ln \frac{\rho_{hk}(\theta_m)}{\beta_h(\theta_m)} - \mathbf{F}_{-2}^{hk}(\theta_m) \left[\frac{\gamma_h(\theta_m)}{\beta_h^2(\theta_m)} + \frac{1}{\rho_{hk}(\theta_m)} \right] \right\} + \mathbf{a}(\mathbf{y}_h)\tau(\mathbf{y}_h) \end{aligned} \quad (3.128)$$

Although this equation shows the same order being used for θ in both the one and two dimensional integrals, this is not a requirement, and a better overall efficiency/accuracy trade-off may be achieved by using different quadrature orders in each case.

Panel-by-panel integration

Direct application of equation 3.128 for all two dimensional basis functions over the patch would then produce the desired result, assuming sufficiently high orders of cubature were used. However, it is more efficient to carry out the process panel by panel, as it was in the case of the off-surface integration. Considering the influence of individual panels, we can classify each panel into one of the following four categories according to the relative position of the collocation point:

- Collocation point is outside panel

- Collocation point is on vertex of panel
- Collocation point is on edge of panel
- Collocation point is inside panel

If the collocation point is outside the panel, the integral of equation 3.67 over that panel contains no singularities and is regular. Although the collocation point lies on the patch that contains the panel, it is off the surface of the panel itself. As such, the method developed in section 3.3.4 for off-surface influence coefficients can be used, employing the same strategies of variable cubature orders and adaptive sub-division. In cases where the collocation point is very close to, but still outside the edge of a panel, the integrand will exhibit a ‘near-singular’ characteristic, and will vary rapidly over the panel. However, this variation will automatically be captured by the panel-subdivision process, and the resulting high density of integration points.

If the collocation point lies inside the panel, or on its edge or vertex, the panel does contain a singularity, and the desingularised form of equation 3.86 must be used. Two sets of integration points are required to evaluate this integral numerically: a polar array of points over the surface of the panel, and an arc of points varying in theta around the collocation point. For the case where the collocation point lies inside the panel, the integration over θ will be from zero to 2π ; for edge or vertex cases the integration will cover a θ span of π and $\frac{\pi}{2}$ respectively. As collocation points always lie in the interior of a patch and never on a patch vertex or edge, the desingularisation will ultimately always include the full range $0 < \theta < 2\pi$; but in the case of a point lying on the vertex or edge of a panel, this range of θ will be split over two to four panels. Figure 3.13 demonstrates typical arrangements of cubature points for the cases of a collocation point on a panel vertex and in the centre of a panel.

Alternative strategies

Whilst the desingularised integration must be used in the immediate vicinity of a collocation point, and the regular integration scheme may be used elsewhere, the




Figure 3.13: Schematic of integration point grids for cases where the collocation point lies on the super-panel surface: (Top) with the collocation point (white) on panel vertices; (Bottom) with the collocation point at a panel centre.

choice of where to change from one to the other has been somewhat arbitrary. Purely for convenience, we have used the desingularised scheme up to the nearest panel boundary, and the regular scheme beyond. Depending on the relative position of the collocation point within a panel, this may result in a very asymmetric division of integration schemes. This in itself should not be a cause for concern, and with sufficiently dense cubature grids accurate influences will be determined. However, where the desingularised region extends a long way from the collocation point, particularly high-order cubature grids would be required to sufficiently model the outfield; and where the extent of the desingularised region is very short, an extremely high degree of subdivision would result in the

neighbouring regular panels. Both cases are undesirable as they result in a high number of integration points with a correspondingly high computational cost.

A more sophisticated scheme could be envisaged that, by subdivision, allowed changes from the desingularised integral to the regular integral to occur within, as well as between panels. The switch could then occur at a distance from the collocation point that provided an efficient balance between cubature density in the desingularised region and degree of subdivision in the regular region.

3.3.6 Influence coefficient validation

Validation of the velocity and potential self-influence coefficients for flat and curved panels can be achieved by conducting a series of convergence studies. For alternative singularity distributions and CSM panel geometries, the influence at a point on the surface of a single CSM panel can be compared to that produced by a Constant Panel representation of the surface comprising multiple panels. As the density of CPM panels increases, convergence of the CPM coefficient towards that of the curved panel CSM solution will verify the solution.

For off-surface influence coefficients, a sample point can be moved progressively closer to the surface of a curved CSM panel. It can then be verified that the influence measured at this point converges to the on-surface influence as the normal distance tends to zero – excepting cases where there is naturally a step change in influence across a panel boundary, such as the tangential velocity change across a doublet panel.

3.3.7 Re-use of panel integration points

Sections 3.3.4 and 3.3.5 have detailed how numeric integration is carried out on a panel-by-panel basis in order to determine the influence of each super-panel in turn. However, by recognising that every panel is a component of several different super-panels, significant efficiency savings can be made.

For a doublet distribution modelled by a bi-cubic spline surface, the influence of a panel P_{pqr} on a collocation point y_h will be evaluated sixteen times, once for each super-panel S_{ij} that includes panel P_r . Although these sixteen influences

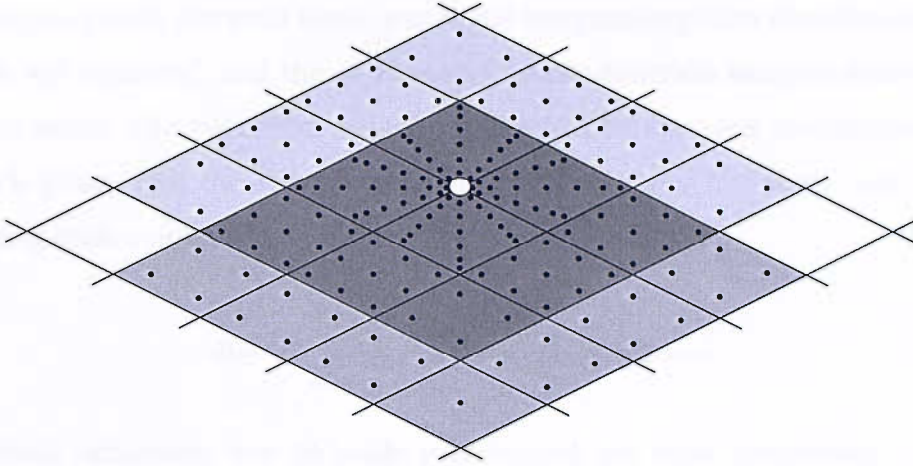


Figure 3.14: Two neighbouring super-panels sharing many of the same integration points (black) for influences relating to the same collocation point (white).

will all be different, as they will have been integrated with different basis function products for each super-panel, they will share the same cubature grids and degree of subdivision, as shown in figure 3.14. Moreover, the kernel function of the integrand describing the unit influence of a point singularity at each integration point is identical for all super-panels.

Recalling equation 3.124

$$q_{hijr} = \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} w_m w_n U_i^{\mu_1}(\xi_{1m}) U_j^{\mu_2}(\xi_{2n}) \mathbf{T}(\xi_{mn}, y_h) J(\xi_{mn}) \quad (3.129)$$

it is seen that the basis-functions for super-panel S_{ij} can be separated from the remainder of the product. It is therefore sensible to use a two-stage approach when constructing the overall influence matrix between super-panels and collocation points. In the first stage, integration points are constructed across each panel, and un-weighted influences are determined for each integration point and a given collocation point.

$$q_{hmnr} = w_m w_n \mathbf{T}(\xi_{mn}, y_h) J(\xi_{mn}) \quad (3.130)$$

Once these sets of integration points and influences have been determined for all panels and the current collocation point, the second stage cycles through

each super-panel. For each super-panel, the integration points from the relevant panels are collected, and the super-panel's basis-function weights determined at each point. The influences are then multiplied by the basis-function weights at each point, and the results summed, adding in any free-term coefficients resulting from a singular integral.

$$q_{hijr} = \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} U_i^{\mu_1}(\xi_{1m}) U_j^{\mu_2}(\xi_{2n}) q_{hmnr} \quad (3.131)$$

The total influences due to each super-panel are thus computed, without evaluating the influence of an integration point on a collocation point more than once. Equation 3.130 represents the more expensive part of the calculation, as this involves evaluation of the geometric surface, its derivatives and then the kernel function, and so considerable time savings are made. The approach is summarised by the algorithms shown in figures 3.18 and 3.19.

3.3.8 Problem configuration for bodies and wakes

A patch with a singularity distribution comprising $N_p \times N_q$ bi-cubic B-spline panels will comprise $N_i \times N_j = (N_p + 3) \times (N_q + 3)$ super-panels and an equal number of unknowns, as seen in section 3.3.1. The most efficient solution to finding these unknowns will involve an exactly determined or square system, with an equal number of equations. Such a system is obtained by associating an equation with each knot of a patch's intrinsic coordinate space, and at the mid-knot points of the panels around the perimeter of the patch. This arrangement is shown in figure 3.15.

In addition to being efficient because it produces an exactly-determined system of equations, this scheme is effective in concentrating constraints at the edges of the patch where there is a greater concentration of spline basis functions; this is also where the biggest variations in doublet strength are likely to be seen. Some tuning of the scheme is possible, for instance by moving the mid-panel collocation points closer to the edges (or on to the trailing edge to coincide with the gradient constraints), but this is unlikely to have a significant impact on the




Figure 3.15: Typical location of constraints on a patch.

solution and should not be necessary. Note that collocation points cannot be moved on to edges which have no wake attached, as the velocity field there is singular; nor should they be moved very close to the edges, otherwise inefficient high-order subdivision would occur during calculation of influence coefficients.

Configuration for thin bodies

On a patch representing a thin body, the equations associated with each point fall into three categories:

- All interior points are Neumann collocation points;
- Points on leading edges (or side edges treated as leading edges) are zero doublet strength constraints;
- Points on trailing edges (edges to which wakes are attached) are zero doublet gradient constraints.

This scheme is shown in figure 3.16. For a B-spline representation of doublet distribution it is not possible to apply a boundary condition of infinite doublet gradient at the leading edge (see section 3.1.5), but a steep gradient will in any case be produced here as a result of the combined effect of the other constraints. Alternatively, the square-root basis function described by equation 3.117 can be




Figure 3.16: Constraints on a patch representing a thin body with a wake. The body and wake are separated to make it clear which constraints belong to which patch, but in reality the wake will of course be joined to the body's trailing edge.

used to replace the edge B-spline basis functions. In this case, the square-root basis function sets the edge value of doublet strength to zero implicitly, and the edge constraints should be removed reducing the number of unknowns and equations uniformly.

Configuration for thick bodies

On a patch representing a thick body, the constraints are as follows:

- All interior points are Dirichlet collocation points;
- One of the two coincident points at each knot on the trailing edge is a Dirichlet collocation point; the other is a zero-pressure jump Kutta condition;
- Points on the remaining edges will generally be attached to other patches, and will form continuity of doublet value and gradient constraints with them.

This scheme is shown in figure 3.17.




Figure 3.17: Constraints on a patch representing a thick body with a wake. Once again, the body and wake are separated only to make it clear which constraints belong to which patch.

Configuration for wakes

On a patch representing a wake, all constraints simply equate the doublet strength to a corresponding doublet strength (or doublet strength delta) on the trailing edge of the wing, according to equations 3.45 and 3.46. Note that although the wake is shown here comprising several streamwise panels, for a steady-state case it can just as well be represented by a single constant strength spline in the stream direction, in which case there would only be one column of constraint equations. In fact it makes little difference which approach is chosen, as a single streamwise panel would in any case be sub-divided automatically for the purpose of determining influence coefficients.

3.4 Implementation details

3.4.1 Matrix assembly and solution

Each patch produces a square set of $N_i \times N_j$ equations corresponding to $N_i \times N_j$ unknowns. Once the equations for each patch have been determined, some of which simply relate unknowns on one patch to unknowns on another, they are combined into a single square matrix for solution via standard techniques.

Although the assembly of patches into a single matrix equation is not a trivial process, it is a problem that has been solved many times before in conventional potential flow and finite element codes, so will not be repeated here.

3.4.2 Main algorithm

Figures 3.18 and 3.19 present the main elements of the algorithms used to find the potential flow solution around a thin body with Neumann boundary conditions. The corresponding algorithms for a thick body with Dirichlet boundary conditions are very similar.

3.4.3 Efficiency improvements

A few methods of improving the efficiency of the scheme have already been discussed, including re-use of integration points, reducing cubature orders in the far-field, etc. In addition to these relatively high-level algorithm design concepts, there are two lower-level implementation oriented schemes that bring such a substantial efficiency benefit that they should be included. A large proportion of the computational effort in running this code is taken up in sampling B-spline curves or surfaces, and both the schemes described below can significantly reduce this effort.

B-Spline polynomial pre-computation

B-spline curves and surfaces are used extensively within the code to represent the singularity distributions, and may also be used to represent the geometry. In all cases, the piecewise-polynomial shape of each basis function that makes up the B-spline is defined only by the order of the B-spline and its knot vector. Using the standard Cox-de-Boor recursion formulae (Rogers & Adams, 1990) to evaluate a point on the B-spline effectively means re-calculating the polynomial coefficients of each basis function every time, with considerable computational overhead. Significant efficiency gains have been made here by converting the Cox-de-Boor recursion formulae to calculate the equivalent polynomial coefficients of each basis function, rather than the value of a sample point on

<p>SolveFlow</p> <p>Initialisation <i>For each patch</i></p> <ul style="list-style-type: none"> ▷ Initialise velocity influence matrix to zero [$3 \times N_c \times N_s$] ▷ Initialise normal velocity influence matrix to zero [$1 \times N_c \times N_s$] <p><i>Next</i></p> <p>Main loop <i>For each collocation point</i></p> <ul style="list-style-type: none"> ▷ Obtain position and surface normal of collocation point <p><i>For each patch</i></p> <ul style="list-style-type: none"> <i>For each panel within the patch</i> <ul style="list-style-type: none"> ▷ Obtain panel integration points, etc. (see PanelCoefficients) <i>Next panel</i> <i>For each super-panel within the patch</i> <ul style="list-style-type: none"> ▷ Obtain super-panel influence (see SuperPanelInfluence) ▷ Store velocity influence of super-panel on collocation point ▷ Compute and store normal component of velocity influence <i>Next super-panel</i> <p><i>Next patch</i></p> <p><i>Next collocation point</i></p> <p>Solve system</p> <ul style="list-style-type: none"> ▷ Combine influences matrices from all patches ▷ Apply additional constraints and boundary conditions ▷ Solve system to determine super-panel weights <p>Post-processing</p> <ul style="list-style-type: none"> ▷ Find collocation point velocities from super-panel weights and stored influences ▷ Determine pressures, forces and moments

Figure 3.18: Outline algorithm for finding potential flow solution around a thin body with Neumann boundary conditions

PanelCoefficients
<p><i>If collocation point is on surface (or edge) of panel</i></p> <ul style="list-style-type: none"> ▷ Generate set of desingularised integration points ▷ Store intrinsic coordinates of each integration point ▷ Store unit velocity influence of each integration point (scaled by quadrature weight and jacobian) ▷ Store 'free-term' influences relating to unit value and gradient <p><i>Else</i></p> <ul style="list-style-type: none"> ▷ Sub-divide panels as necessary (recursive) <p><i>For each sub-panel</i></p> <ul style="list-style-type: none"> ▷ Generate set of regular integration points ▷ Store intrinsic coordinates of each integration point ▷ Store unit velocity influence of each integration point (scaled by quadrature weight and jacobian) <p><i>Next sub-panel</i></p> <p><i>End</i></p>

SuperPanelInfluence
<p>Collect integration points from all panels within super-panel</p> <ul style="list-style-type: none"> ▷ Combine regular and de-singularised lists ▷ Obtain list of intrinsic coordinates ▷ Obtain corresponding list of scaled unit influences <p>Multiply influences by shape functions and sum</p> <ul style="list-style-type: none"> ▷ Determine super-panel shape function weight at each integration point ▷ Multiply scaled unit influences by weights ▷ Sum weighted scaled influences for all integration points <ul style="list-style-type: none"> ▷ Determine super-panel value and gradients at collocation point ▷ Multiply all panels' free-term influences by values and gradients <ul style="list-style-type: none"> ▷ Add integration point and free-term influences ▷ Return total influence of super-panel

Figure 3.19: Outline algorithm of PanelCoefficients and SuperPanelInfluence routines

the spline. Subsequently when a sample point is required, the basis functions can be evaluated as simple polynomials, without any further need for recursion formulae.

To obtain the polynomial form of the B-spline basis functions $U_{i,b}$, we start with the standard Cox-de-Boor recursion formulae written in expanded form

$$U_{i,1}(\xi) = \begin{cases} 1 & \text{if } Z_i \leq \xi < Z_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (3.132a)$$

and

$$U_{i,b}(\xi) = \frac{\xi U_{i,b-1}(\xi) - Z_i U_{i,b-1}(\xi)}{Z_{i+b-1} - Z_i} + \frac{Z_{i+b} U_{i+1,b-1}(\xi) - \xi U_{i+1,b-1}(\xi)}{Z_{i+b} - Z_{i+1}} \quad (3.132b)$$

where Z_i are the elements of the knot vector.

We wish to obtain the basis function of order b in the piecewise polynomial form

$$U_{i,b}(\xi) = \sum_{\kappa=0}^{b-1} A_{\kappa,m,i,b} \xi^\kappa \quad Z_m \leq \xi < Z_{m+1} \quad (3.133)$$

where to clarify the indices of $A_{\kappa,m,i,b}$

- κ is the power of ξ in the polynomial expansion
- m is the index of the segment of the piecewise polynomial
- i is the number of the basis function
- b is the order of the basis function

Considering a first order (piecewise constant) B-spline, equation 3.132a implies

$$A_{0,m,i,1} = \begin{cases} 1 & \text{if } i = m \\ 0 & \text{otherwise} \end{cases} \quad (3.134a)$$

For higher order B-splines $b \geq 2$, by considering the coefficients of powers of ξ within the recursion relationship of equation 3.132a, we can obtain

$$A_{\kappa,m,i,b} = \frac{A_{\kappa-1,m,i,b-1} - Z_i A_{\kappa,m,i,b-1}}{Z_{i+b-1} - Z_i} + \frac{Z_{i+b} A_{\kappa,m,i+1,b-1} - A_{\kappa-1,m,i+1,b-1}}{Z_{i+b} - Z_{i+1}} \quad (3.134b)$$

where $A_{-1,m,i,b}$ is defined as zero.

The derivatives of the basis functions can be obtained directly as from the polynomial expansions. For example

$$\frac{\partial U_{i,b}}{\partial \xi} = \sum_{\kappa=1}^{b-1} \kappa A_{\kappa,m,i,b} \xi^{\kappa-1} \quad Z_m \leq \xi < Z_{m+1} \quad (3.135)$$

and

$$\frac{\partial^2 U_{i,b}}{\partial^2 \xi} = \sum_{\kappa=2}^{b-1} \kappa(\kappa-1) A_{\kappa,m,i,b} \xi^{\kappa-2} \quad Z_m \leq \xi < Z_{m+1} \quad (3.136)$$

In practice, separate polynomial coefficients are stored for values, first derivatives and second derivatives to produce the fastest possible B-spline evaluations.

B-Spline point pre-computation

Further efficiency gains can be made by recognising that in the majority of cases, a particular sample point on a B-spline surface is evaluated more than once. This is particularly true if B-spline surfaces are used to represent the geometry (and if not, similar efficiency gains can probably be made with whatever function is used). For example, although the re-use described in section 3.3.7 prevents recalculation of integration points on a panel relating to a specific collocation point, it is likely that the same integration points will be used for many other collocation points. Even if panel-subdivision and variable cubature lead to different sets of integration points for different collocation points, there will be significant overlap between the cases. Similarly, if the same B-spline surface orders and knot vectors are used for geometry and singularity distributions, every integration point will duplicate a sample on the B-spline surface.

By employing B-spline polynomial pre-computation as described in the previous section, the cost of each B-spline sample is significantly reduced, but

it is still not negligible, and efficiency can be increased by storing results. This requires a look-up facility to be generated for a specific B-spline surface that can return a sample value as a function of the two intrinsic coordinates of the surface. This could either be generated beforehand for likely sample points (e.g. all orders of cubature within each panel, possibly including some sub-divided panels), or results of every B-spline sample could be added to an expanding look-up facility. For geometric surfaces, it will be better to store three dimensional physical coordinates of each point rather than just the basis-function weights of each sample.

3.5 Post-processing of velocities, pressures and forces

Once singularity distributions over each patch are determined, surface pressures and total lift and induced drag forces can be calculated directly via Bernoulli's equation. In contrast to conventional panel and vortex lattice methods, no interpolation is required at this stage, and the resulting pressure distributions can be considered exact with respect to singularity distributions. The accuracy of total aerodynamic forces is therefore subject only to the number of cubature points chosen over which to integrate the pressures.

3.5.1 Surface doublet gradients

The surface gradient $\nabla_{\mathbf{x}}\mu(\boldsymbol{\xi})$ is required for the calculation of surface velocities, and to evaluate this gradient, we introduce a third intrinsic coordinate ξ_3 which is defined by

$$\frac{\partial \mathbf{x}}{\partial \xi_3} = k \frac{\partial \mathbf{x}}{\partial \xi_1} \times \frac{\partial \mathbf{x}}{\partial \xi_2} \quad (3.137)$$

such that it is always normal to the surface on which ξ_1 and ξ_2 lie. The constant k is introduced purely to demonstrate in the final result that the scaling of the off-surface dimension is arbitrary; in practice, a value $k = 1$ can be used.

We can then write the surface gradient of μ with respect to ξ in terms of known quantities

$$\nabla_{\xi}\mu(\xi) = \begin{Bmatrix} \frac{\partial\mu}{\partial\xi_1} \\ \frac{\partial\mu}{\partial\xi_2} \\ 0 \end{Bmatrix} = \begin{bmatrix} \frac{\partial x_1}{\partial\xi_1} & \frac{\partial x_2}{\partial\xi_1} & \frac{\partial x_3}{\partial\xi_1} \\ \frac{\partial x_1}{\partial\xi_2} & \frac{\partial x_2}{\partial\xi_2} & \frac{\partial x_3}{\partial\xi_2} \\ \frac{\partial x_1}{\partial\xi_3} & \frac{\partial x_2}{\partial\xi_3} & \frac{\partial x_3}{\partial\xi_3} \end{bmatrix} \begin{Bmatrix} \frac{\partial\mu}{\partial x_1} \\ \frac{\partial\mu}{\partial x_2} \\ \frac{\partial\mu}{\partial x_3} \end{Bmatrix} = \mathbf{G} \begin{Bmatrix} \frac{\partial\mu}{\partial x_1} \\ \frac{\partial\mu}{\partial x_2} \\ \frac{\partial\mu}{\partial x_3} \end{Bmatrix} \quad (3.138)$$

In this equation, the third element of $\nabla_{\xi}\mu(\xi)$ is defined as zero because the relevant gradient here is that on the surface to which the ξ_3 axis is normal. The quantity we originally required, the surface gradient of μ with respect to \mathbf{x} , is therefore

$$\nabla_{\mathbf{x}}\mu(\xi) = \begin{Bmatrix} \frac{\partial\mu}{\partial x_1} \\ \frac{\partial\mu}{\partial x_2} \\ \frac{\partial\mu}{\partial x_3} \end{Bmatrix} = \mathbf{G}^{-1} \begin{Bmatrix} \frac{\partial\mu}{\partial\xi_1} \\ \frac{\partial\mu}{\partial\xi_2} \\ 0 \end{Bmatrix} \quad (3.139)$$

The matrix \mathbf{G}^{-1} can be written as an analytic expression containing only the derivatives of \mathbf{x} with respect to ξ_1 and ξ_2 , so that no matrices need to be inverted within the post-processing code. This inverse expression is rather lengthy, so is omitted from this text; however, when it is written out it is seen that the constant k only appears in the third column of \mathbf{G}^{-1} , which is multiplied by zero, confirming that its choice was indeed arbitrary.

3.5.2 Field point velocities

Velocities at general points within the fluid region are calculated using the same scheme as that used to determine collocation point influences, this time summing up contributions from all super-panels and multiplying by the now known super-panel weights. Thus, combining equations 3.113 and 3.124, the

velocity induced by a doublet field on a point \mathbf{y} is

$$\mathbf{q}^\mu(\mathbf{y}) = \sum_{i=1}^{N_i} \sum_{j=1}^{N_j} B_{ij}^\mu \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} w_m w_n U_i^{\mu 1}(\xi_{1m}) U_j^{\mu 2}(\xi_{2n}) \mathbf{T}^\mu(\boldsymbol{\xi}_{mn}, \mathbf{y}) J(\boldsymbol{\xi}_{mn}) \quad (3.140)$$

where the same strategies of desingularisation, panel subdivision, variable cubature, re-use of integration points etc. are employed. Similarly, the velocity induced by a source distribution on the patch is

$$\mathbf{q}^\sigma(\mathbf{y}) = \sum_{i=1}^{N_i} \sum_{j=1}^{N_j} B_{ij}^\sigma \sum_{m=1}^{N_m} \sum_{n=1}^{N_n} w_m w_n U_i^{\sigma 1}(\xi_{1m}) U_j^{\sigma 2}(\xi_{2n}) \mathbf{T}^\sigma(\boldsymbol{\xi}_{mn}, \mathbf{y}) J(\boldsymbol{\xi}_{mn}) \quad (3.141)$$

The total velocity at a field point \mathbf{y} is therefore

$$\mathbf{q}(\mathbf{y}) = \mathbf{Q}^\infty + \sum_1^{N_B} \sum_1^{N_T} [\mathbf{q}^\mu(\mathbf{y}) + \mathbf{q}^\sigma(\mathbf{y})] \quad (3.142)$$

where the sum is over the N_T patches on each of the N_B bodies.

In comparison to the computational cost of evaluating a collocation point influence, a velocity evaluation could be more expensive for a field point just off a patch surface, or cheaper for a point distant from all surfaces. However, note that off-surface field point velocities are only required for specific types of post-processing such as the generation of streamlines for visualisation purposes.

3.5.3 Surface velocities

For a thick body solved with Dirichlet boundary conditions, the velocity on the outer surface of a body can be found more directly by recalling that the internal velocity potential was set to equal the free-stream potential (equation 3.15). The potential on the outer surface can then be found via equation 3.10, which equates the potential jump across the surface to the local doublet strength. The outer surface potential is therefore written as the sum of the internal potential and the

potential jump:

$$\begin{aligned}\phi^T(\mathbf{x}_-) &= \phi^\infty + [\phi^T(\mathbf{x}_-) - \phi^T(\mathbf{x}_+)] \\ &= \phi^\infty - \mu(\mathbf{x})\end{aligned}\tag{3.143}$$

where \mathbf{x} , \mathbf{x}_- and \mathbf{x}_+ are all functions of the same value of ξ . Since the source distribution is configured to remove the normal component of the free-stream velocity, the surface velocities found from the gradient of the above potential can be written

$$\begin{aligned}\mathbf{q}(\mathbf{x}_-) &= \nabla_{\mathbf{x}}\phi^T(\mathbf{x}_-) = \nabla_{\mathbf{x}}\phi^\infty - \nabla_{\mathbf{x}}\mu(\mathbf{x}) \\ &= \mathbf{Q}^\infty - [\mathbf{Q}^\infty \cdot \mathbf{n}(\mathbf{x})] \mathbf{n}(\mathbf{x}) - \nabla_{\mathbf{x}}\mu(\mathbf{x})\end{aligned}\tag{3.144}$$

where $\nabla_{\mathbf{x}}\mu(\mathbf{x})$ is given by equation 3.139.

For a thin body, no such shortcut is available, and on-surface velocities must be found in the same way as collocation point influences, i.e. by use of equation 3.142 making appropriate substitutions for super-panels where desingularised influences are required (see section 3.2.3). However, the velocity produced in this way is the ‘mid-surface’ velocity, to which an additional component must be added. This discontinuous tangential velocity component results from the local surface doublet gradient and is given by

$$\mathbf{q}^*(y_\pm) = \mp \frac{1}{2} \nabla_{\mathbf{x}}\mu(\mathbf{y})\tag{3.145}$$

where $\nabla_{\mathbf{x}}\mu(\mathbf{x})$ is given by equation 3.139.

Surface velocity calculations are trivial at collocation points, where influences have already been calculated, but elsewhere on the surface the computation is expensive. Depending on relative requirements of run-time and accuracy, it may be appropriate to only evaluate mid-surface velocities at collocation points, and interpolate these to other surface points at which velocities are required. The tangential component due to the surface doublet gradient would then be evaluated directly at each sample point, and added to the interpolated

mid-surface component. It would then be prudent to subtract any resulting velocity component parallel to the local surface normal, which may result from interpolation errors. This strategy is quite reasonable, as the doublet gradient component will generally vary more quickly over the surface than the mid-surface component.

3.5.4 Surface pressures and forces

Once surface velocities have been obtained, it is straightforward to determine pressures via Bernoulli's equation

$$p(\mathbf{y}_{\pm}) = p_0 - \frac{1}{2}\rho |\mathbf{q}(\mathbf{y}_{\pm})|^2 \quad (3.146)$$

where p_0 is the static pressure. For thin bodies, upper and lower surface pressures are usually required, whereas for thick bodies only the outer surface pressures $p(\mathbf{y}_-)$ are relevant.

Forces can then be obtained by integrating the pressures acting in the direction of the local surface normal. For example, integrating over a thick body

$$\mathbf{F}_B = \int_{S_B} p(\mathbf{y}_-) \mathbf{n}(\mathbf{y}) dA \quad (3.147)$$

By use of a standard ($N_m \times N_n$) cubature rule (e.g. Legendre-Gauss, Engels (1980)) on each of the ($N_p \times N_q$) panels on each of the N_T patches, this can be evaluated numerically as

$$\mathbf{F}_B = \sum_k^{N_T} \sum_i^{N_p} \sum_j^{N_q} \sum_m^{N_m} \sum_n^{N_n} w_m w_n p(\mathbf{y}_{-kijmn}) \mathbf{n}(\mathbf{y}_{kijmn}) J(\mathbf{y}_{kijmn}) \quad (3.148)$$

For thin surfaces, the integration is carried out for upper and lower surfaces. Moments are obtained in a similar fashion.

3.5.5 Leading edge suction forces

At the leading edge (and other non-wake edges) of a thin aerofoil there is a suction force which can be thought of as an infinite pressure acting over an infinitesimal

area, as discussed in section 3.1.5. This force must be added to the surface loads already obtained to provide estimates of total lift and induced drag. Calculation of this force follows the method presented by Hancock & Garner (1975).

At the leading edge, the potential jump across the surface of thin aerofoil has the form

$$\Delta\phi = C\sqrt{s} \quad (3.149)$$

where s is the distance to the leading edge and C is the leading edge suction coefficient. The gradient of the potential jump is clearly singular as $s \rightarrow 0$. The force associated with this singularity can be approximated by

$$F = \int_{LE} \frac{\pi\rho C^2}{16} ds \quad (3.150)$$

This force is assumed to act in the tangent plane of the lifting surface and perpendicular to the leading edge itself.

Within the CSM, this force can be estimated by fitting square-root basis functions of the form

$$\mu = \Delta\phi = C\sqrt{s} \quad (3.151)$$

to the continuous doublet distribution at a series of stations along the leading edge, and integrating via a one dimensional quadrature rule. Alternatively, if square-root basis functions are employed in the definition of the doublet surface (see section 4.1.5), the suction coefficients C can be determined directly from the corresponding basis function weights.

3.5.6 Forces calculated in the Trefftz plane

In an alternative method of load calculation, the momentum equation is applied to a control volume surrounding the entire system of lifting bodies. The control volume is chosen to be a large rectangular prism, extending well away from the locality of the bodies. The perturbation velocities induced by the wing therefore vanish everywhere except on the plane through which the wake passes, known as the Trefftz plane. When the sides of the control volume are aligned with the free-stream velocity, the momentum equation reduces to simple line integrals

traversing the wake. For a straight wake section, the lift and induced drag in the Trefftz plane are given by (Katz & Plotkin, 2001)

$$F_L = \rho Q_\infty \int_C \mu(x) dx \quad (3.152)$$

$$F_{D_i} = -\frac{\rho}{2} \int_C \mu(x) q(x) dx \quad (3.153)$$

where $q(x)$ is the downwash induced by the wake's vorticity. For a curved wake, where the direction of the lift is known only to be parallel to the Trefftz plane (and perpendicular to the free-stream), the corresponding vector equations are

$$\mathbf{F}_L = \rho \mathbf{Q}_\infty \times \int_C \mu(\xi_1) \frac{\partial \mathbf{x}}{\partial \xi_1} d\xi_1 \quad (3.154)$$

$$\mathbf{F}_{D_i} = -\frac{\rho}{2} \int_C \mu(\xi_1) \mathbf{q}_i(\xi_1) \times \frac{\partial \mathbf{x}}{\partial \xi_1} d\xi_1 \quad (3.155)$$

where $\mathbf{q}_i(\xi_1)$ is the velocity induced parallel to the Trefftz plane, and ξ_1 is the intrinsic surface coordinate spanning the wake.

Within the CSM these forces are evaluated numerically using a one dimensional quadrature rule. Since the Trefftz plane method calculates the total fluid dynamic force generated by the system of lifting bodies, a separate evaluation of the leading edge suction force is not required.

Chapter 4

Results and analysis

4.1 Parameter studies

Before the theory and implementation of the CSM presented in chapter 3 can be employed to solve a potential flow problem, a configuration must be specified to define how the panels and collocation points are set out over the body or bodies in the flow, and what basis functions are used in the doublet surface representation. In this section a number of parameter studies are presented that systematically investigate panel layout, collocation point layout and basis function schemes. The aim is to provide an understanding of some characteristics of the CSM, and to determine an optimal configuration for the analysis of a typical body in a free-stream.

Whilst it is tempting for these studies to choose a regular body for which there is a known solution, such as the rectangular flat plate described section 4.4, this approach is likely to lead to a very restricted understanding of the method. Many aspects of typical potential flow methods are not encountered until surfaces are distorted and panels are skewed, so an irregular three dimensional curved surface was considered a more useful specimen. The main feature of the CSM that sets it apart from other continuous methods is its ability to solve thin-body problems with a velocity formulation; as such, a thin surface is chosen.

With these considerations in mind, it is appropriate to use a mainsail geometry of an America's Cup yacht as the test body, which is shown in figure 4.1. A wake with a fixed, pre-defined geometry that extends approximately 100 times the chord length is modelled. The wake leaves the sail's trailing edge smoothly, but curves towards the free-stream direction within one chord length. A uniform free-stream with a nominal angle-of-attack of 18° is applied throughout.




Figure 4.1: The sail body used for the parameter study, with axis coordinates shown in m. The grey grid contains lines of constant ξ_1 and ξ_2 . The two labelled points are sample points at which surface pressures are calculated.

Different panel and collocation point configurations are tested against criteria of accuracy and convergence rate for a number of typical post-processed fluid dynamics quantities. These are:

- Total lift and induced drag coefficients evaluated in the Trefftz plane
- Total lift and induced drag coefficients calculated by pressure integration over the body
- Edge suction forces that arise on non-wake edges
- Point pressure samples at two discrete points on the sail (see figure 4.1)

The evaluation of on- and off-surface velocities is covered in sections 4.5.5 and 4.5.6.

Also shown in the graphs that follow is the matrix condition number, which provides a relative indication of the solution's sensitivity to errors. A low condition number implies a well-conditioned matrix whose solution will be relatively insensitive to errors in its coefficients. A very high condition number corresponds to a near-singular matrix, indicative of a poorly formed solution.

Results are plotted against the total number of panels used to represent the doublet surface over the sail. In all cases, the ratio of chordwise panels to spanwise panels has been kept constant for a particular convergence curve, thus increasing panel densities along both axes simultaneously. Each curve represents a comprehensive set of around twenty results: for example, in configurations with $(N \times N)$ panels, results are shown for solutions with 4, 9, 16... 400 sail panels.

Many of the graphs that follow show a horizontal shaded region which is labelled 'assumed 1% error band'. For the force or coefficient that is presented by the graph, this shaded region represents the range:

$$\text{Converged result} - 1\% \longrightarrow \text{Converged result} + 1\%$$

As the potential flow around the curved sail body has not been solved analytically, the exact results are not in fact known. However, they can be estimated with reasonable confidence by examining the convergence of vortex

lattice and CSM solutions with very high panel densities, where the specific combination of simulations used to determine these converged values has varied depending on the quantity being measured.

For the case of the leading-edge suction force, accurate values could not be determined with confidence. Here, a dotted line is used to indicate what is thought to be the approximate answer. (Note that the total of all edge suction forces can be determined accurately from the difference in the Trefftz plane and the pressure integration forces; however, as some of the studies in this section relate differently to the leading and side edges, the leading-edge suction force is a more useful quantity to compare.)

In all cases, for the calculation of edge suction force the method of fitting a square-root curve to a point on the boundary of the first and second column of panels is used, as described in the following section.

4.1.1 Edge suction forces via a square-root fit

Two methods of determining edge-suction forces were described in chapter 3. For certain types of basis functions that contain leading-edge square-root characteristics, it is possible to determine edge-suction forces directly from basis function weights, and this approach is investigated in section 4.1.5. For more standard basis function sets, a square-root fit to the doublet surface shape is required.

With a conventional vortex lattice or surface panel code it is natural to fit this square-root function to the doublet strength at the first row of collocation points: with constant doublet strength over each panel, there is in fact little choice. However, as the CSM produces a curved doublet sheet, it is less obvious what we should fit to. It could be argued that the fit should be carried out very close to the leading edge where the square-root profile assumption is most applicable, but in practice it might be more reliable to fit further downstream where distance and doublet strength are larger, and numerical errors might be reduced. (The fit could of course be over a region, rather than at a point, possibly using a least-

squared error approach, but it was not considered likely that this would improve the results, and such a method was not investigated.)

Edge suction forces will exist on the leading edges and the side edges of the sail geometry, as shown in figure 4.3. In this section, the leading edge suction force alone is considered, although results will generally be applicable to other non-wake edges.

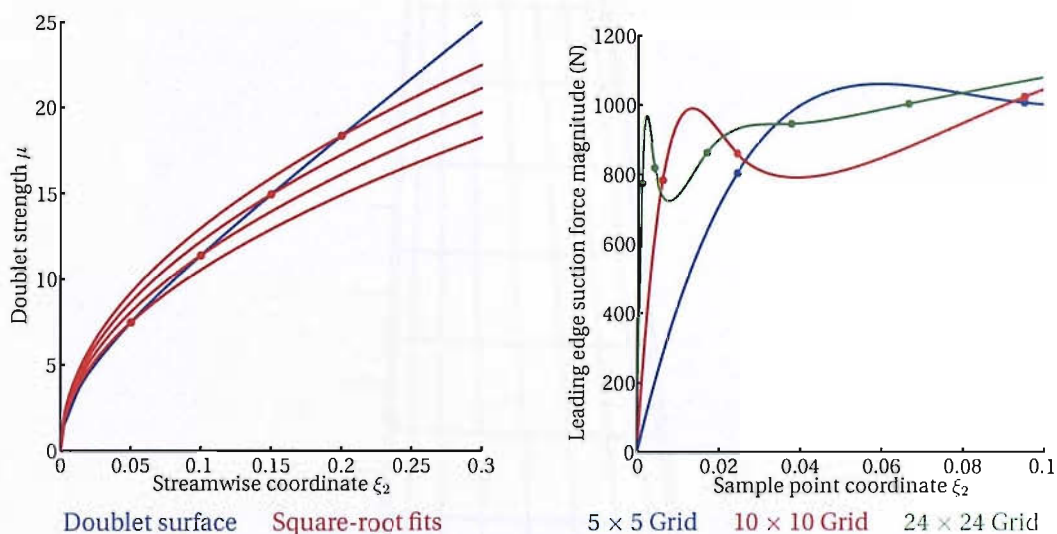


Figure 4.2: (Left) The fitting point affects the magnitude of the square-root fit. (Right) The influence of the fitting point and panel density on the magnitude of the square-root fit.

The result of fitting to different points in the vicinity of the leading edge is shown in figure 4.2 (left), and it is clear that estimates of edge suction force will vary significantly depending on which point is chosen. This dependency is shown more explicitly in figure 4.2 (right), which shows how the magnitude of the leading edge suction force varies with the position of the sample point. Three curves are shown, indicating how the characteristic changes depending on the panel density. It is seen that the estimated edge suction force is strongly dependent on both the panel density and the position of the point used to fit the square-root curve to.

The dots on the curves show the positions of the collocation points, and some correlation is seen between the curve shapes and these point positions; it therefore seems sensible to base the choice of sample point in relation to




Figure 4.3: Suction forces on non-wake edges. The lengths of the blue arrows are proportional to the suction force magnitude per unit length. These were calculated by the fitting method with parameters thought to be reasonably accurate.

collocation point positions. For instance, the first or second collocation point could be chosen as the sample point.

Figure 4.4 shows how the choice of position affects the convergence of the edge force with increasing numbers of panels. The dotted line shows the converged force magnitude obtained from the vortex lattice results, but as seen in section 4.5, it is likely that this result is quite inaccurate. As such, it is not known which of the above solutions for the CSM is the most accurate. We can see though that the solution formed by sampling at 0.25 of the distance to second collocation point converges much more quickly than the others.

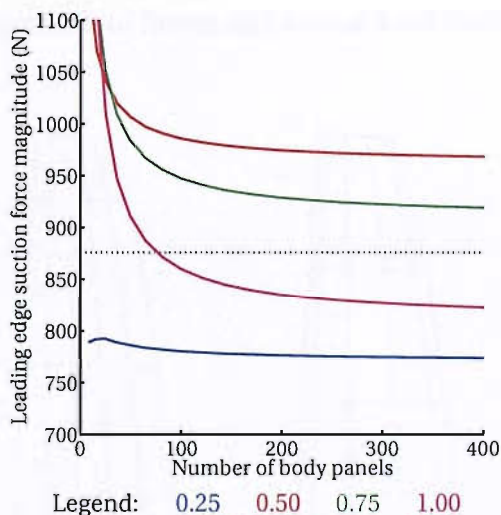


Figure 4.4: Effect of sample point position on leading edge suction force convergence. The numbers corresponding to the curves represent the sample point position as a fraction of the distance from the leading edge to the second collocation point.

4.1.2 Panel spacing

In general, the form of knot vector with repeated end knots described in section 3.3.1 has been used, and here we consider how the spacing of the internal knots in the chordwise and spanwise vectors affect the solution. Specifically, we compare a linear knot spacing, where each panel has the same (intrinsic) length, with a cosine knot spacing, where panels are concentrated towards both edges.

Excluding the repeated end knots, the two forms of knot vector are given by:

$$Z_i = \frac{i}{N_i} \quad i \in \{0, 1 \dots N_i\} \quad (4.1)$$

for the linear form, and

$$Z_i = \frac{1}{2} \left\{ 1 - \cos\left(\frac{i\pi}{N_i}\right) \right\} \quad i \in \{0, 1 \dots N_i\} \quad (4.2)$$

for the cosine form.

The four panel layouts shown in figure 4.5 all have the same grid size, but they have different combinations of linear and cosine knot vectors in their chord and span.




Figure 4.5: Panel spacings (from left to right): cosine-cosine; cosine span; cosine chord, linear-linear

The chordwise spacing will strongly influence the shape of the panels near the leading and trailing edges; it is therefore not surprising that convergence of the pressure sample at point A, close to the leading edge, is almost entirely controlled by the chordwise spacing, as seen in figure 4.6 (middle-left). In this case, the linear chordwise knot vector produces a more quickly convergent solution.

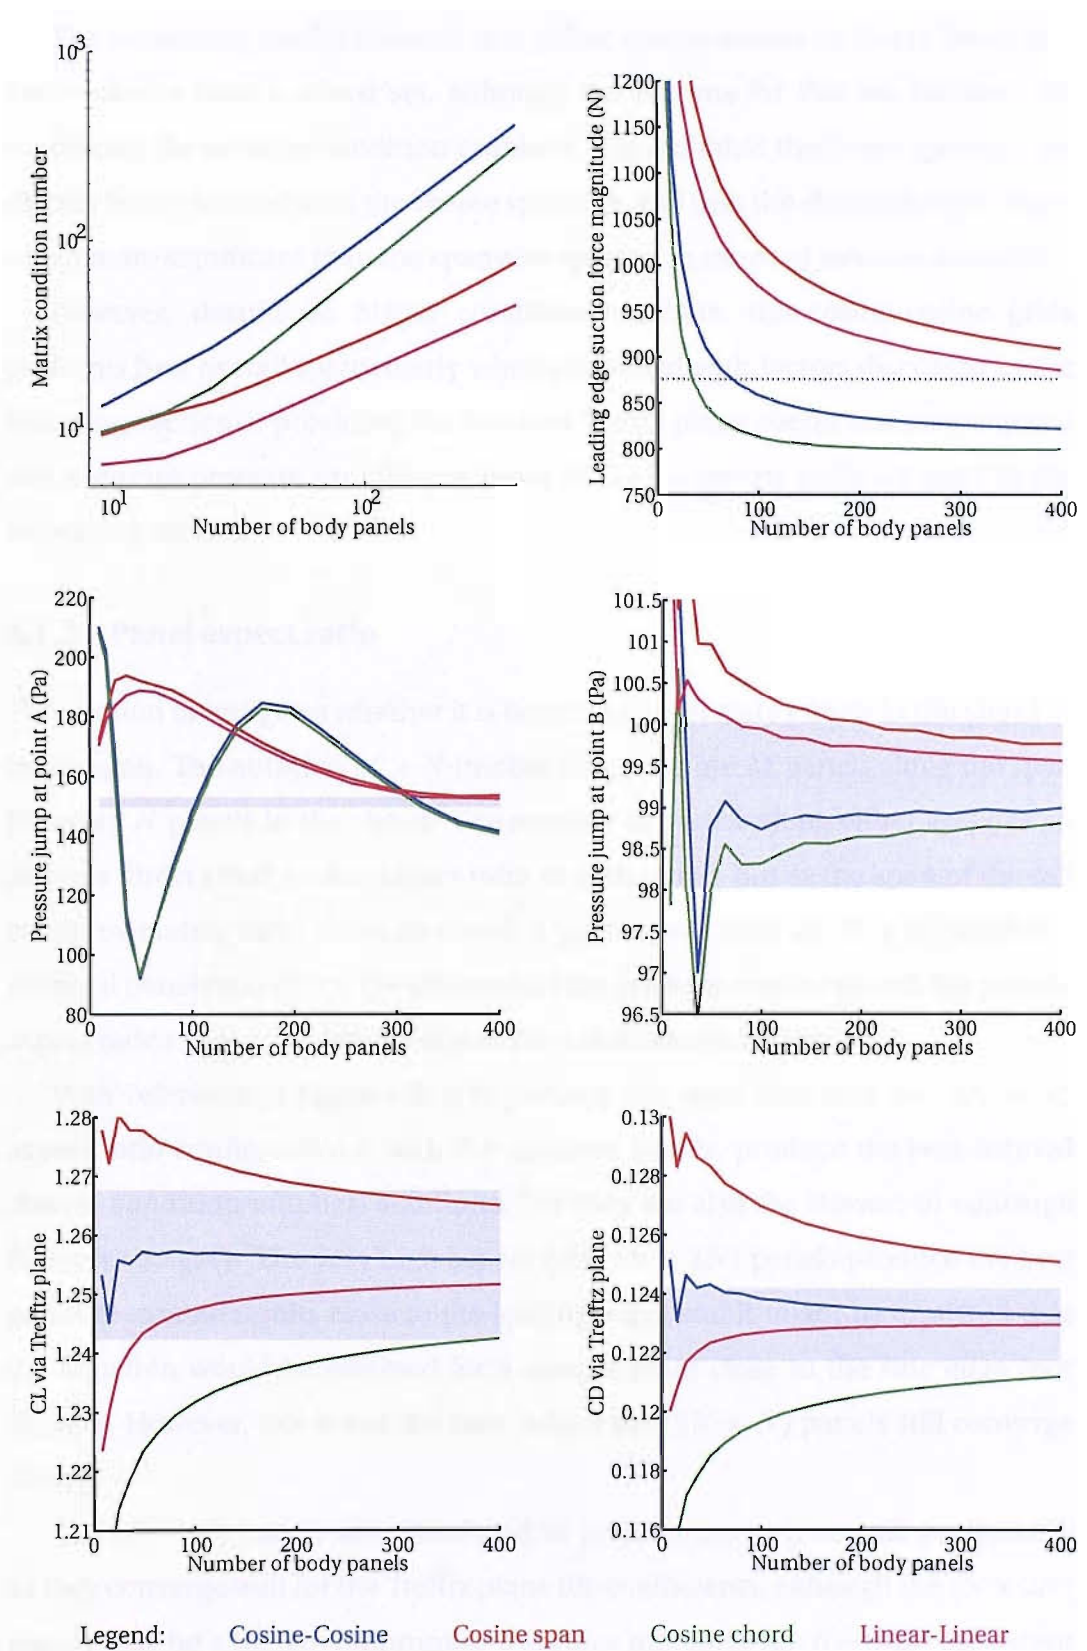


Figure 4.6: Influence of panel spacing on convergence

The remaining results indicate that either cosine-cosine or linear-linear is a better choice than a mixed set, although the reasons for this are unclear. By examining the solution condition numbers, it is seen that the linear spacings are slightly better formed than the cosine spacings, and that the chordwise spacing is much more significant than the spanwise spacing in terms of solution stability.

However, despite its higher condition numbers, the cosine-cosine grids performs best overall (particularly when combined with factors discussed in the following sections), providing the quickest Trefftz plane coefficient convergence and accurate pressure sampling at point B. Cosine-cosine grids are used in the remaining studies.

4.1.3 Panel aspect ratio

This section investigates whether it is better to concentrate panels in the chord or in the span. The notation $M \times N$ implies that there are M panels along the span for every N panels in the chord. The number of panels along either axis has an obvious direct effect on the aspect ratio of each panel, but as the span of the sail is approximately three times its chord, a 'grid aspect ratio' of $(N \times N)$ implies a nominal panel ratio of (3×1) – although as the grids are cosine spaced, the panels' aspect ratios vary considerably across the sail as shown in figure 4.7.

With reference to figure 4.8, it is perhaps not surprising that the $(3N \times N)$ aspect ratio configurations, with the squarest panels, produce the best-formed (lowest condition number) solutions; but they are also the slowest to converge in every category. The very high aspect ratio $(N \times 3N)$ panels produce the best pressure sample results close to the leading edge, and it might be expected that the situation would be reversed for a sample point close to the side edge (not shown). However, this is not the case, where the $(3N \times N)$ panels still converge slowly.

The $(N \times N)$ panels are considered to perform the best overall, particularly as they converge well for the Trefftz plane lift-coefficients, although the $(N \times 2N)$ panels may be a better compromise if surface pressures are the most important output.




Figure 4.7: Grid aspect ratios (from left to right): $(N \times 3N)$, $(N \times 2N)$, $(N \times N)$, $(2N \times N)$

4.1.4 Collocation point positions

For the results shown in previous sections collocation points have been located on spline knots, with additional points at the leading and side edges as discussed in section 3.3.8 and shown in figure 3.16. The chordwise arrangement is also shown in figure 4.9 (top), where it is seen that in general with a cosine-spaced knot vector, collocation points do not fall at the maxima of basis functions.

A number of alternative collocation point schemes were investigated, of which the results of two are shown here. In the first, a collocation point is located at the maximum of each basis function, as shown in figure 4.9 (middle). Whilst this produces a smooth doublet surface and improves the accuracy of the pressure sample near the leading edge of the sail, the convergence of lift and drag coefficients is generally made worse (see figure 4.10).

With the standard collocation point arrangement, the very poor convergence of the pressure jump at sample point A results from the doublet surface being significantly affected by the position of the column of collocation points closest to

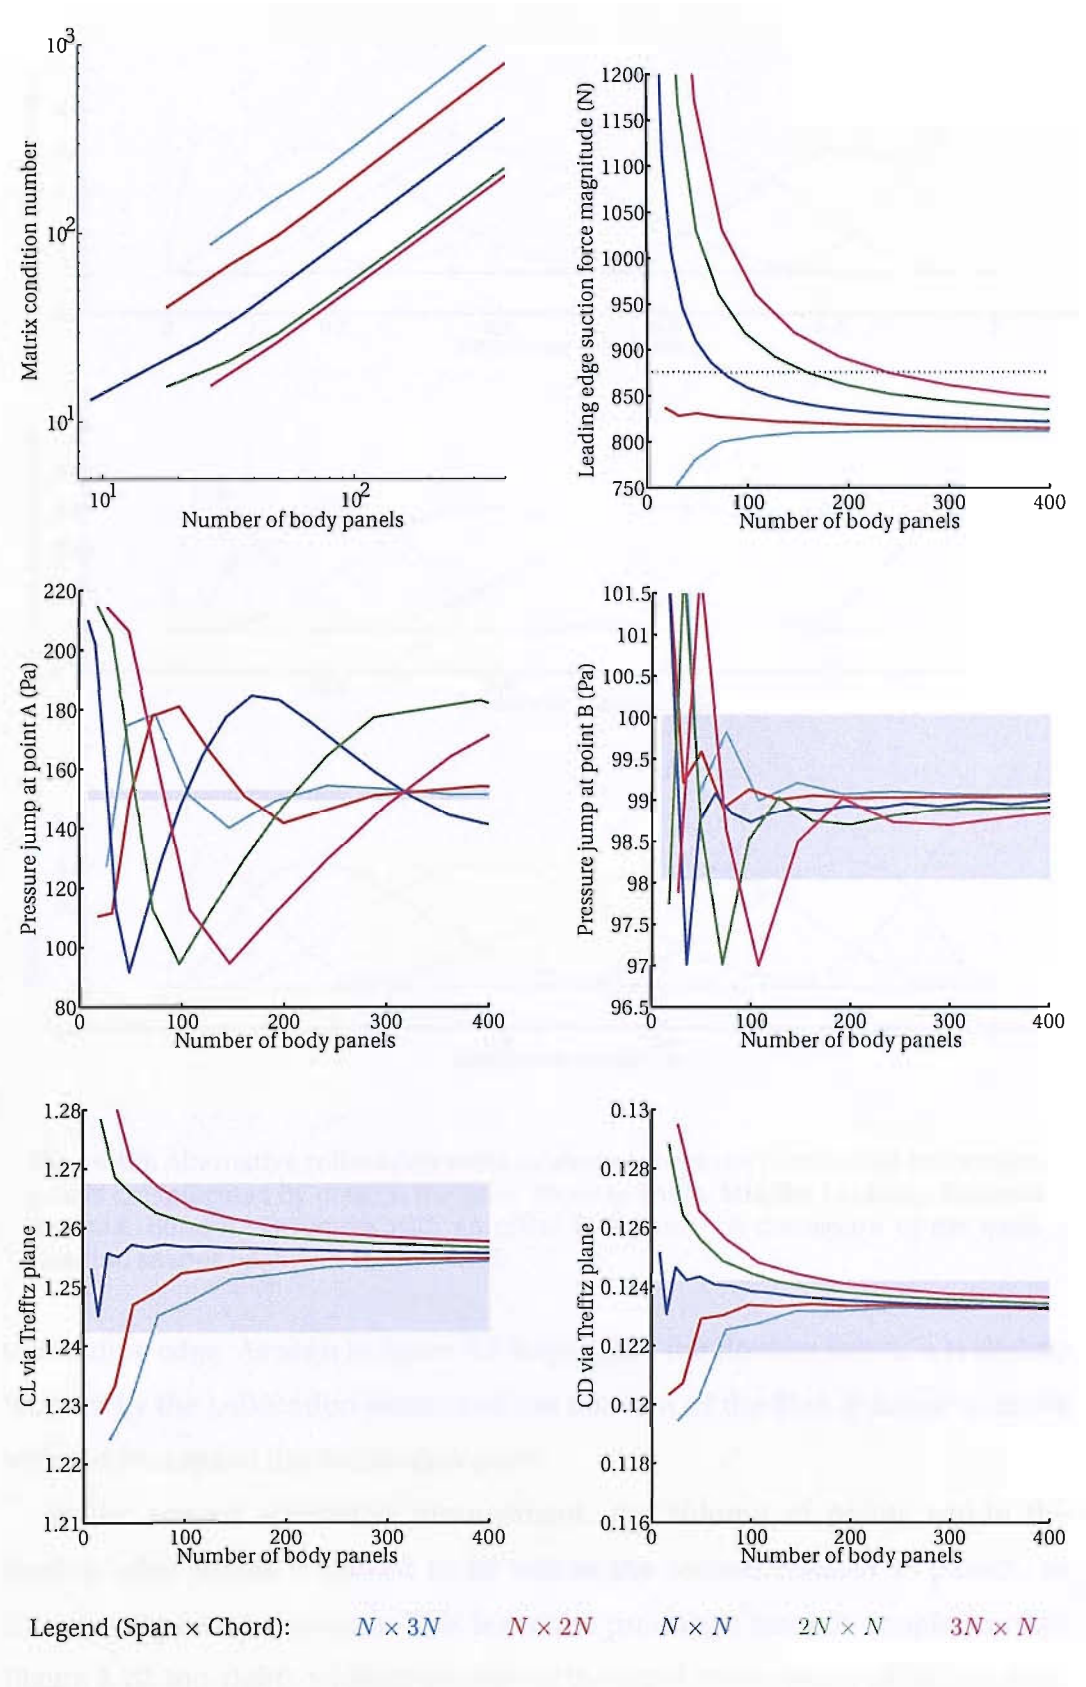


Figure 4.8: Influence of aspect ratio on convergence

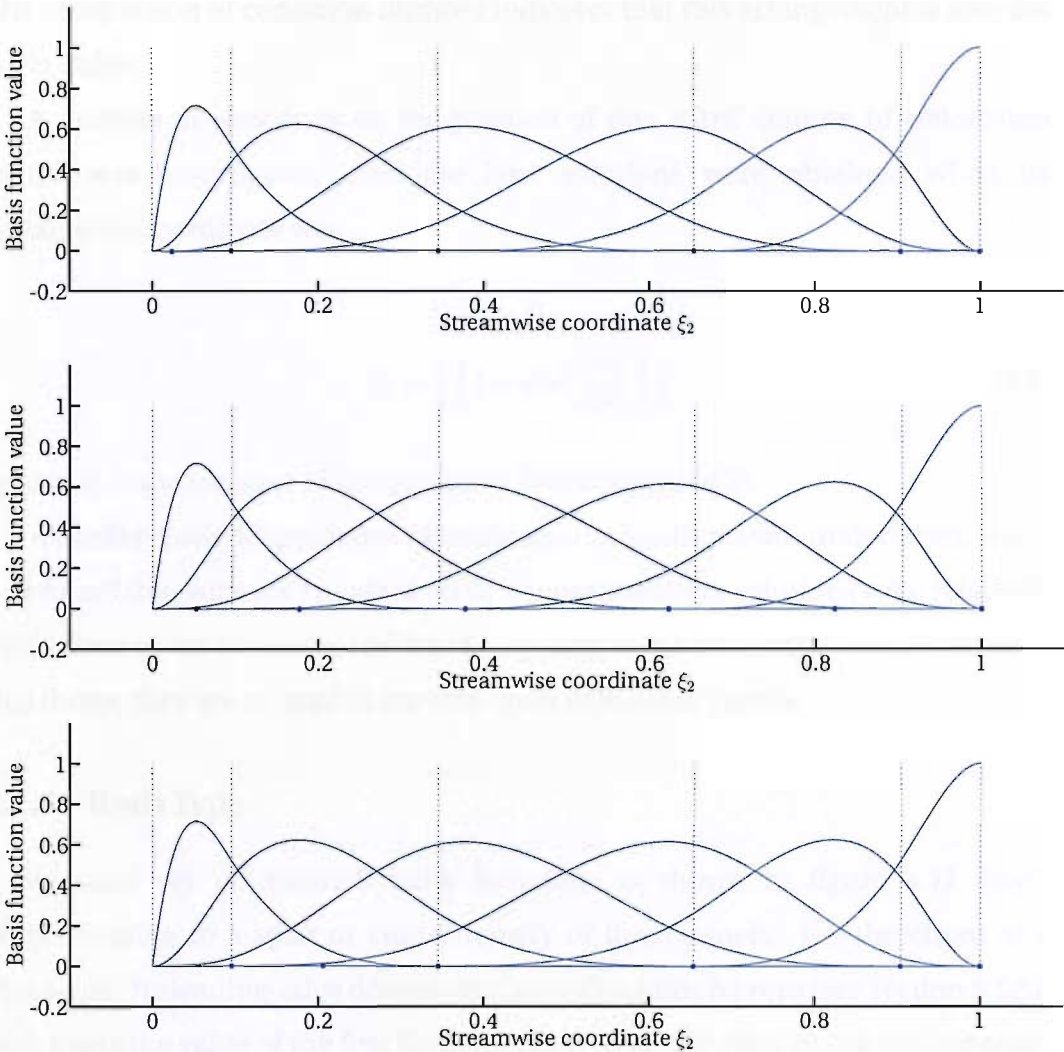


Figure 4.9: Alternative collocation point schemes, where the positions of collocation points are indicated by dots on the axes. Top: On knots, Middle: On basis-function maxima, Bottom: On knots with an offset first point. A discussion of the basis function shapes follows in section 4.1.5.

the leading edge. As seen in figure 4.9 (top-right), the doublet solution is slightly 'kinked' by the collocation point, and the position of the kink is found to move with the position of the collocation point.

In the second alternative arrangement, the column of points within the leading edge panels is moved to lie within the second column of panels, as shown in figure 4.9 (bottom). This is seen to produce a smooth doublet surface (figure 4.10, top-right), whilst maintaining the rapid convergence of lift and drag.

The comparison of condition number indicates that this arrangement is also the most stable.

A number of variations on the position of this ‘extra’ column of collocation points was investigated, and the best solutions were obtained when its streamwise coordinate was

$$\xi_2 = \frac{1}{2} \left\{ 1 - \cos\left(\frac{3\pi}{2N_i}\right) \right\} \quad (4.3)$$

where N_i is the number of unique knots (see equation 4.2).

A similar study of spanwise collocation point positions was undertaken, and it was found that with the standard on-knot configuration, solutions were relatively insensitive to the placement of the extra points at the two ends. For the results in this thesis, they are located at the mid-span of the side panels.

4.1.5 Basis Type

A standard set of B-spline basis functions is shown in figure 4.11 (top), corresponding to a span or chord density of three panels. For the chord of a thin body, the leading-edge doublet surface value must be zero (see section 3.1.5), which sets the value of the first basis function weight to zero. At the trailing edge, the doublet strength gradient must be zero (see section 3.1.6), constraining the last two basis function weights to be equal. The first basis function can therefore be discarded, and the last two combined, as shown in figure 4.11 (bottom).

In the span, the basis functions at both (non-wake) edges can also be discarded to leave the set shown in figure 4.12. Reducing the number of basis functions by two in both the chord and span produces significant gains in computational efficiency.

In section 3.3.2, the possible benefits of modifying the standard set of basis functions was discussed. In particular, the use of a basis function with a square-root characteristic at the leading edge might be well-suited to thin body problems, where it is known that the true doublet distribution will approximate this curve. One possible advantage of using such characteristics is that the leading-edge

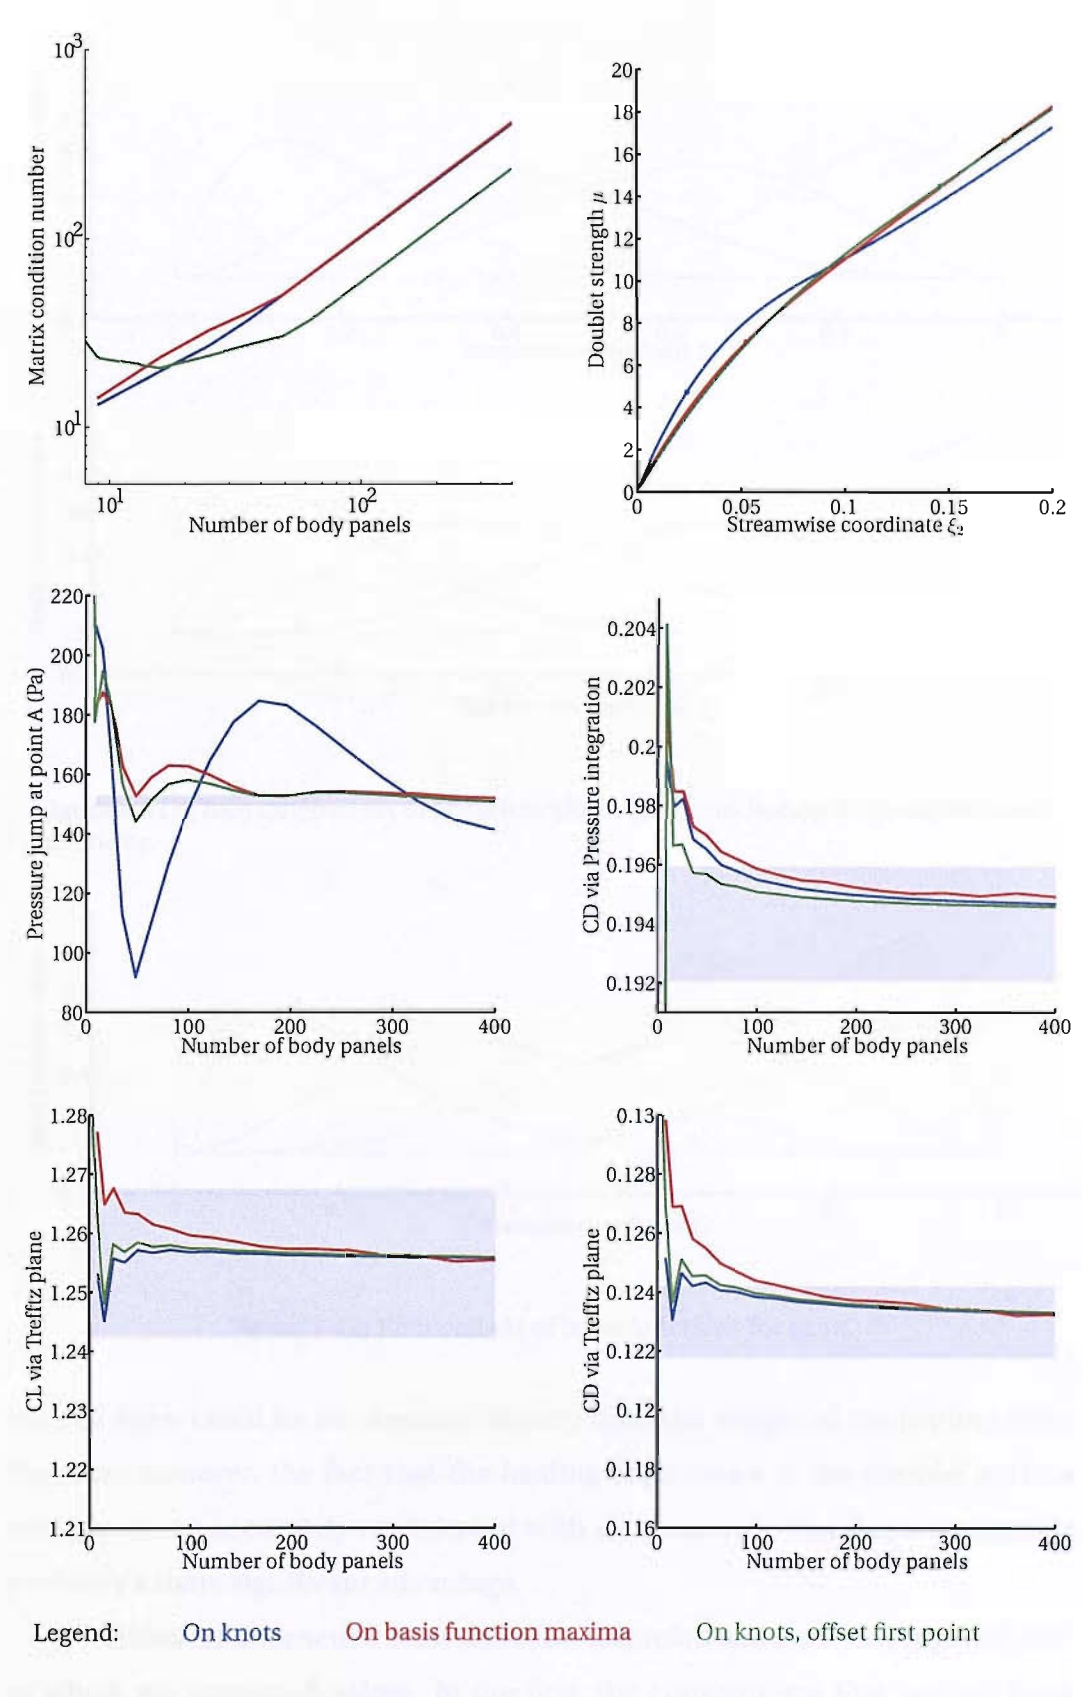


Figure 4.10: Influence of collocation point positions on convergence





Figure 4.11: (Top) Original set of basis functions. (Bottom) Reduced chordwise basis functions.




Figure 4.12: Reduced set of basis functions for span.

suction force could be determined directly from the weight of the leading basis function; however, the fact that the leading edge shape of the doublet surface could be more accurately represented with a square-root type basis function is probably a more significant advantage.

A number of alternative basis function formulations were investigated, two of which are presented below. In the first, the conventional first (active) basis

function $U_2(\xi)$ is replaced by a basis function $U'_2(\xi)$ that spans the same two panels, but has a leading-edge square-root characteristic

$$U'_2(\xi) = \sqrt{\xi} \{U_1(\xi) + U_2(\xi)\} \quad (4.4)$$

In the second formulation, the standard basis function $U_2(\xi)$ is maintained, but an additional basis function $U'_1(\xi)$ is introduced at the leading edge, defined by

$$U'_1(\xi) = \sqrt{\xi}U_1(\xi) \quad (4.5)$$

In this case, the term $U_1(\xi)$ is the B-spline basis function that was previously discarded because it had a finite leading-edge value. The standard and modified basis function sets are shown in figure 4.13.

The results presented in figure 4.14 show that the scheme that replaces the existing first basis function has inferior convergence rates in every respect. It is thought that whilst the doublet surface shape at the very leading edge may be improved, the shape further downstream is adversely affected by the modified basis function.

The second scheme, which introduces a new basis function, in general produces similar results to the standard scheme. The exception is in the pressure integration of drag, where the second scheme has more rapid convergence. A significant disadvantage of this scheme, however, is that for a given panel grid the number of collocation points and basis functions is increased, thus raising overall computation time.

Square-root type basis function weights

If basis-functions with square-root characteristics at the leading edge are used, the weight assigned to the first basis function should be directly related to the magnitude of the leading edge force. In figure 4.15 (left), this direct method is compared to the fitting method for the two types of modified basis function, and it is seen that it does not produce good estimates for suction forces. The reason for this is shown in the right-hand panel of the figure. Although the doublet curve at

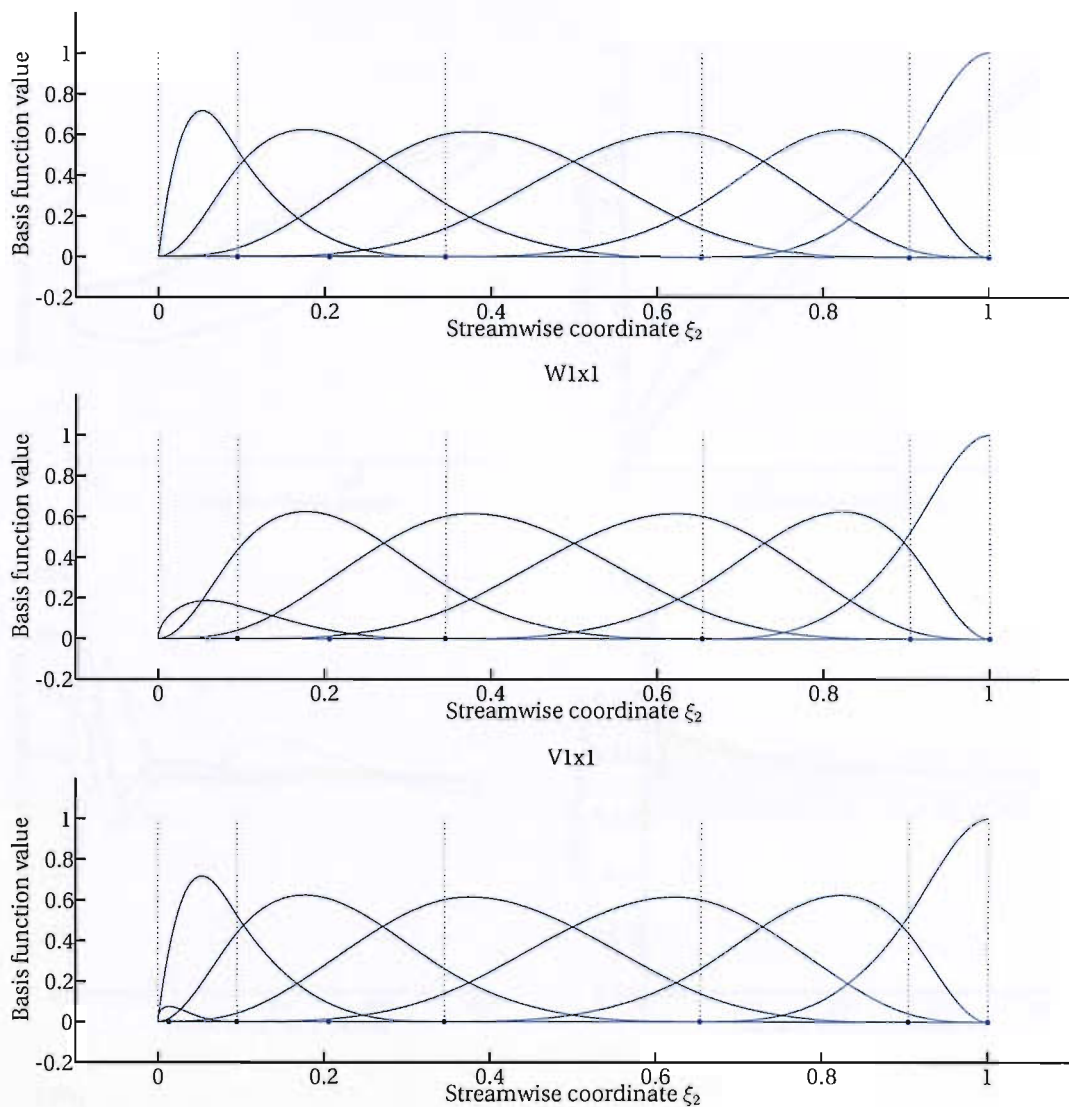


Figure 4.13: Basis function sets for a five panel solution. (Top) Original basis functions. (Middle) First active basis function replaced by similar function with leading-edge square-root characteristic. (Bottom) Additional basis function. The fact that the new basis functions have a relatively low magnitude resulting from the square-root factor is not significant, since they will be scaled as required by the solution weights.

the very leading edge is dominated by the first (square-root type) basis function, by the time the first collocation point is reached (shown by the first blue dot), the second and third basis functions contribute significantly to the overall doublet strength; hence, within the limited resolution provided by the discrete collocation

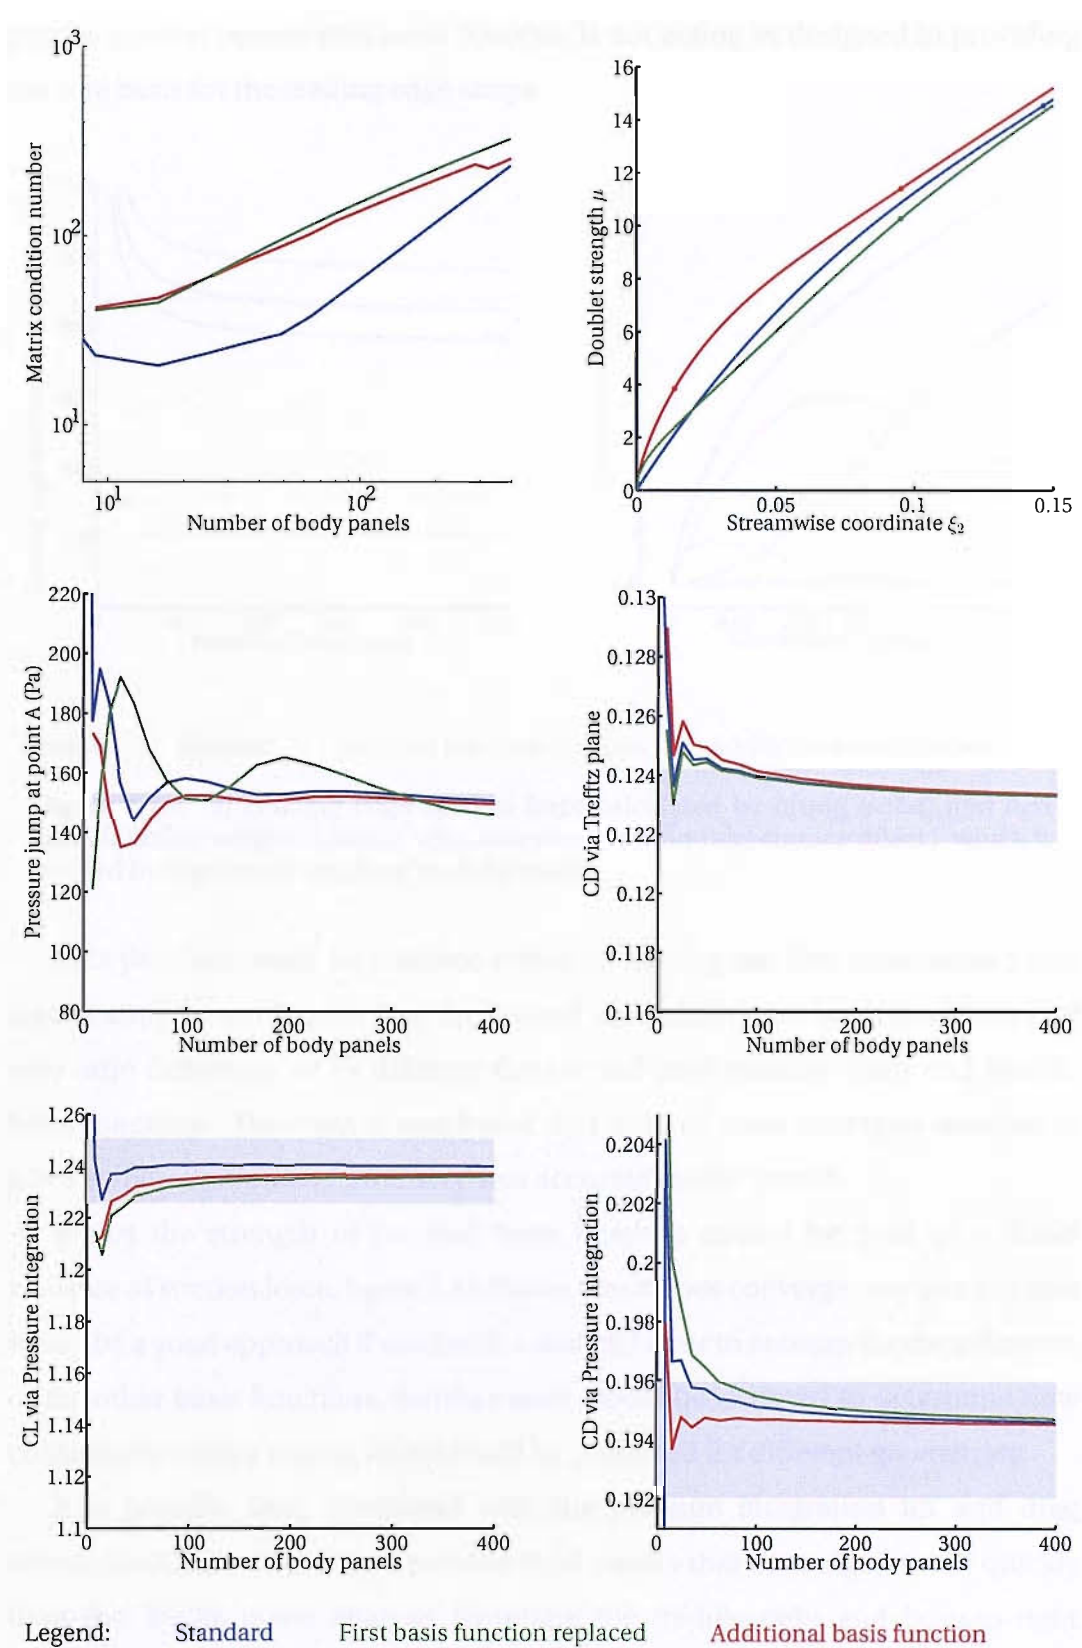



Figure 4.14: Influence of modified basis functions on convergence

points, the first square-root basis function is not acting as designed in providing the sole basis for the leading edge shape.



Legend: Standard Modified first basis function Additional basis function

Figure 4.15: (a) Leading edge suction force calculated by fitting (solid) and first basis function weight (dashed). (b) Cross-section of doublet surface (black), which is formed by the sum of weighted basis functions.

This problem could be resolved either by moving the first collocation point significantly forward, such that the second and subsequent basis functions had very little influence, or by deleting the second (and possibly third and fourth) basis functions. However, it was found that both of these strategies resulted in a less stable solution and produced less accurate results overall.

Whilst the strength of the first basis function cannot be used as a direct measure of suction force, figure 4.15 shows that it does converge very quickly, and it may be a good approach if used with a scaling factor to account for the influence of the other basis functions. Further work would be required to determine how consistently such a scaling factor could be predicted for different geometries.

It is possible that, combined with the pressure integration lift and drag calculations, this would then provide total results that converged more quickly than the Trefftz plane analysis (compare the middle-right and bottom-right graphs of figure 4.14). However, the disadvantage of increasing the size of the solution described above probably outweighs the benefits seen in using the

modified basis function scheme, and a the standard basis function set is used in the remainder of the thesis.

4.2 Optimal configuration

The graphs presented in this section are the result of over 500 simulations used to investigate a number of configuration options. These results in fact represent only a ‘short-list’ of configurations, drawn from a wider field of options that was investigated. For the particular sail body under investigation, it was found that an efficient, accurate and stable configuration was defined by:

- A cosine-cosine panel grid spacing;
- An $(N \times N)$ panel density;
- Collocation points on spline knots with an additional column of points within the second panel from the leading-edge, and additional rows of points at the mid-span of the side-edge panels;
- A standard set of B-spline basis functions.

It is expected that many of the conclusions on panel and collocation point configurations drawn in this study will apply in general to other bodies, but as with conventional potential flow methods, it would be prudent to carry out some form of parameter study with any new type of geometry to confirm its applicability and efficiency.

4.3 Computation time

Figure 4.16 shows how computation time increases with the number of body panels in the configuration. With a conventional panel method the run-time order is approximately constant at $O(N^2)$, since the number of influence calculations is equal to the number of panels multiplied by the same number of collocation points (see Bernasconi, 2007). The run-time order for the CSM is less straightforward to predict, as it depends on how many panels are sub-divided,

and how many integration points are required in each sub-panel. A fit line of order $N^{1.9}$ shows that for a typical number of panels the CSM scales similarly to a conventional panel code, but at very high panel numbers this order is seen to increase.

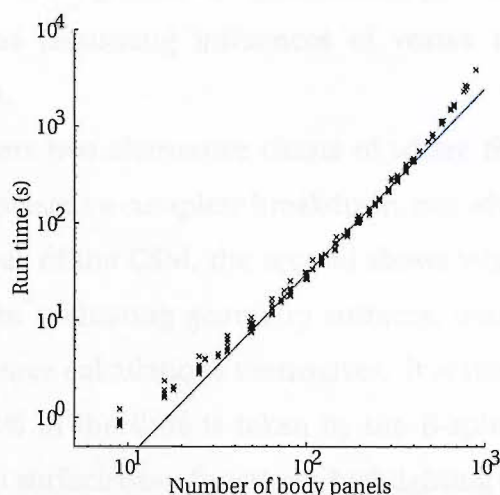


Figure 4.16: Variation of run-time with number of panels. The fit line is of the form $t \propto N^{1.9}$.

Although it is not the aim of this thesis to describe the implementation of the CSM in detail, it is interesting to outline where the run-time is spent for a typical thin-body solution with a grid of (5×5) body panels, with the same body and wake geometry as described in section 4.1. For this solution:

- The base configuration contains 3 patches: the body contains 25 panels; the initial wake contains 25 panels, and a long straight wake extension contains 5 panels.
- Following subdivision, there are a total of 215 panels and sub-panels in the body, 129 in the wake, and 63 in the wake extension. This includes every level of the hierarchy, such that panels are counted in addition to their sub-panels. Figure 4.17 illustrates the sub-panel divisions in the body and main wake.
- For the standard, non-singular influences 9,425 integration points are required.

- For the singular influences 36,864 surface (two dimensional) integration points are required, with 9,216 line (one dimensional) integration points.

In comparison, a (5×5) vortex lattice grid with a similar wake and wake extension would have 55 panels, 25 collocation points, and require 3,150 influence calculations (assuming influences of vortex lines are re-used for neighbouring panels).

Figure 4.18 presents two alternative charts of where the overall run-time is spent. Both charts represent a complete breakdown, but where the first is divided by the functional areas of the CSM, the second shows what proportion of time is spent specifically in evaluating geometry surfaces, doublet surfaces, and in performing the influence calculations themselves. It is interesting to see in this second chart that 50% of the time is taken by the B-spline evaluations of the geometry and doublet surface basis functions, highlighting the need for this small section of code to be optimised as much as possible. In a conventional panel code, neither of these sections are necessary (except possibly in initialising the geometry). Once the location of the integration points has been determined, and their basis weights evaluated, actual computation of the regular or singular velocity influences accounts for less than 10% of overall run-time. The significant proportion of effort taken in general program logic and data handling would be significantly reduced if the code were to be re-written in a compiled language, rather than the interpreted one (MATLAB) used in this case.

In the top chart, it is seen that slightly more time is spent in singular influence coefficient calculations than in standard influence calculations, even when the panel-subdivision process is included with the latter. Although for a given collocation point there are far more regular influences than singular ones, it is reasonable that more run-time is spent in evaluating singular influences, as these will have a stronger influence on ultimate solution accuracy.

4.4 Validation of the CSM

The optimal configuration established in section 4.2 is applied here to the case an inclined rectangular flat plate of aspect ratio 2.0, combined with a semi-infinite




Figure 4.17: Panels (thick lines) and sub-panels (thin lines) in a (5×5) grid. The body is shown in black and the wake in grey. The long wake extension patch is not shown.

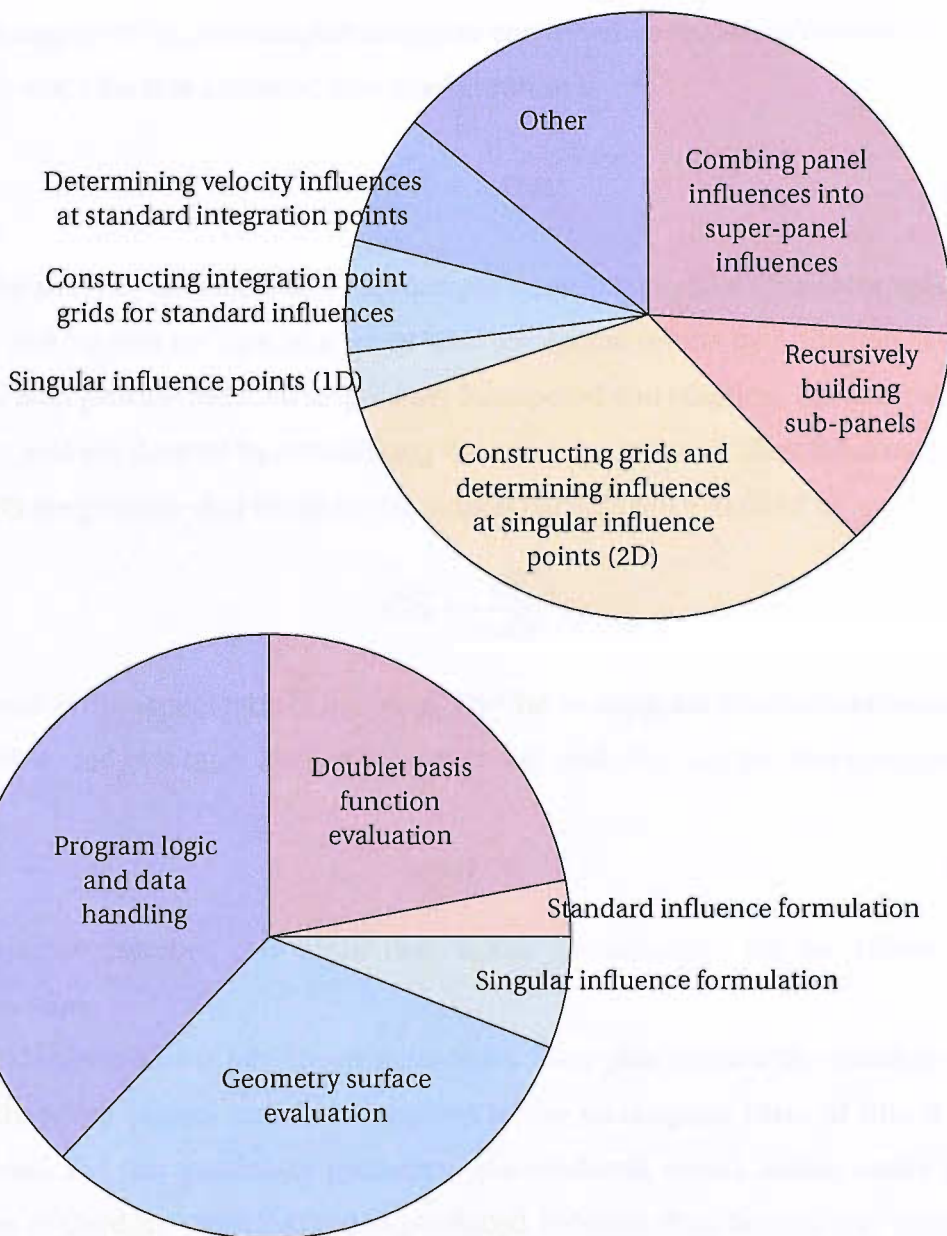


Figure 4.18: Breakdown of overall run-time for (5×5) panel solution. (Top) Breakdown by functional areas. (Bottom) Breakdown by lower-level code area, showing that 50% of time is spend in evaluation of B-spline coefficients for geometry and doublet surfaces.

undeformed flat wake aligned with the free-stream. This enables a validation of the CSM to be made against one of the few thin-surface geometries for which an almost exact solution is known. This solution is provided by Labrujere & Zandbergen (1973), who employ extensive convergence studies to determine that the lift slope for this potential flow configuration is

$$\frac{dC_L}{d\alpha} = 2.4744 \quad (4.6)$$

The study by Labrujere and Zandbergen does not consider figures for induced drag, and for this we turn to a set of semi-analytical results by Anderson (1936). These data provide induced drag values for tapered and elliptical, flat and twisted wings, and are derived by considering the wing as a series of chordwise sections. Results are presented in terms of the induced drag factor e defined by

$$C_{D_i} = \frac{C_L^2}{\pi \mathcal{R} e} \quad (4.7)$$

where \mathcal{R} is the aspect ratio of the wing. The flat rectangular plate considered here is represented by a taper ratio of 1.0, zero twist, and $\mathcal{R} = 2.0$, for which Anderson gives

$$e = 1.00 \quad (4.8)$$

This factor matches the value that would be obtained for an elliptic lift distribution.

Anderson's results have been derived for wing planforms with rounded tips, and therefore cannot strictly be applied to the rectangular plate of this study. However, for this particular geometry, conventional vortex lattice codes (e.g. Fiddes & Gaydon, 1996) have also produced induced drag factors very close to $e = 1.00$, supporting the use of this value a benchmark.

The results of the CSM convergence study are plotted in figure 4.19. Lift and induced drag, calculated both in the Trefftz plane and by pressure integration, are plotted against the number of panels in the plate body.

It is seen that both methods of lift calculation converge at the same rate, but to slightly different asymptotic values. This difference is a result of the leading-edge

suction force contribution, which is included in the Trefftz plane calculation, but not in the pressure integration. Although this force could have been estimated (see section 4.1.1), the difficulties in determining it accurately would obscure the overall results.




Figure 4.19: Lift and induced drag for the flat rectangular plate calculated in the Trefftz plane and via pressure integration. The benchmark results of Labrujere & Zandbergen, and Anderson, are shown as dashed black lines.

In contrast to the lift results, the Trefftz plane induced drag results are very different from those obtained via pressure integration. Since the inflow angle for the solution is small, the leading edge suction force, acting in the plane of the wing, contributes much more to (negative) drag than it does to lift.

Whilst the pressure distributions should not be used to determine overall forces, lift and induced drag calculated in the Trefftz plane converge well towards the results of Labrujere & Zandbergen, and Anderson. For this geometry, the validity of the CSM method is thereby confirmed within the accuracy of the available comparisons. Further validation, for a more complex geometry, is provided in section 4.5.

4.5 Comparison of CSM and Vortex Lattice results

An important test of the CSM is to compare its performance with that of a standard vortex lattice method, and such a comparison is presented in this section. The test problem on which the two methods are compared is the irregular curved sail described in section 4.1, with a pre-defined wake shape. The CSM configurations that were identified by the parameter study to give the best overall performance are used in the comparison, thus presenting the method working at its best. Specifically, the panels are configured on $(N \times N)$ cosine-cosine grids, using standard B-spline bi-cubic basis functions (discarding/merging the edge basis functions as appropriate, but not using any square-root characteristics). Collocation points are also cosine-spaced, with the optimal arrangement of additional points as described in section 4.1.4.

In order to provide a fair comparison, a similar parameter study was carried out for the VLM to identify the optimum configuration for the sail body. This exercise, the details of which form appendix C, found very similar results to the CSM study. The best arrangement of VLM panels is also configured on a $(N \times N)$ cosine-cosine grid, with cosine-cosine spaced collocation points. It was found that projecting collocation points on to the bi-linear geometry of the quads bounded by vortex lines produced significantly better results than setting them on the curved surface of the original sail shape.

The Vortex Lattice is a well-established potential flow method, and the particular implementation used as a benchmark here (Bernasconi, 2007) has itself been validated against the Labrujere & Zandbergen (1973) case described in section 4.4. The VLM and the CSM were both written in MATLAB v7, and run on the same desktop PC (3.0GHz processor, 1Gb RAM). Although in a commercial environment it is likely that a potential flow code would be written in a compiled language, it is considered that run-times for the two methods would be reduced by similar factors if compiled. In the author's experience, a numerical application of this type is likely to run around ten times quicker if written in C, than if written in MATLAB, assuming both are efficiently coded.

Both the VLM and the CSM were optimised to a reasonable extent. However, where there is considerable scope for further efficiency gains with the CSM, there are few areas left for improvement with the much simpler and shorter code of the VLM.

4.5.1 Basis for comparison

Although it has been convenient in previous sections to refer to particular CSM and VLM solutions by the number of panels that comprise the body, this is not a very useful basis on which to compare the two methods. As the panels of the CSM are higher-order than those of the VLM, the number of unknowns and the computational cost of the solution will always be higher for the same number of panels.

An end-user of either method will generally require the most accurate solution possible within the constraints of either computation time or available memory. With an average PC, computation time is more likely to be the driving factor, but with very large solutions matrix sizes may become limiting. The solution matrices of both the VLM and CSM are square with $(N \times N)$ elements, where N is the number of unknowns in the problem. Results are therefore compared against axes of both computation-time and N .

4.5.2 Run time and condition number

Figure 4.20 (left) compares the run-time for the two methods. For the same number of unknowns, the CSM is in the region of 50 times slower. The run-times of both methods are approximately proportional to the square of the number of unknowns, although at very low panel numbers the general overheads in setting up the problem, etc., influence the VLM's run-times.

Figure 4.20 (right) shows that the condition number of the CSM solution matrices are generally lower than those corresponding to the VLM, indicating that the VLM would be slightly more sensitive to numerical errors. However, the difference between the two methods in this respect is not very significant.

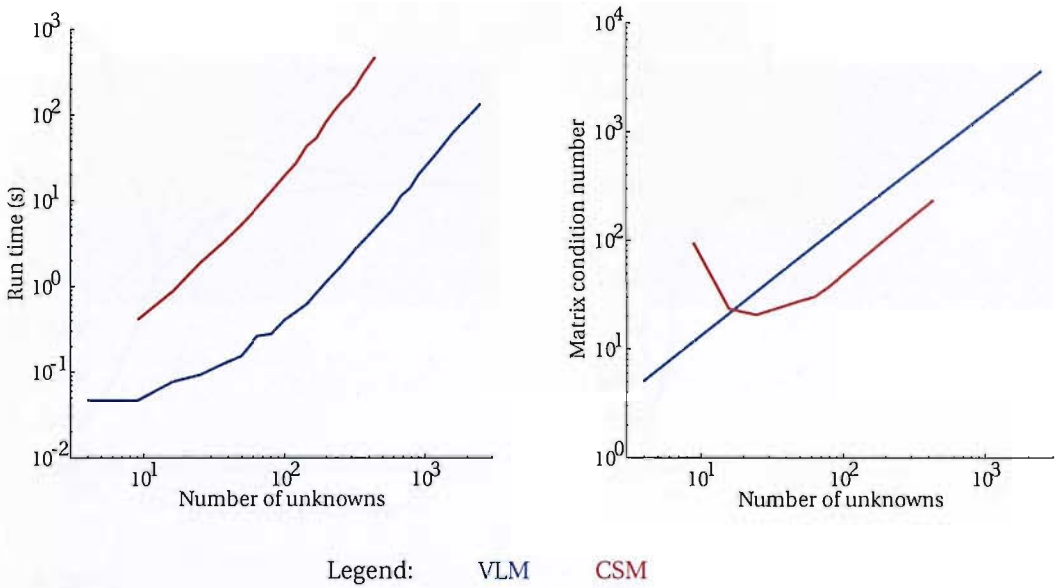


Figure 4.20: Comparison of (Left) run-time, and (Right) matrix condition numbers

4.5.3 Lift and drag coefficients

Figure 4.21 compares the lift and induced drag calculated by a Trefftz plane analysis for both methods. Most importantly, both methods converge to almost identical solutions, confirming the ability of the CSM to predict these fundamental aerodynamic coefficients. Both methods provide solutions within a 1% error band with very small panels, and converge towards the assumed solution at a similar rate with respect to the number of unknowns. However, as the computational cost per panel is so much lower with the VLM, it provides lift and drag results around ten times more quickly than the CSM, for a comparable accuracy.

Lift and induced drag can also be calculated by direct pressure integration, and comparative results are shown in figure 4.22. Once again it is clear that both methods produce solutions which converge to the same value. However, unlike the Trefftz plane analysis, the CSM lift value converges much more quickly with number of unknowns than the VLM value, and even on a run-time basis the CSM is considerably more efficient.

As pressure integration produces an overall force magnitude and direction which is subsequently resolved into lift and drag components, it might be

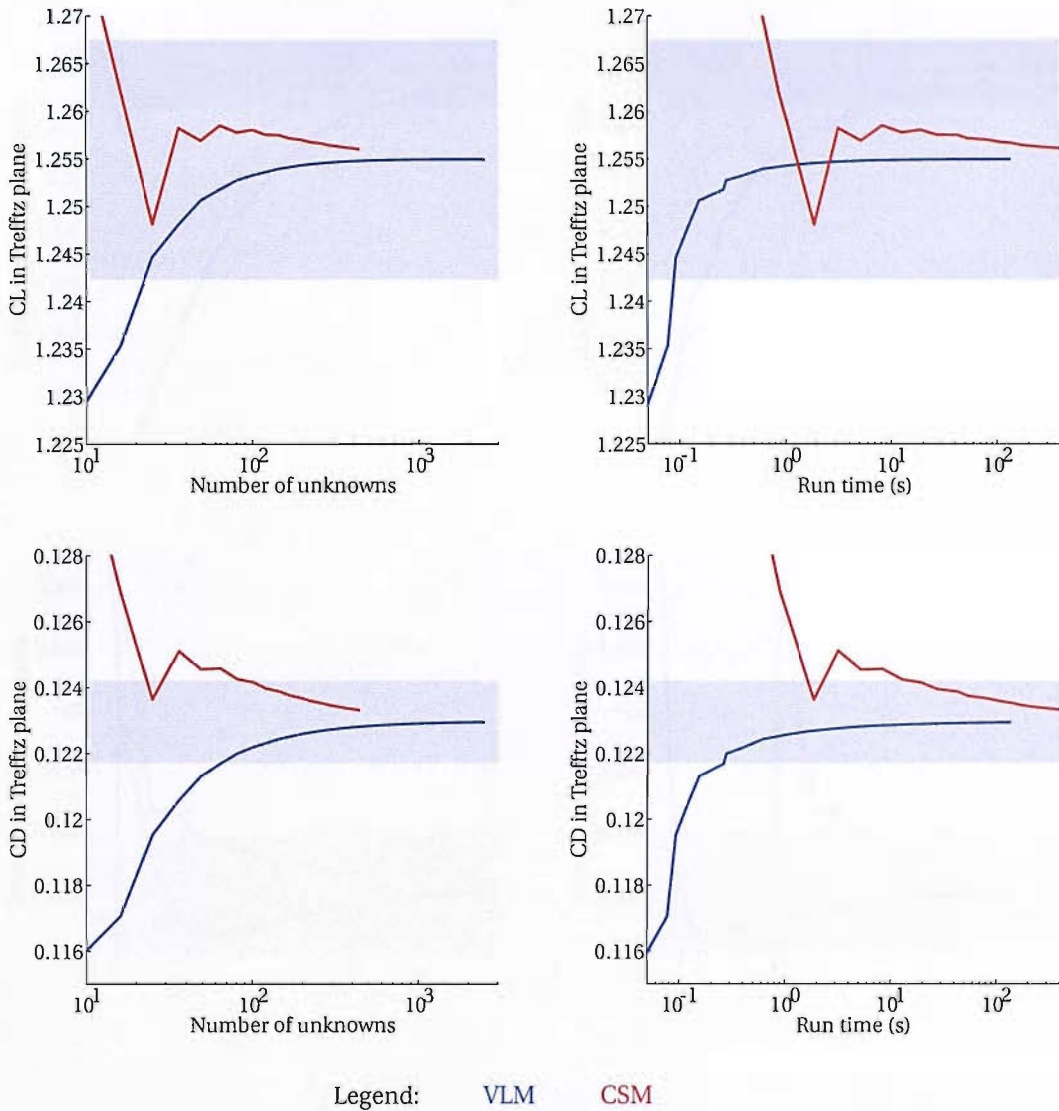


Figure 4.21: Comparison of CL and CD calculated in the Trefftz plane

expected that the drag results would exhibit similar convergence rates to the lift results. However, this is not the case, and where the CSM's drag is slower to converge than its lift, the VLM's drag is much quicker than its lift.

One significant difference between the pressure integration methods is that the forces integrated over the CSM panels are all pointing in exactly the right direction, normal to the local surface, whereas those in the VLM are normal to the quadrilateral panels that approximate the body's geometry. It is possible the errors in force direction caused by the geometric approximation are partly

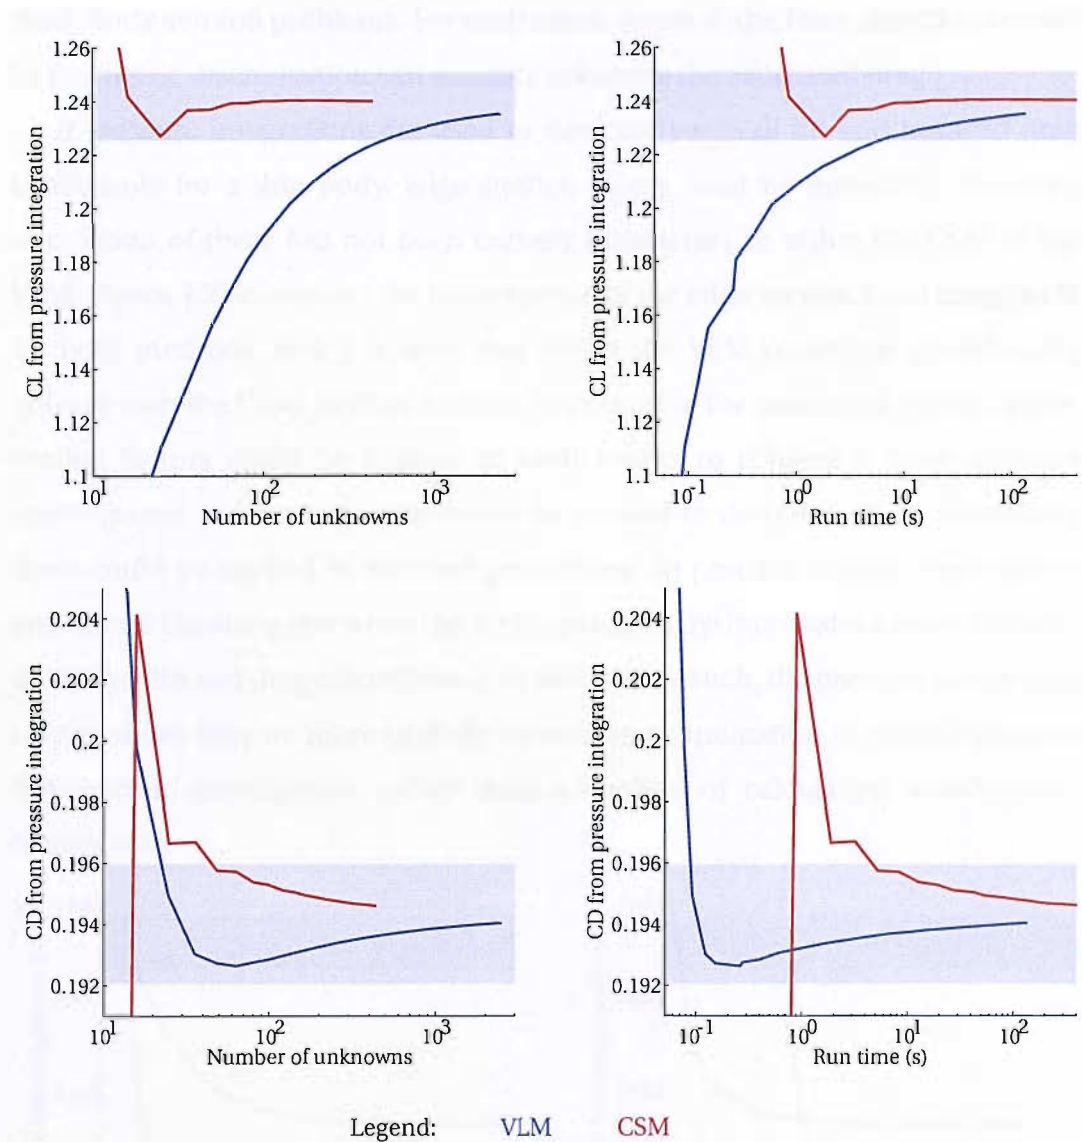


Figure 4.22: Comparison of CL and CD calculated by pressure integration

responsible for the slow convergence of the VLM's lift. A detailed inspection of the CSM's results did not reveal a clear explanation of the lift converging more quickly than the drag.

[It is interesting to note that whilst pressure integration is a notoriously inaccurate method of determining induced drag coefficients in a constant panel method, this VLM test has produced very good drag results. This is partly because with this thin body there is no high-curvature leading edge, as is seen in many

thick-body aerofoil problems. For such cases, errors in the force direction caused by geometric discretisation can strongly influence the estimated drag.]

If pressure integrations are used to determine overall lift and induced drag coefficients for a thin body, edge suction forces must be included. However, calculation of these has not been entirely satisfactory in either the CSM or the VLM. Figure 4.23 compares the convergence of the edge suction force magnitude for both methods, and it is seen that whilst the VLM converges considerably quicker than the CSM, neither method converges to the (assumed) correct value. Scaling factors could be applied to both results to achieve a more accurate convergence, but further work would be needed to determine how universally these could be applied to different geometries. In practice though, there seems little benefit in doing this when the Trefftz plane analysis provides a more accurate method of lift and drag calculation in both cases. As such, the pressure integration results above may be more usefully viewed as an indication of overall pressure distribution convergence, rather than a method of calculating aerodynamic coefficients.




Figure 4.23: Comparison of edge suction force magnitude

4.5.4 Surface pressures

Figure 4.24 presents a comparison of surface pressure differences evaluated at two discrete points on the sail body. The two points, whose locations are shown in figure 4.1, were chosen to be representative of regions of slowly varying pressure (point A), and quickly varying pressure (point B). In both cases, the CSM requires far fewer unknowns than the VLM to converge to a solution within 1% of the (assumed) correct value. For point B, at the centre of the sail, the run-time is also significantly less. Point A, close to the leading edge of the sail presents more of a problem to convergence, particularly for the CSM. Whilst the CSM still converges with a lower number of unknowns, it approximately matches the VLM in run-time.

It is interesting to note that the VLM is significantly slower to converge for these two discrete points than it was in converging to a net pressure difference for the whole sail, indicating that in general some local VLM errors cancel out when integrating results over the body. In particular, the saw-tooth shaped curves of the VLM on point A suggest that odd and even panel numbers affect the sample point differently (this central point is in the centre of a panel for odd-numbered grids, and at a vertex between four panels for even-numbered grids).

4.5.5 On-surface velocities

The boundary conditions imposed at collocation points require that the flow velocity normal to the surface is zero. As the VLM and CSM both produce square systems of equations, these collocation point boundary conditions will be satisfied exactly. Over the remainder of the surface, between collocation points, the zero penetration condition will generally not hold.

Figure 4.25 shows how the component of velocity normal to the curved surface varies along a section through the sail. The velocities produced by the VLM only satisfy the zero normal flow condition at the collocation points, and are singular at the vortex lines which divide the panels. In contrast, the CSM produces velocities which are approximately correct throughout, except at the leading edge of the body. At this edge, the true doublet distribution which is known to have

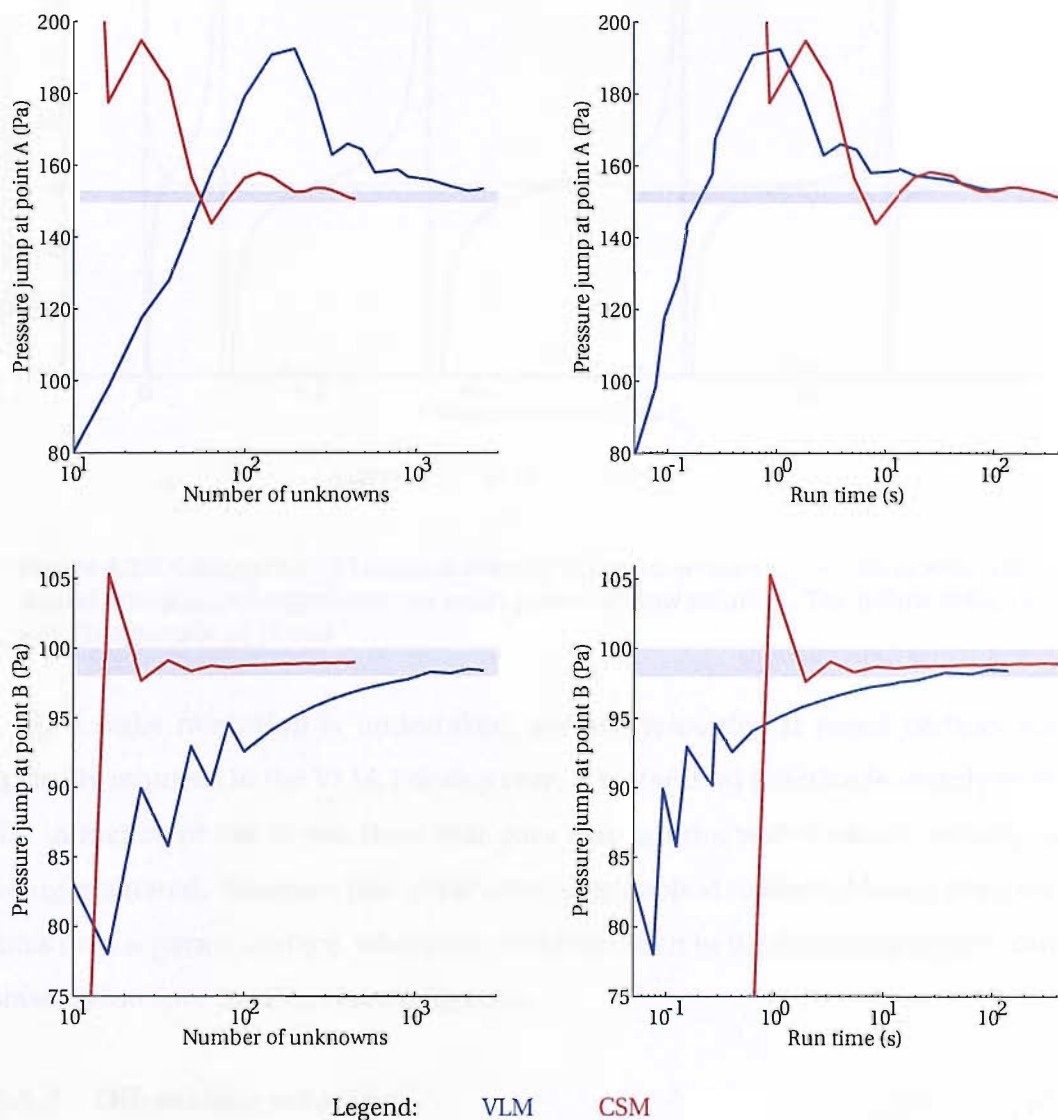


Figure 4.24: Comparison of pressure jumps at sample points A and B

a square-root characteristic is not accurately modelled by the standard B-spline basis functions, resulting in the singular velocity seen.

Whilst the graph usefully illustrates a fundamental difference of the two methods, the singular nature of the VLM surface velocity characteristic does not generally cause any problems with the method if the wake shape is fixed. The normal velocity component is not a quantity that is often required in post-processing, as it is assumed to be zero throughout; and if other components are required, they can be interpolated between values at collocation points.




Figure 4.25: Comparison of normal velocity on cross-section ($\xi_1 = 0.5$) of sail. The dotted line at $q_n = 0$ represents the exact potential flow solution. The inflow velocity has a magnitude of 10 m/s

If a wake relaxation is undertaken, surface velocities at panel vertices are typically required in the VLM. For this case, it is standard practice to simply omit the influence of the vortex lines that pass through the vertex whose velocity is being evaluated. However, this ‘trick’ cannot be applied to the influence of vortex lines on a separate surface, where the problems seen in the following section can severely hamper a stable wake relaxation.

4.5.6 Off-surface velocities

Figure 4.26 shows a comparison of off-surface streamlines constructed for the VLM (top) and CSM (bottom) solutions. The streamlines start from the same place in both cases, spaced evenly along a line approximately perpendicular to the leading edge of the body, and to the inflow direction. In both cases, the streamlines were constructed using a variable-step Runge-Kutta (4, 5) formula (Dormand & Prince, 1980), where sample-point velocities $q(t, y) = q(y)$ were evaluated in the normal way.

Looking first at the outermost windward and leeward streamlines, it is seen that they follow approximately similar paths for both the VLM and CSM solutions. These streamlines flow about 1 m from the sail’s surface, which is far enough from





Figure 4.26: Comparison of streamlines from the VLM (top) with a (10×10) panel grid, and the CSM (bottom) with a (3×3) panel grid.

the body to be relatively unaffected by the discretisation of the grids (the chord length of the sail is approximately 8 m at this span).

The streamlines closer to the sail's surface flow smoothly over the CSM solution, and as they leave the body are distributed evenly into two groups, one on the windward side and one on the leeward side. This is the correct behaviour, and has been achieved with a low resolution (3×3) panel grid. The only visible error is that one streamline slightly penetrates the leading edge of the sail, from the windward side to the leeward side. This corresponds to the error in on-surface velocity seen in figure 4.25.

In comparison, the corresponding streamlines over the VLM solution are highly irregular, and oscillate across the sail surface as they are affected by each vortex line comprising the panel boundaries. In addition, two of the windward side streamlines begin to form a vortex roll-up around one of the chordwise panel boundaries. Downstream of the sail, there is no clear division between windward and leeward streamlines, and they are unevenly distributed. Some alternate starting points to those shown also resulted in some streamlines being 'captured' by a vortex line, as illustrated in figure 4.27. We can conclude that streamlines can only be reliably generated for the VLM at some distance (at least one panel length) from the body.

4.5.7 Summary of VLM and CSM comparison

Numerical results from both methods have shown that the CSM requires fewer unknowns, or a smaller solution matrix, to achieve results of a given accuracy for the most commonly required post-processed quantities. This includes total aerodynamic coefficients evaluated in the Trefftz plane, total pressure differences integrated over the body surfaces, specific local surface pressures, and field-point velocities. The only exception is edge suction forces, for which both the VLM and CSM perform quite poorly.

However for most applications, more important than the a comparison by matrix size is a comparison by run-time. For a given matrix size, the VLM is around 50 times quicker than the current implementation of the CSM, and this




Figure 4.27: A streamline ‘captured’ by a vortex line with the VLM.

makes the two methods much closer competitors when compared on run-time basis.

It was found that total lift coefficients, when calculated in the Trefftz plane, could be calculated significantly more quickly by the VLM than the CSM. This shows that the high-order panels of the CSM have not drastically increased the accuracy of the spanwise distribution of vorticity modelled in the solution. It is in the more local results that the CSM has the advantage over the VLM, where the ability to directly evaluate surface pressures and velocities, without the need for interpolation, leads to quicker convergence for these quantities.

If on-surface velocities or streamlines are required, the continuous solutions of the CSM produce much better results than those of the VLM, even at very low panel densities. As will be shown in chapter 5, it is the ability of the CSM to accurately predict on- and off-surface velocities that should give it an advantage in generating stable wake relaxations.

Chapter 5

Wake relaxation

5.1 The requirement for relaxing the wake

The wake shed from a lifting body is a shear layer modelled as a thin geometric surface containing a doublet distribution. As described in section 3.1.6, the doublet distribution is dependent on that of the lifting body, and for a steady-state case is constant in the streamwise direction. Whilst it is often acceptable to prescribe an approximate wake shape for a particular case, ideally the wake will be allowed to ‘relax’ to assume a geometry parallel to the flow downstream of the body.

The geometry of the wake usually has a small but significant influence on the potential solution for the upstream body, and a larger influence on the flow pattern in the vicinity of the wake itself. As such, wake relaxation schemes within potential flow models are inevitably iterative: typically an initial wake shape is assumed, and a first approximation for a doublet solution is found; a new wake shape is defined by tracing streamlines from the trailing edge of the body, and this shape is then used in the following iteration. In a successful wake relaxation, the doublet solution and wake shape will both converge over a number of iterations.

In this chapter, a wake relaxation scheme applicable to smooth geometric surfaces is introduced. The scheme is demonstrated for the simple case of flow over a flat inclined plate, and the more complex example of flow over two closely overlapping sails. Some problems with the basic scheme are identified, and more complex strategies to deal with those problems are then discussed.

5.2 Wake relaxation principles for the CSM

5.2.1 Wake sample points

The wake relaxation strategies adopted in the CSM are based on representing the wake geometry as one or more connected B-spline patches. These patches are updated by considering the movement of a grid of points lying within the wake over a number of iterations. The grid of *wake sample points* $\mathbf{v}_{f,g}$, with intrinsic coordinates $\xi_{f,g}$, are designated by spanwise index f and streamwise index g . The points lie at knot intersections of the B-spline geometric grid, with additional rows of points at the mid-span of edge panels, as shown in figure 5.1.

5.2.2 Staged relaxation process

Within each relaxation iteration, the geometry is updated in stages, initially fixing the shape of the wake closest to the lifting body, then progressing downstream. At each stage, the geometry is updated by considering the movement of one

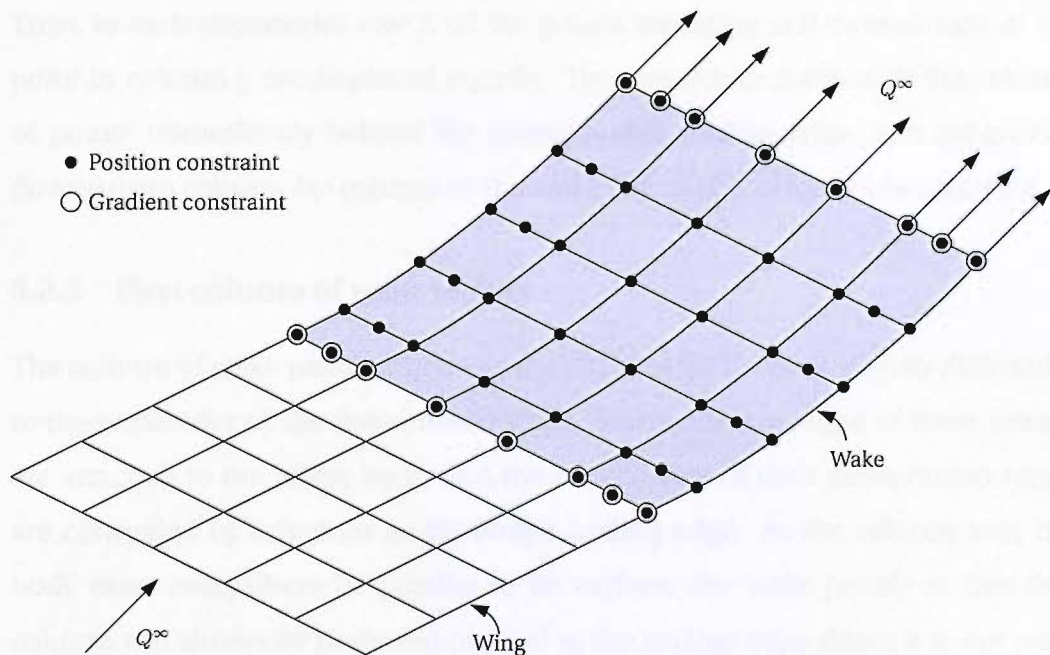


Figure 5.1: Constraints used to define a wake geometry in a relaxation step. Note the short first column of wake panels aligned with the lifting body.

spanwise column g of points. As the wake relaxes, a new position $\mathbf{v}'_{f,g}$ for each point is determined by evaluating the velocity $\mathbf{q}_{f,g-1}$ at the point's upstream neighbour, then projecting downstream from that upstream point's position $\mathbf{v}_{f,g-1}$ in the direction of the velocity. The distance projected is equal to the initial distance between the two points, thus maintaining a constant spacing between grid points throughout the relaxation. This is effectively a first-order Euler integration step, as follows:

$$\mathbf{v}'_{f,g} = \mathbf{v}_{f,g-1} + \frac{\mathbf{q}_{f,g-1}}{|\mathbf{q}_{f,g-1}|} |\mathbf{v}_{f,g}^0 - \mathbf{v}_{f,g-1}^0| \quad (5.1)$$

The displacement of the point $\mathbf{v}_{f,g}$ is therefore

$$\begin{aligned} \Delta_{f,g} &= \mathbf{v}'_{f,g} - \mathbf{v}_{f,g} \\ &= \mathbf{v}_{f,g-1} + \frac{\mathbf{q}_{f,g-1}}{|\mathbf{q}_{f,g-1}|} |\mathbf{v}_{f,g}^0 - \mathbf{v}_{f,g-1}^0| - \mathbf{v}_{f,g} \end{aligned} \quad (5.2)$$

Once the displacements for the points in column g are determined, they are applied not only to those points, but also to all points downstream of column g . Thus, in each streamwise row f , all the points including and downstream of the point in column g are displaced equally. This procedure starts with the column of points immediately behind the lifting body's trailing edge, and progresses downstream column-by-column to the end of the wake, as shown in figure 5.2.

5.2.3 First column of wake panels

The column of wake panels adjoining the lifting body behaves slightly differently to the remainder of the downstream wake. The upstream edges of these panels are attached to the lifting body, and the coordinates of their downstream edges are controlled by velocities at the body's trailing edge. As the velocity over the body must everywhere be parallel to its surface, the wake panels in this first column will always be projected parallel to the trailing edge slope; it is not until the second column of panels that the wake can begin to 'roll-up'. For this reason, it is important that the first column of panels is quite short, so as not to artificially over-constrain the relaxing wake shape.




Figure 5.2: The wake relaxation process, applied to a two dimensional cross-section of the wing and wake. Movement of the first column of panels, constrained to remain in the plane of the lifting body, is not shown. In (a), the velocity is sampled at point 2, and the distance that point 3 needs to move to lie on the velocity projection is found. In (b), points 3-5 are moved by this amount. In (c), a modified wake geometry is fitted to the new point coordinates, applying gradient constraints at the upstream and downstream ends. In (d), the velocity is sampled at point 3, and points 4-5 are moved fitting a new wake shape in (e). Finally in (f), point 5 alone is moved, and another wake shape fitted. This completes one iteration; more iterations may be required to achieve convergence.

5.2.4 Far-field wake extension

For a steady-state solution, the wake should extend to infinity; in practice, the wake can be modelled efficiently by attaching a prismatic far-field wake patch to the downstream edge of the more accurately modelled relaxing wake patch. This far-field wake extension has a uniform cross-section that matches the shape of the standard wake's edge, and is extruded along the free-stream direction from that edge. As the singularity distribution is constant along the streamwise direction of the entire wake, including this far-field extension, the vorticity carried by the extension patch is also aligned with the free-stream.

Since there are closed-form expressions for the velocity and potential influences of semi-infinite vortex lines (Katz & Plotkin, 2001), the far-field extension could be modelled by a series of such lines, parallel to the free-stream and with strengths per unit span equal to the local spanwise doublet gradient. However, it is more straightforward to approximate the extension as a standard (but very simple) B-spline patch, linear in the stream direction, and extending a long way downstream rather than to infinity. In this way, the standard influence procedures can be used.

5.2.5 Geometry constraints

At each stage, the displacements of the wake points given by equation 5.2 provide geometric constraints for fitting a B-spline patch. In addition to these position constraints, three other types of constraints are employed in order to produce an exactly determined system of equations:

- The beginning of the wake must be connected to the wake-shedding edge of the lifting body. For wake and body patches with equal distributions of knots, the specification of position constraints at knots (and mid-knot points of end panels) on the wake edge ensures a seamless join.
- The Kutta condition (see section 3.1.6) requires flow to leave the trailing edge of an aerofoil smoothly: this provides a column of gradient constraints at the beginning of the wake. The constraints align the initial wake direction

with the trailing-edge of a thin aerofoil, or with the angle-bisector of a thick aerofoil.

- To create a seamless join between the standard wake patch and the far-field wake extension, a column of gradient constraints along the near-field patch's downstream edge constrains the wake to be aligned with the free-stream direction.

For a wake comprising a single patch, figure 5.1 shows a system in which position constraints formed by velocity projections are combined with the special position and gradient constraints above to form an exactly determined system by which the wake surface can be defined.

If a body sheds a wake from two or more edges, wake patches will meet along a line emanating from a corner of the body. In this situation, the mid-span position constraints within the adjoining rows of wake panels are replaced by equations constraining the spanwise gradients and curvatures to be equal across the join, thus forming a continuous wake surface.

5.2.6 Velocity sampling

Fundamental to the wake relaxation strategy is the evaluation of velocities within the wake. In general this is straightforward, and is achieved by the method detailed in section 3.5.3 for thin body mid-surface velocities. However, at the edge of a wake, where the spanwise doublet gradient is non-zero, the surface velocity is strictly undefined¹; this poses a problem for the wake-edge constraints specified in figure 5.1.

One solution would be to simply move these position constraints inboard of the wake edge. In practice though, this turns out to be unsatisfactory: if moved too far inboard, the relaxation process – unconstrained at the wake's edge – becomes less stable; and if moved only slightly inboard, the large amount of subdivision required to achieve an accurate velocity sample is computationally very inefficient.

¹Although strictly undefined, the velocity converges to a finite value as this point is approached, and it seems likely that further research may reveal a more elegant formulation than that proposed here.

In reality this wake edge is the centre of a vortex shed from the edge of the lifting body, and a pragmatic solution to the problem can be found by sampling the off-surface velocities at two points diametrically opposite the vortex centre, and taking their mean value as an estimate for the velocity at the vortex core. These off-surface velocities can be found via equation 3.142, and the two sample points can conveniently be defined on the local surface normal at a distance either side of the wake edge of the order 10^{-4} times the trailing edge span.

5.2.7 Surface fitting

Each of the constraints described specifies the position or gradient of a new wake geometry at specific intrinsic coordinates ξ_{fg} , or relates the gradient or curvature on one patch to that on another. The standard B-spline formulae (or the pre-computed versions presented in section 3.4.3) provide equations for the geometric constraints at ξ_{fg} in terms of unknown control point coordinates. The sets of equations for all wake patches are then combined into a single square matrix for solution via standard techniques, which provides B-spline control points and hence the newly relaxed wake geometries.

5.2.8 Singularity distributions during wake relaxation

Whilst the process for relaxing the geometric surfaces has been described in detail, little mention has been made of the singularity distributions that lie on the wake patches. In fact, as long the singularity distribution and the geometry are anchored to the same underlying intrinsic coordinates, deformation of the geometry has no effect on the singularity distribution. As the geometric grids relax, points on the wakes move such that the lines $\xi_1 = \text{const.}$ become streamlines emanating from the wake-shedding edges of the lifting bodies. As the doublet distributions within the wakes are defined by the same intrinsic coordinates ξ as the geometric surfaces, they are tied to the geometry such that when tracing a streamline, the doublet strength remains constant along it.

Within a conventional vortex lattice model, wake relaxations produce vorticity density changes as vortex lines of fixed strengths become closer together or

further apart. Within the CSM, changes in wake geometry effectively expand or compress the doublet distribution across the wake surface. This causes a change in doublet strength gradient with respect to distance along a wake cross-section, which corresponds to the change in vorticity seen within a vortex lattice.

5.2.9 Summary

Figure 5.3 presents the main elements of the process used to iteratively solve the potential flow around a body with wake relaxation. Criteria used to halt the main iteration cycle can be based on the convergence of either the potential flow solution over the body, or the geometry of the relaxing wake.

This scheme is similar to that employed in several simpler vortex lattice codes, and is an efficient means of obtaining a convergent wake shape. The scheme can be thought of as fitting a surface to a set of streamlines emanating from the trailing edge of the lifting body; the extra step of having to move the wake in streamwise stages is required because if the streamlines were allowed to flow off the surface of the wake, they would experience strong spanwise flow velocities in a direction that changes across the wake surface. In tracing the wake, we need to trace streamlines within the shear layer itself, not those lying either side of it.

An alternative approach used in some codes is to sample the velocities at all panel vertices simultaneously, and then move the wake in one step, rather than in stages. For this approach to be stable, all wake panels must have equal streamwise length, otherwise the geometry of the semi-relaxed wake becomes highly dependent on the relative panel sizes. For steady potential flow problems, efficiency can usually be gained by increasing the length of wake panels far from the lifting body; as such, this alternative scheme is better suited to unsteady problems which require regular wake panel grids.

5.3 Example of wake relaxation for a flat inclined plate

Figure 5.4 shows selected steps from the first two iterations of a wake relaxation, for the case of a flat inclined plate. The plate, with a span twice its chord, comprises a linear distribution of 3 spanwise and 5 chordwise panels. The wake,

SolveAndRelax

Initialisation

▷ Initialise wing patches (see **SolveFlow**, figure 3.18)

For each wake patch

▷ Prescribe initial wake shape

Next wake patch

Main loop

Repeat

▷ Obtain singularity distributions for current geometry (see **SolveFlow**)

For each column of wake points $g : 2 \rightarrow N_g$

For each row of wake points $f : 1 \rightarrow N_f$

▷ Evaluate coordinates of point $y_{f,g}$ at $\xi_{f,g}$

▷ Evaluate coordinates of upstream point $y_{f,g-1}$ at $\xi_{f,g-1}$

▷ Evaluate velocity $\mathbf{q}_{f,g-1}$ at upstream point

▷ Determine $\Delta_{f,g}$ (see equation 5.2)

For each column of wake points $g' : g \rightarrow N_g$

▷ Move point $y_{f,g'}$ by $\Delta_{f,g}$

Next column g'

Next row f

▷ Build equations comprising constraints from all wake patches

▷ Solve system to determine relaxed wake geometry

▷ Modify wake extension patch to fit wake geometry

Next column g

Until wake shapes or singularity distributions have converged

Post-processing

▷ Determine pressures, forces and moments

Figure 5.3: Outline algorithm for iteratively solving the potential flow around a body with wake relaxation

which is twice the chord length of the plate, has the same panel arrangement, and is connected to a far-field wake patch which extends beyond the standard wake a further 20 chord lengths of the plate. In this example, the first relaxation step is omitted: this is the step in which the second column of wake points can move in-plane only, and it has little effect on the relaxed wake shape. Images (b) to (e) show the stages of the first wake relaxation, in which the beginning of a vortex shape can be seen to propagate downstream, culminating in image (e) which shows the geometry after one complete relaxation iteration. After two complete iterations (f), the vortex roll-up is slightly more pronounced, and at this stage the geometry has almost converged: only slight differences are seen in the subsequent two iterations (not shown), and by the fifth the geometry is completely stable.

This example demonstrates the success of the wake relaxation strategy, and its ability to provide stable wake shapes with rapid convergence. The example is typical of many simple bodies shedding single wakes.

5.4 Example of wake relaxation for two overlapping sails

Figure 5.5 shows the relaxation of wakes shed from the mainsail and genoa of an America's Cup class yacht. In this example, linear sail and wake grids have been employed, with 6 panels in the span and 5 in the chord of either sail.

Image (a) shows the wake shape after the first full relaxation step. There is a significant roll-up in the wakes shed from the head of both sails, but a symmetry plane representing the hull and sea surface have prevented roll-ups occurring along the foot edges of the wakes. There is a slight penetration of the genoa wake through the mainsail, just downstream of the genoa's head. This is a result of the first column of wake panels being constrained to align with the trailing surface of the genoa, as described in section 5.2.3, and could be eliminated by the use of a shorter column of wake panels. However, because of the continuous representation of singularity surfaces within the CSM, the penetration does not create an instability in the wake modelling, allowing the relaxation to continue.

Images (b)-(e) show different views of the wake shape after five complete relaxation iterations. Whilst the wake surface immediately downstream of the




Figure 5.4: Selected wake relaxation steps for an inclined flat plate, with plate and wake each comprising a 3×5 panel grid. (a) Initial geometry, (b) Iteration 1: step 2, (c) Iteration 1: step 3, (d) Iteration 1: step 4, (e) Iteration 1 complete, (f) Iteration 2 complete

sails is similar to that in image (a), further downstream the roll-up has developed to the point where the top panels of the wake have folded through themselves. This has happened because the definitions of the B-spline wake surfaces, with low numbers of panels, do not allow sufficient freedom to model the true nature of the wake roll-ups.




Figure 5.5: Relaxation of wakes behind a genoa and mainsail. Image (a) shows the wake after the first relaxation step, and images (b)–(e) show different views of the wake shape after five relaxation steps.

5.4.1 Pressure and force variation with wake relaxation

Figure 5.6 shows the differential surface pressures over the sails before any wake relaxations (a), and after five relaxation steps (b). Although the two pressure distributions are very similar, slight differences can be seen, for example at the top of the mainsail's leading edge.

The effect that the relaxation steps have on total loads are more apparent than they are on the pressure distributions, as shown in figure 5.7. It is seen here that total loads are reduced by up to 8% through the course of the five relaxation steps.




Figure 5.6: Pressure distributions before relaxation (a), and after five relaxation steps (b). Both are plotted with the same colour scale, where blue indicates a low difference between pressures on the two sides of the sail, and red a high difference.

The folding-in of the wakes in the top panels presents a problem to convergence of the wake geometries, and this is seen in the slight oscillation of total loads in the last few iterations. The variations seen here of around 0.5% would be acceptable in many situations, but within a highly-tuned design optimisation, simulation noise such as this can be problematic. The following section proposes means by which the basic wake roll-up modelling process can be investigated and improved.

5.5 Alternative wake relaxation schemes

All higher-order methods have greater per-panel computational costs than their low-order counterparts, and must therefore rely on fewer panels to improve upon overall computation times for a given accuracy. Whilst we have seen that this can




Figure 5.7: Convergence of total aerodynamic loads generated by the sails with a relaxing wake. Iteration 1 corresponds to the original unrelaxed wake geometry.

be achieved with known geometries, the introduction of relaxing wakes increases the demands of a high-order method: not only must it produce potential solutions with a low number of panels, but it must also determine complex wake geometries with fewer wake sample points. Although it is strictly impossible to model the behaviour of a vortex core exactly with *any* finite geometry, the problem becomes ever more difficult as the resolution of wake samples decreases.

The CSM has the advantage that its wake relaxations are generally more stable than those of a vortex-lattice, because it avoids the singularities associated with discrete vortex lines. However, if there are too few sample points to approximate the roll-up of a wake into a vortex, or if the geometry definition does not comprise enough degrees of freedom to model it, problems like those seen with the overlapping sails can arise.

This final section considers three alternative schemes to improve the modelling of wake roll-ups. The aim of all of them is to improve the accuracy and stability of wake relaxations which can present modelling difficulties with more complex geometries.

5.5.1 Use of wake cross-sections

The alternative wake relaxation strategies are investigated by considering the development of spanwise cross-sections of a wake as they flow downstream, which enables a detailed two dimensional analysis to be undertaken before implementation in three dimensions. Appendix D provides full details of the two dimensional modelling assumptions, the three alternative wake relaxation strategies, results in two dimensions, and their extension into three dimensions.

5.5.2 Knot merging

This wake control scheme is implemented by testing for potential instabilities at each iteration, and if they occur, simplifying the geometry of the wake by reducing the number of control points at its extremities. This effectively creates a concentrated vortex core at the wake edge.

The knot merging scheme has shown some success in two dimensions, increasing the stability of wake relaxations, and preventing any ‘tangled’ wakes developing. However, whilst reducing the geometric complexity of wake edges, the merged edge knots also constrain the wake edge to be initially straight, a property which is clearly inaccurate at the centre of a wake vortex. Full details are provided in section D.2.

5.5.3 Truncation and redistribution

The truncation and redistribution scheme solves the above problem by separating from the main geometry any wake edges that are very tightly curved, and modelling them as an independent vortex core. This allows the remaining wake elements to form an accurate curved shape between wake sample points.

In this strategy, knots within the geometric B-splines are redistributed after every time-step to ensure sufficient flexibility to model the wake across its span. Regions at its extremities that become too tightly rolled-up are ‘cut away’ from the wake geometry and replaced by a concentrated line vortex element. This core is not attached explicitly to the wake edge.

Results in two dimensions are excellent, as shown in figure 5.8, with stable and accurate vortex shapes developing well downstream. The redistribution of spline knots at each relaxation iteration ensures there is sufficient freedom towards wake edges for the roll-up outside of the vortex core to be modelled explicitly. However, an unfortunate consequence of redistributing geometric spline knots is that the direct correspondence between intrinsic geometric and singularity surface coordinates is lost, requiring a mapping to be maintained between the two. Whilst this is quite straightforward in two dimensions, it is surprisingly complex in three dimensions, and much of the simplicity of the CSM method is lost in the three dimensional implementation of this scheme, as discussed in section D.3. It also requires a rather high density of wake sample points to operate well, and was not found to be a practical method for three dimensional modelling.



Figure 5.8: A wake cross section developed with the truncation and redistribution scheme after 50 steps. An explicitly modelled vortex core at the centre of each roll-up (not visible in the figure) contains the vorticity that has been ‘truncated’ from the visible wake sheet.

5.5.4 Framework shape functions

In this final scheme, a B-spline surface with a large number of control points is employed, capable of modelling a detailed roll-up geometry. In order to maintain a low computational cost, only a low density of wake sample points is defined, and tailored parametric spiral-like shape functions are fitted to the wake sample points, providing a framework to which the B-spline surface can be fitted. Full details are given in section D.4.

The method of employing tailored shape functions has shown promise in two dimensions, as shown in figure 5.9. It requires only a small number of wake points, and is the simplest scheme to implement in three dimensions. However, whilst the benefits of this scheme have been demonstrated in principle, further

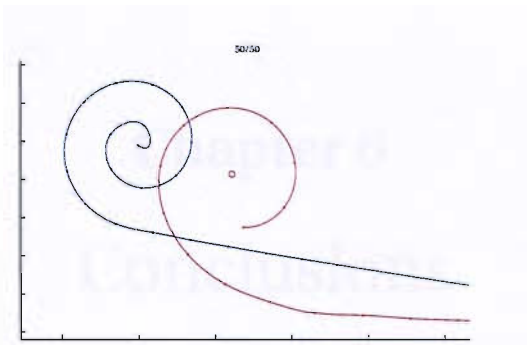


Figure 5.9: A comparison of a framework shape-function scheme with the truncation and redistribution scheme

work will be required to develop one or more parametric shape functions that could reliably model a wake roll-up with limited wake points, whilst remaining flexible enough to be used within a variety of geometric and flow situations. The truncation and redistribution scheme, which works exceptionally well in two dimensions, would be a good tool to use as a benchmark in further developing such shape functions.

5.5.5 Literature on vortex modelling

The development of wake vortices in two dimensions is in fact an area that has attracted considerable research. For example, Moore (1974), Fink & Soh (1978), Hoeijmakers & Vastra (1983), Krasny (1987) and Lamarre & Paraschivoiu (1992) consider discrete schemes for modelling wake roll-up into a vortex core, and discuss issues of instability similar to those described here. Comparisons are made to the analytical vortex spiral geometry of Kaden (1931), demonstrating the applicability of simple shape functions to the problem. It is expected that a review of work in this area would reveal methods that are highly relevant to the modelling of three dimensional wake roll-up within the CSM.

Chapter 6

Conclusions

A detailed survey of higher-order methods has been undertaken, and several schemes with linear or quadratic singularity or geometry surface representations have been investigated. It was found that although these methods alleviate some issues associated with constant panel methods, they fail to eradicate the main problems caused by the discontinuous modelling of geometry and singularity surfaces. A promising scheme is presented by Maniar (1995), although this only models thick bodies and Dirichlet boundary conditions; as such, it cannot model wakes.

A new higher-order method has been developed that models continuous source and doublet singularity distributions over three-dimensional curved surfaces. The singular on-surface influence coefficients are treated by a robust desingularisation algorithm, whereas off-surface coefficients are calculated by means of an efficient subdivision and variable cubature scheme. The Continuous Surface Method allows Dirichlet or Neumann boundary conditions to be applied to solve potential flow problems over thick bodies, thin surfaces and wakes modelled as doublet sheets.

The CSM has a number of advantages over a CPM:

- Curved geometries are represented exactly, so changing the number or distribution of unknowns in a solution does not physically change the shape of the body being analysed.
- Singularity distributions are continuous, and the layout of B-spline knots has only a weak influence on the possible solution shapes; this is in contrast to CPMs with which the panel layout pre-defines the positions of steps in the solution.

- Pressure distributions and the leading edge suction force are evaluated directly over the surface, without resorting to interpolation or extrapolation.
- Velocities sampled on or close to surfaces suffer no irregularities, discontinuities or singularities. In terms of wake relaxation, this means that wake surfaces do not ‘blow-up’ in the way they can do with a constant panel method.

6.1 Results with fixed geometries

Numerical results comparing the CSM to a conventional vortex lattice code have demonstrated the accuracy of the method for the example case of a curved sail with prescribed wake geometry. It was also found that the arrangement of knots in the singularity surface has less effect on the solution than the arrangement of panels in a constant panel method. The aim of less grid dependency has therefore been achieved.

In general, the CSM required fewer unknowns, or a smaller solution matrix, to achieve results of a given accuracy for the most commonly required post-processed quantities. However, for most applications, more important than a comparison by matrix size is a comparison by run time. For a given number of panels, the VLM is around 50 times quicker than the current implementation of the CSM, and this makes the two methods much closer competitors when compared on a run-time basis.

Total lift coefficients, when calculated in the Trefftz plane, are calculated more quickly by the VLM than by the CSM for the sail and wake example. This demonstrates that the high-order panels of the CSM have not significantly increased the accuracy of the spanwise distribution of vorticity modelled in the solution. It is in the more local results that the CSM has major advantages over the VLM, where the ability to directly evaluate surface pressures, without interpolation, leads to quicker convergence. This is an important result for fluid-structure interaction problems, such as the flow over deformable surfaces. For example, in the case of a yacht’s sail, the pressure distribution determines the

shape of the lifting surface, and a fluid-structure iteration cycle is required to obtain converged surface geometries, lift, and induced drag.

6.2 Wake relaxation

A wake relaxation scheme has been developed, and convergent solutions have been obtained for simple three dimensional cases such as the wake streaming from a flat plate. In more complex situations, such as the example of two overlapping sails with wakes, the wake relaxation algorithm converges to within 0.5% of total aerodynamic load, but fails to obtain a completely stable wake geometry.

The high accuracy but high cost per unknown of the CSM means that in practice, it will be used with a small number of unknowns. For instance, a 20×20 panel grid on a CPM would be matched in accuracy by a panel grid of around 5×5 with the CSM. It is this low panel density that poses a problem to the wake relaxation: it is difficult to model an accurate and convergent wake roll-up with a wake that has only five or six position constraints across its span. Whilst the geometric surface of the wake could contain more unknowns than the singularity distribution, the velocity samples are very costly to evaluate.

Three alternative wake relaxation schemes have been proposed to resolve the issues described. However, whilst the first two have both proved robust in a two dimensional prototype, neither was found to be stable in three dimensions. A third scheme employs an analytic or empirical model of wake roll-up geometries to create a framework of shape-functions. Only a low density of position constraints is required to parameterise the wake framework, over which a standard B-spline surface is fitted. Although further work is required to investigate suitable framework shape-functions, the scheme has shown promise in two dimensions and is straightforward to implement in three dimensions. In the continuing absence of a robust CPM solution to design problems involving closely interacting wakes, it is hoped that a future development of this third scheme, in conjunction with the existing CSM, will ultimately provide a definitively stable and highly accurate design tool.

Appendix A

Power series expansions for three dimensional influence coefficients

A.1 Series expansions for basic functions

In this appendix we develop the expressions needed to construct $F_{-1}(\theta)$ and $F_{-2}(\theta)$, the expansion of the kernel function, together with $\beta(\theta)$ and $\gamma(\theta)$, the expansion for the boundary equation $\alpha(\epsilon, \theta)$.

To avoid clutter, once a function dependent on certain coordinates has been introduced, its functional notation will frequently be omitted; for example, $\mathbf{r}(\mathbf{y}, \mathbf{x})$ will be written simply as \mathbf{r} . Also, magnitudes of vector quantities will be represented as $r = |\mathbf{r}|$.

A.1.1 Expansion of powers of \mathbf{r}

We start with a Taylor expansion of the vector from the singular point \mathbf{y} to a point \mathbf{x} on the integration surface, in terms of the intrinsic coordinates

$$\begin{aligned} \mathbf{r}(\mathbf{y}, \mathbf{x}) = \mathbf{x} - \mathbf{y} = & \left[\frac{\partial \mathbf{x}}{\partial \xi_1} (\xi_1 - \eta_1) + \frac{\partial \mathbf{x}}{\partial \xi_2} (\xi_2 - \eta_2) \right] \\ & + \left[\frac{\partial^2 \mathbf{x}}{\partial^2 \xi_1} \frac{(\xi_1 - \eta_1)^2}{2} + \frac{\partial^2 \mathbf{x}}{\partial \xi_1 \partial \xi_2} (\xi_1 - \eta_1)(\xi_2 - \eta_2) + \frac{\partial^2 \mathbf{x}}{\partial^2 \xi_2} \frac{(\xi_2 - \eta_2)^2}{2} \right] + \dots \quad (\text{A.1}) \end{aligned}$$

where all derivatives are evaluated at $\boldsymbol{\eta}$, image of \mathbf{y} in the intrinsic coordinate plane. By employing the same transformation to polar coordinates that we did in section 3.2.3, we can then write

$$\mathbf{r} = \rho \mathbf{A}(\theta) + \rho^2 \mathbf{B}(\theta) + O(\rho^3) \quad (\text{A.2a})$$

where

$$\mathbf{A}(\theta) = \left. \frac{\partial \mathbf{x}}{\partial \xi_1} \right|_{\xi=\eta} \cos \theta + \left. \frac{\partial \mathbf{x}}{\partial \xi_2} \right|_{\xi=\eta} \sin \theta \quad (\text{A.2b})$$

and

$$\mathbf{B}(\theta) = \left. \frac{\partial^2 \mathbf{x}}{\partial^2 \xi_1} \right|_{\xi=\eta} \frac{\cos^2 \theta}{2} + \left. \frac{\partial^2 \mathbf{x}}{\partial \xi_1 \partial \xi_2} \right|_{\xi=\eta} \cos \theta \sin \theta + \left. \frac{\partial^2 \mathbf{x}}{\partial^2 \xi_2} \right|_{\xi=\eta} \frac{\sin^2 \theta}{2} \quad (\text{A.2c})$$

In the expansions of F_{-1} and F_{-2} we will also require some integer power of $r = |\mathbf{r}|$. We start with the first power, and with reference to equation A.2a, we note that

$$r \rightarrow \mathbf{r} \cdot \left(\frac{\mathbf{A}}{A} \right) \quad \text{as } \rho \rightarrow 0 \quad (\text{A.3})$$

We can therefore write

$$r = \rho A + \rho^2 \frac{\mathbf{A} \cdot \mathbf{B}}{A} + O(\rho^3) \quad (\text{A.4})$$

Positive integer powers of $|\mathbf{r}|$ can quickly be obtained as

$$r^n = \rho^n A^n \left(1 + n\rho \frac{\mathbf{A} \cdot \mathbf{B}}{A^2} \right) + O(\rho^{n+2}) \quad (\text{A.5})$$

and by employing a Taylor series expansion of the form $\frac{1}{1+x}$ it is seen that in fact equation A.5 also holds for negative integers. In particular, we will require

$$\frac{1}{r} = \frac{1}{\rho A} - \frac{\mathbf{A} \cdot \mathbf{B}}{A^3} + O(\rho) \quad (\text{A.6a})$$

$$\frac{1}{r^3} = \frac{1}{\rho^3 A^3} - \frac{3\mathbf{A} \cdot \mathbf{B}}{\rho^2 A^5} + O\left(\frac{1}{\rho}\right) \quad (\text{A.6b})$$

$$\frac{1}{r^5} = \frac{1}{\rho^5 A^5} - \frac{5\mathbf{A} \cdot \mathbf{B}}{\rho^4 A^7} + O\left(\frac{1}{\rho^3}\right) \quad (\text{A.6c})$$

A.1.2 Expansion of the local Jacobian vector

An expansion of the local Jacobian vector $\mathbf{J}(\xi) = \mathbf{n}(\xi)J(\xi)$ (product of the Jacobian determinant and surface normal) is required for doublet distributions.

Following a similar path to the expansion of \mathbf{r} , we can write this as

$$\mathbf{J}(\boldsymbol{\xi}) = \mathbf{J}_0 + \rho \mathbf{J}_1(\theta) + O(\rho^2) \quad (\text{A.7a})$$

where

$$\mathbf{J}_0 = \mathbf{J}(\boldsymbol{\eta}) \quad (\text{A.7b})$$

and

$$\mathbf{J}_1(\theta) = \left. \frac{\partial \mathbf{J}}{\partial \xi_1} \right|_{\boldsymbol{\xi}=\boldsymbol{\eta}} \cos \theta + \left. \frac{\partial \mathbf{J}}{\partial \xi_2} \right|_{\boldsymbol{\xi}=\boldsymbol{\eta}} \sin \theta \quad (\text{A.7c})$$

The kernels associated with source distributions do not contain a surface normal, as sources are scalar rather than vector singularities. For these cases, we therefore need an expansion of the scalar Jacobian determinant, which we write as

$$J(\boldsymbol{\xi}) = J_0 + \rho J_1(\theta) + O(\rho^2) \quad (\text{A.8a})$$

where

$$J_0 = J(\boldsymbol{\eta}) \quad (\text{A.8b})$$

We do not need to derive an explicit expression for $J_1(\theta)$ since it does not appear in any final expressions, as will be seen in sections A.2.2 and A.2.4.

A.1.3 Expansion of the singularity distribution

Finally, an expansion of the singularity distribution is required. If this distribution is unknown, and represented by shape- or basis-functions, the expansion of these functions will be of identical form to that presented below.

$$\tau(\boldsymbol{\xi}) = \tau_0 + \rho \tau_1(\theta) + O(\rho^2) \quad (\text{A.9a})$$

where

$$\tau_0 = \tau(\boldsymbol{\eta}) \quad (\text{A.9b})$$

and

$$\tau_1(\theta) = \left. \frac{\partial \tau}{\partial \xi_1} \right|_{\boldsymbol{\xi}=\boldsymbol{\eta}} \cos \theta + \left. \frac{\partial \tau}{\partial \xi_2} \right|_{\boldsymbol{\xi}=\boldsymbol{\eta}} \sin \theta \quad (\text{A.9c})$$

A.2 Kernel function expansions

A.2.1 Potential influence of a doublet distribution

The scalar kernel function is

$$T(\mathbf{y}, \mathbf{x}) = -\frac{1}{4\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{r^3(\mathbf{y}, \mathbf{x})} \right\} \quad (\text{A.10})$$

and by defining

$$F(\rho, \theta) = \rho T(\mathbf{y}, \mathbf{x}(\boldsymbol{\xi}(\rho, \theta))) \mu(\rho, \theta) J(\rho, \theta) \quad (\text{A.11})$$

we obtain

$$F(\rho, \theta) = -\frac{1}{4\pi} \rho \left\{ \frac{\mathbf{r}(\rho, \theta) \cdot \mathbf{J}(\rho, \theta)}{r^3(\rho, \theta)} \right\} \mu(\rho, \theta) \quad (\text{A.12})$$

Substituting equations A.2a, A.6b, A.7a, and A.9a, for $\mathbf{r}(\rho, \theta)$, $r^3(\rho, \theta)$, $\mathbf{J}(\rho, \theta)$ and $\mu(\rho, \theta)$ respectively, the following expansion is easily obtained

$$\mathbf{F}(\rho, \theta) = \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \quad (\text{A.13a})$$

where

$$\mathbf{F}_{-1}(\theta) = -\frac{1}{4\pi} \left\{ \frac{\mathbf{A} \cdot \mathbf{J}_0}{A^3} \right\} \mu_0 = 0 \quad (\text{A.13b})$$

since the normal Jacobian vector $\mathbf{J}_0(\mathbf{y})$ is perpendicular to the surface tangent vector $\mathbf{A}(\theta)$. The integration can therefore be accomplished without a one dimensional integral relating to \mathbf{F}_{-1} or \mathbf{F}_{-2} , and the transformation of integration variables to a polar reference frame around the singular point is sufficient to deal with the weak singularity.

A.2.2 Potential influence of a source distribution

The scalar kernel function is

$$T(\mathbf{y}, \mathbf{x}) = -\frac{1}{4\pi} \left\{ \frac{1}{r(\mathbf{y}, \mathbf{x})} \right\} \quad (\text{A.14})$$

and by defining

$$F(\rho, \theta) = \rho T(\mathbf{y}, \mathbf{x}(\boldsymbol{\xi}(\rho, \theta))) \sigma(\rho, \theta) J(\rho, \theta) \quad (\text{A.15})$$

we obtain

$$F(\rho, \theta) = -\frac{1}{4\pi}\rho \left\{ \frac{1}{r(\rho, \theta)} \right\} \sigma(\rho, \theta) \quad (\text{A.16})$$

Substituting equations A.6a, A.8a and A.9a for $r^{-1}(\rho, \theta)$, $J(\rho, \theta)$ and $\sigma(\rho, \theta)$ respectively, it is seen that there are no terms of order ρ^{-2} or ρ^{-1} . In common with the potential influence of a doublet distribution, the weakly singular source distribution potential can be evaluated with just a two dimensional integral over a polar reference frame.

A.2.3 Velocity influence of a doublet distribution

The kernel function is

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = -\frac{1}{4\pi} \left\{ \frac{\mathbf{n}(\mathbf{x})}{r^3(\mathbf{x})} - \frac{3[\mathbf{r}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})]\mathbf{r}(\mathbf{x})}{r^5(\mathbf{x})} \right\} \quad (\text{A.17})$$

and by defining

$$\mathbf{F}(\rho, \theta) = \rho \mathbf{T}(\mathbf{y}, \mathbf{x}(\xi(\rho, \theta))) \mu(\rho, \theta) J(\rho, \theta) \quad (\text{A.18})$$

we obtain

$$\mathbf{F}(\rho, \theta) = -\frac{1}{4\pi}\rho \left\{ \frac{\mathbf{J}(\rho, \theta)}{r^3(\rho, \theta)} - \frac{3[\mathbf{r}(\rho, \theta) \cdot \mathbf{J}(\rho, \theta)]\mathbf{r}(\rho, \theta)}{r^5(\rho, \theta)} \right\} \mu(\rho, \theta) \quad (\text{A.19})$$

Substituting equations A.2a, A.6b, A.6c, A.7a, and A.9a, for $\mathbf{r}(\rho, \theta)$, $r^3(\rho, \theta)$, $r^5(\rho, \theta)$, $\mathbf{J}(\rho, \theta)$ and $\mu(\rho, \theta)$ respectively, and noting that $\mathbf{A} \cdot \mathbf{J}_0 = 0$, the following expansion is obtained by simple algebra

$$\mathbf{F}(\rho, \theta) = \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} + \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \quad (\text{A.20a})$$

where

$$\mathbf{F}_{-2}(\theta) = -\frac{1}{4\pi} \left\{ \frac{\mathbf{J}_0}{A^3} \right\} \mu_0 \quad (\text{A.20b})$$

$$\mathbf{F}_{-1}(\theta) = -\frac{1}{4\pi} \left\{ \left[\frac{\mathbf{J}_1}{A^3} - 3 \frac{(\mathbf{A} \cdot \mathbf{B})\mathbf{J}_0 + (\mathbf{B} \cdot \mathbf{J}_0)\mathbf{A} + (\mathbf{A} \cdot \mathbf{J}_1)\mathbf{A}}{A^5} \right] \mu_0 + \frac{\mathbf{J}_0}{A^3} \mu_1 \right\} \quad (\text{A.20c})$$

A.2.4 Velocity influence of a source distribution

The kernel function is

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = \frac{1}{4\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x})}{r^3(\mathbf{y}, \mathbf{x})} \right\} \quad (\text{A.21})$$

and by defining

$$\mathbf{F}(\rho, \theta) = \rho \mathbf{T}(\mathbf{y}, \mathbf{x}(\xi(\rho, \theta))) \sigma(\rho, \theta) J(\rho, \theta) \quad (\text{A.22})$$

we obtain

$$\mathbf{F}(\rho, \theta) = \frac{1}{4\pi} \rho \left\{ \frac{\mathbf{r}(\rho, \theta) J(\rho, \theta)}{r^3(\rho, \theta)} \right\} \sigma(\rho, \theta) \quad (\text{A.23})$$

Substituting equations A.2a, A.6b, A.8a and A.9a for $\mathbf{r}(\rho, \theta)$, $r^3(\rho, \theta)$, $J(\rho, \theta)$ and $\sigma(\rho, \theta)$ respectively, the following expansion is easily obtained

$$\mathbf{F}(\rho, \theta) = \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \quad (\text{A.24a})$$

where

$$\mathbf{F}_{-1}(\theta) = \frac{1}{4\pi} \left\{ \frac{J_0 \mathbf{A}}{A^3} \right\} \sigma_0 \quad (\text{A.24b})$$

There are no terms of order ρ^{-2} , because the kernel function in this case is strongly singular, rather than hypersingular.

A.3 Expansion of integration limit α

Integration of the functions $\mathbf{F}_{-1}(\theta)$ and $\mathbf{F}_{-2}(\theta)$ on the polar representation of the intrinsic coordinate frame is limited around the singular point by the function $\alpha(\epsilon, \theta)$. This function describes the equation of the boundary of e_ϵ , the vanishing region excluded from the integration.

The contour of $\alpha(\epsilon, \theta)$ is given by

$$\epsilon = r \quad (\text{A.25})$$

which by substituting equation A.4 can be written

$$\epsilon = \rho A(\theta) + \rho^2 \frac{\mathbf{A}(\theta) \cdot \mathbf{B}(\theta)}{A(\theta)} + O(\rho^3) \quad (\text{A.26})$$

We require an expansion for ρ , and write this

$$\rho = \epsilon\beta(\theta) + \epsilon^2\gamma(\theta) + O(\epsilon^3) \quad (\text{A.27})$$

By substituting equation A.27 into A.26, we obtain

$$\begin{aligned} \epsilon &= [\epsilon\beta + \epsilon^2\gamma]A + \epsilon^2\beta^2\frac{\mathbf{A} \cdot \mathbf{B}}{A} + O(\epsilon^3) \\ &= (\beta A)\epsilon + \left(\gamma A + \beta^2\frac{\mathbf{A} \cdot \mathbf{B}}{A}\right)\epsilon^2 + O(\epsilon^3) \end{aligned} \quad (\text{A.28})$$

and by equating powers of ϵ the explicit expressions we require are found to be

$$\beta(\theta) = \frac{1}{A(\theta)} \quad (\text{A.29a})$$

$$\gamma(\theta) = -\frac{\mathbf{A}(\theta) \cdot \mathbf{B}(\theta)}{A^4(\theta)} \quad (\text{A.29b})$$

Appendix B

Power series expansions for two dimensional influence coefficients

B.1 Series expansions for basic functions

In this appendix we develop the expressions needed to construct F_{-1}^m and F_{-2}^m , the expansion of the kernel function, together with β_m and γ_m , the expansion for the boundary equation $\alpha(\epsilon, \theta)$. To avoid clutter, once a function dependent on certain coordinates has been introduced, its functional notation will frequently be omitted; for example, $\mathbf{r}(\mathbf{y}, \mathbf{x})$ will be written simply as \mathbf{r} . Also, magnitudes of vector quantities will be represented as $r = |\mathbf{r}|$. Each of the expansions mentioned above need to be determined for both $m = 1$ and $m = 2$. The m subscript is omitted here, and it is assumed that all variables apply to the relevant section of the curve (i.e. either one side of \mathbf{y} with $m = 1$, or the other side with $m = 2$).

B.1.1 Expansion of powers of \mathbf{r}

We start with a Taylor expansion of the vector from the singular point \mathbf{y} to a point \mathbf{x} on the integration surface, in terms of the intrinsic coordinates

$$\begin{aligned}\mathbf{r}(\mathbf{y}, \mathbf{x}) &= \mathbf{x} - \mathbf{y} = \frac{\partial \mathbf{x}}{\partial \xi} (\xi - \eta) + \frac{\partial^2 \mathbf{x}}{\partial^2 \xi} \frac{(\xi_1 - \eta_1)^2}{2} + \dots \\ &= \rho \mathbf{A} + \rho^2 \mathbf{B} + O(\rho^3)\end{aligned}\tag{B.1}$$

where all derivatives are evaluated at η , image of \mathbf{y} in the intrinsic coordinate plane.

In the expansions of \mathbf{F}_{-1} and \mathbf{F}_{-2} we will also require some integer power of $r = |\mathbf{r}|$. We start with the first power, and with reference to equation B.1, we note

that

$$r \rightarrow \mathbf{r} \cdot \left(\frac{\mathbf{A}}{A} \right) \quad \text{as } \rho \rightarrow 0 \quad (\text{B.2})$$

We can therefore write

$$\begin{aligned} r &= \left| \rho A + \rho^2 \frac{\mathbf{A} \cdot \mathbf{B}}{A} \right| + O(\rho^3) \\ &= |\rho| A \left(1 + \rho \frac{\mathbf{A} \cdot \mathbf{B}}{A^2} \right) + O(\rho^3) \end{aligned} \quad (\text{B.3})$$

Positive integer powers of $|\mathbf{r}|$ can quickly be obtained as

$$r^n = |\rho|^n A^n \left(1 + n\rho \frac{\mathbf{A} \cdot \mathbf{B}}{A^2} \right) + O(\rho^{n+2}) \quad (\text{B.4})$$

and by employing a Taylor series expansion of the form $\frac{1}{1+x}$ it is seen that in fact equation B.4 also holds for negative integers. In particular, we will require

$$\frac{1}{r^2} = \frac{1}{\rho^2 A^2} - \frac{2\mathbf{A} \cdot \mathbf{B}}{\rho A^4} + O(1) \quad (\text{B.5a})$$

$$\frac{1}{r^4} = \frac{1}{\rho^4 A^4} - \frac{4\mathbf{A} \cdot \mathbf{B}}{\rho^3 A^6} + O\left(\frac{1}{\rho^2}\right) \quad (\text{B.5b})$$

B.1.2 Expansion of the local Jacobian vector

An expansion of the local Jacobian vector $\mathbf{J}(\xi) = \mathbf{n}(\xi)J(\xi)$ (product of the Jacobian determinant and surface normal) is required for doublet distributions.

Following a similar path to the expansion of \mathbf{r} , we can write this as

$$\mathbf{J}(\xi) = \mathbf{J}_0 + \rho \mathbf{J}_1 + O(\rho^2) \quad (\text{B.6a})$$

where

$$\mathbf{J}_0 = \mathbf{J}(\eta) \quad (\text{B.6b})$$

and

$$\mathbf{J}_1 = \left. \frac{\partial \mathbf{J}}{\partial \xi} \right|_{\xi=\eta} \quad (\text{B.6c})$$

The kernels associated with source distributions do not contain a surface normal, as sources are scalar rather than vector singularities. For these cases, we therefore need an expansion of the scalar Jacobian determinant, which we write as

$$J(\xi) = J_0 + \rho J_1 + O(\rho^2) \quad (\text{B.7a})$$

where

$$J_0 = J(\eta) \quad (\text{B.7b})$$

We do not need to derive an explicit expression for J_1 since it does not appear in any final expressions, as will be seen in sections B.2.2 and B.2.4.

B.1.3 Expansion of the singularity distribution

Finally, an expansion of the singularity distribution is required. If this distribution is unknown, and represented by shape- or basis-functions, the expansion of these functions will be of identical form to that presented below.

$$\tau(\xi) = \tau_0 + \rho \tau_1 + O(\rho^2) \quad (\text{B.8a})$$

where

$$\tau_0 = \tau(\eta) \quad (\text{B.8b})$$

and

$$\tau_1 = \left. \frac{\partial \tau}{\partial \xi} \right|_{\xi=\eta} \quad (\text{B.8c})$$

B.2 Kernel function expansions

B.2.1 Potential influence of a doublet distribution

The scalar kernel function is

$$T(\mathbf{y}, \mathbf{x}) = -\frac{1}{2\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{r^2(\mathbf{y}, \mathbf{x})} \right\} \quad (\text{B.9})$$

and by defining

$$F(\rho) = T(\mathbf{y}, \mathbf{x}(\boldsymbol{\xi}(\rho)))\mu(\rho)J(\rho) \quad (\text{B.10})$$

we obtain

$$F(\rho) = -\frac{1}{2\pi} \left\{ \frac{\mathbf{r}(\rho) \cdot \mathbf{J}(\rho)}{r^2(\rho)} \right\} \mu(\rho) \quad (\text{B.11})$$

Substituting equations B.1, B.5a, B.6a, and B.8a, for $\mathbf{r}(\rho)$, $r^2(\rho)$, $\mathbf{J}(\rho)$ and $\mu(\rho)$ respectively, the following expansion is easily obtained

$$\mathbf{F}(\rho) = \frac{\mathbf{F}_{-1}}{\rho} + O(1) \quad (\text{B.12a})$$

where

$$\mathbf{F}_{-1} = -\frac{1}{2\pi} \left\{ \frac{\mathbf{A} \cdot \mathbf{J}_0}{A^2} \right\} \mu_0 = 0 \quad (\text{B.12b})$$

since the normal Jacobian vector $\mathbf{J}_0(\mathbf{y})$ is perpendicular to the surface tangent vector $\mathbf{A}(\theta)$. The integration can therefore be accomplished without separating \mathbf{F}_{-1} or \mathbf{F}_{-2} from \mathbf{F} , and the integration in two parts either side of \mathbf{y} is sufficient to deal with the weak singularity.

B.2.2 Potential influence of a source distribution

The scalar kernel function is

$$T(\mathbf{y}, \mathbf{x}) = -\frac{1}{2\pi} \ln r(\mathbf{y}, \mathbf{x}) \quad (\text{B.13})$$

and by defining

$$F(\rho) = T(\mathbf{y}, \mathbf{x}(\boldsymbol{\xi}(\rho)))\sigma(\rho)J(\rho) \quad (\text{B.14})$$

we obtain

$$F(\rho) = -\frac{1}{2\pi} \rho \ln r(\rho)\sigma(\rho) \quad (\text{B.15})$$

Substituting equations B.1, B.7a and B.8a for $r(\rho)$, $J(\rho)$ and $\sigma(\rho)$ respectively, it is seen that there are no terms of order ρ^{-1} . In common with the potential influence of a doublet distribution, the weakly singular source distribution potential can be evaluated by simply splitting the integral into two parts either side of the weak singularity.

B.2.3 Velocity influence of a doublet distribution

The kernel function is

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = -\frac{1}{2\pi} \left\{ \frac{\mathbf{n}(\mathbf{x})}{r^2(\mathbf{x})} - \frac{2[\mathbf{r}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{x})]\mathbf{r}(\mathbf{x})}{r^4(\mathbf{x})} \right\} \quad (\text{B.16})$$

and by defining

$$\mathbf{F}(\rho) = \mathbf{T}(\mathbf{y}, \mathbf{x}(\boldsymbol{\xi}(\rho)))\mu(\rho)J(\rho) \quad (\text{B.17})$$

we obtain

$$\mathbf{F}(\rho) = -\frac{1}{2\pi} \left\{ \frac{\mathbf{J}(\rho)}{r^2(\rho)} - \frac{2[\mathbf{r}(\rho) \cdot \mathbf{J}(\rho)]\mathbf{r}(\rho)}{r^4(\rho)} \right\} \mu(\rho) \quad (\text{B.18})$$

Substituting equations B.1, B.5a, B.5b, B.6a, and B.8a, for $\mathbf{r}(\rho)$, $r^2(\rho)$, $r^4(\rho)$, $\mathbf{J}(\rho)$ and $\mu(\rho)$ respectively, and noting that $\mathbf{A} \cdot \mathbf{J}_0 = 0$, the following expansion is obtained by simple algebra

$$\mathbf{F}(\rho) = \frac{\mathbf{F}_{-2}(\theta)}{\rho^2} + \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \quad (\text{B.19a})$$

where

$$\mathbf{F}_{-2}(\theta) = -\frac{1}{2\pi} \left\{ \frac{\mathbf{J}_0}{A^2} \right\} \mu_0 \quad (\text{B.19b})$$

$$\mathbf{F}_{-1}(\theta) = -\frac{1}{2\pi} \left\{ \left[\frac{\mathbf{J}_1}{A^2} - 2 \frac{(\mathbf{A} \cdot \mathbf{B})\mathbf{J}_0 + (\mathbf{B} \cdot \mathbf{J}_0)\mathbf{A} + (\mathbf{A} \cdot \mathbf{J}_1)\mathbf{A}}{A^4} \right] \mu_0 + \frac{\mathbf{J}_0}{A^2} \mu_1 \right\} \quad (\text{B.19c})$$

B.2.4 Velocity influence of a source distribution

The kernel function is

$$\mathbf{T}(\mathbf{y}, \mathbf{x}) = \frac{1}{2\pi} \left\{ \frac{\mathbf{r}(\mathbf{y}, \mathbf{x})}{r^2(\mathbf{y}, \mathbf{x})} \right\} \quad (\text{B.20})$$

and by defining

$$\mathbf{F}(\rho) = \mathbf{T}(\mathbf{y}, \mathbf{x}(\boldsymbol{\xi}(\rho)))\sigma(\rho)J(\rho) \quad (\text{B.21})$$

we obtain

$$\mathbf{F}(\rho) = \frac{1}{2\pi} \left\{ \frac{\mathbf{r}(\rho)J(\rho)}{r^2(\rho)} \right\} \sigma(\rho) \quad (\text{B.22})$$

Substituting equations B.1, B.5a, B.7a and B.8a for $\mathbf{r}(\rho)$, $r^2(\rho)$, $J(\rho)$ and $\sigma(\rho)$ respectively, the following expansion is easily obtained

$$\mathbf{F}(\rho) = \frac{\mathbf{F}_{-1}(\theta)}{\rho} + O(1) \quad (\text{B.23a})$$

where

$$\mathbf{F}_{-1}(\theta) = \frac{1}{2\pi} \left\{ \frac{J_0 \mathbf{A}}{A^2} \right\} \sigma_0 \quad (\text{B.23b})$$

There are no terms of order ρ^{-2} , because the kernel function in this case is strongly singular, rather than hypersingular.

B.3 Expansion of integration limit α

Integration of the functions \mathbf{F}_{-1} and \mathbf{F}_{-2} on the intrinsic coordinate frame is limited around the singular point by the function $\alpha(\epsilon)$. This function describes the equation of the boundary of e_ϵ , the vanishing region excluded from the integration.

The contour of $\alpha(\epsilon)$ is given by

$$\epsilon = r \quad (\text{B.24})$$

which by substituting equation B.3 can be written

$$\epsilon = |\rho|A + \rho|\rho| \frac{\mathbf{A} \cdot \mathbf{B}}{A} + O(\rho^3) \quad (\text{B.25})$$

We require an expansion for $|\rho|$, and write this as

$$|\rho| = \epsilon\beta + \epsilon^2\gamma + O(\epsilon^3) \quad (\text{B.26})$$

By substituting equation B.26 into B.25, we obtain

$$\begin{aligned} \epsilon &= (\epsilon\beta + \epsilon^2\gamma) A + \epsilon^2\beta^2 \frac{\mathbf{A} \cdot \mathbf{B}}{A} \frac{\rho}{|\rho|} + O(\epsilon^3) \\ &= (\beta A)\epsilon + \left(\gamma A + \beta^2 \frac{\mathbf{A} \cdot \mathbf{B}}{A} \frac{\rho}{|\rho|} \right) \epsilon^2 + O(\epsilon^3) \end{aligned} \quad (\text{B.27})$$

and by equating powers of ϵ the explicit expressions we require are found to be

$$\beta = \frac{1}{A} \tag{B.28a}$$

$$\gamma = -\frac{\mathbf{A} \cdot \mathbf{B}}{A^4} \frac{\rho}{|\rho|} \tag{B.28b}$$

Appendix C

Vortex lattice parameter study

This appendix presents the results of a vortex lattice method parameter study, designed to find the optimal configuration of panels and collocation points for predicting flow over a specific body. This body is the thin mainsail surface with fixed wake described in section 4.1 and shown in figure 4.1. The geometry of the sail and wake were defined as bi-cubic B-spline surfaces, over which skewed quadrilateral elements were constructed to form the vortex lattice. The configuration parameters investigated were:

- *Panel spacing:* Grids of regularly spaced ‘Linear-Linear’ panels were compared with ‘Cosine-Cosine’ grids in which panels are concentrated towards surface edges. Expressions for the panel boundaries are equivalent to those for knots given in equations 4.1 and 4.2. For the linear grids, vortex lines were located 25% downstream of the panel boundaries, and collocation points 75% downstream of the panel boundaries, following the method of Katz & Plotkin (2001). For the cosine grids, vortex lines were located on panel boundaries, with collocation points on mid-cosine intervals, following the method of Lan (1974).
- *Panel aspect ratio:* The ratio of panel density in the span direction to that in the chord direction was varied from $(4N \times N)$ to $(N \times 2N)$. Grids with different panel aspect ratios, for both linear and cosine spacings, are shown in figure C.1.
- *Collocation point normal position:* Collocation points were either located on the curved B-spline surface used to define the vortex line grids, or

projected on to the skewed quad elements formed by the vortex lines. The two schemes are shown in figure C.2.

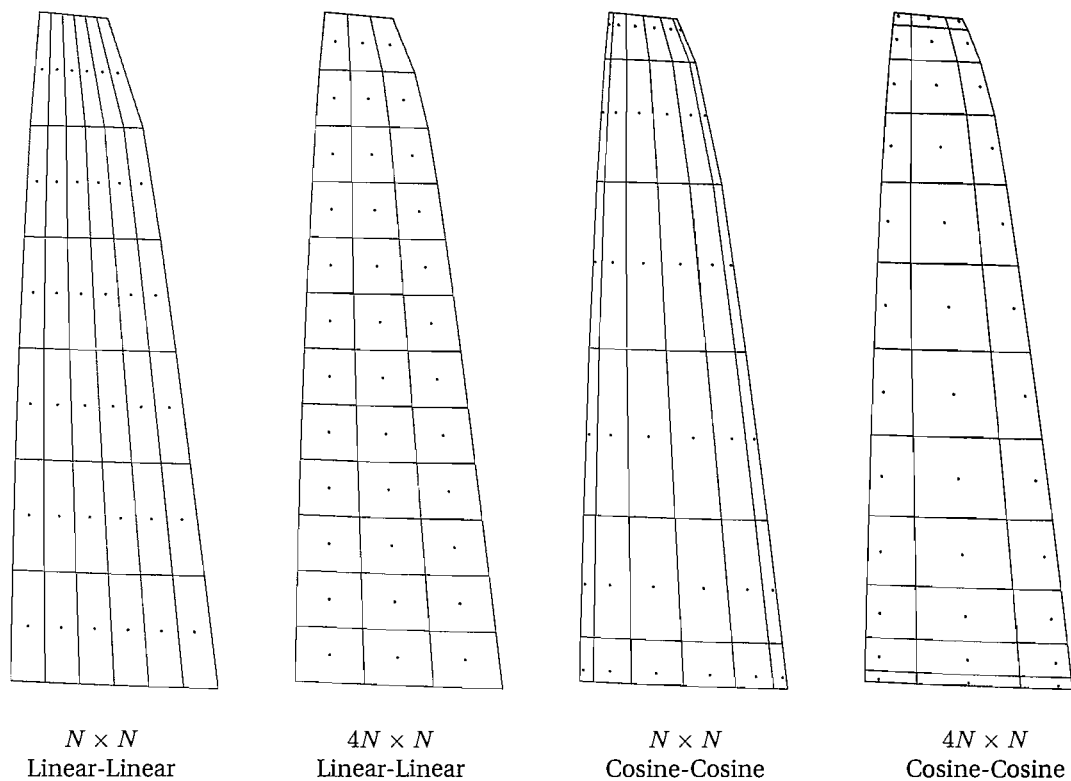


Figure C.1: Different panel spacings and aspect ratios




Figure C.2: Collocation points projected on to the skewed quadrilateral panels (left), and collocation points on the original curved surface (right)

C.1 Effect of configuration on edge suction forces

Figure C.3 shows how the contribution of the edge suction forces to C_L and C_D converges with the number of panels, and how it changes with panel aspect ratio. For the case shown, of cosine grids with collocation points projected on to the vortex lattice panels, it is seen that aspect ratio has little effect on the solution, which apparently converges to an incorrect value for all curves (the ‘assumed correct’ value has been calculated by comparing Trefftz plane force estimates with pressure summation estimates).


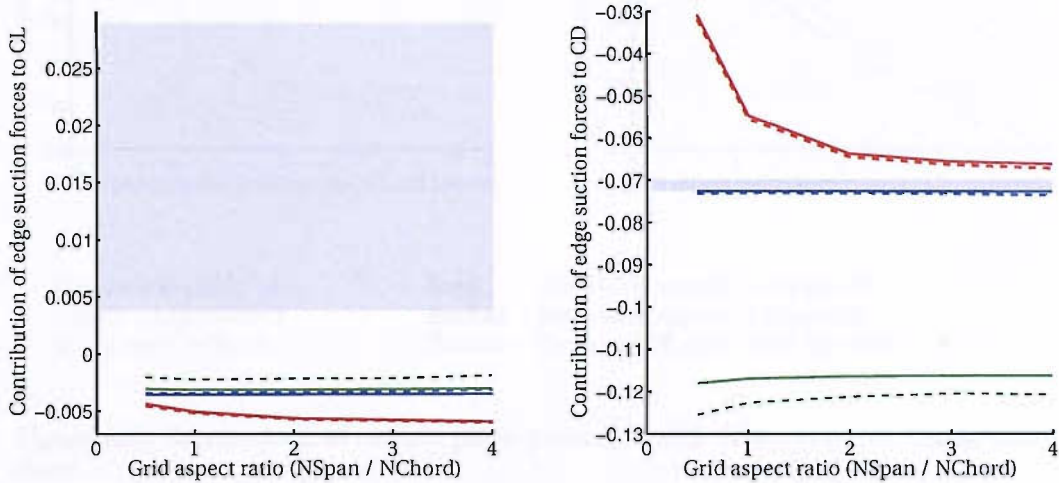


Figure C.3: Edge suction force convergence with different panel aspect ratios, with cosine grids and collocation points on panels

The results of figure C.3 are shown again in figure C.4, in which the edge suction force contributions are plotted against aspect ratio. The solid blue curve shows the ‘converged’ result for high panel densities, whilst the dashed curve shows the corresponding results for an intermediate panel density. The red curves show the equivalent results for the case where collocation points are on the curved surface, and the green curves show the results for linear grids. The contribution of edge suction forces to C_L is very small in all cases, but a significant variation is seen in the contributions to C_D . Whilst the cosine

grids with collocation points on panels produce relatively accurate and quickly converging solutions, moving them on to the curved surface dramatically reduces accuracy and introduces variation with aspect ratio. Using linear grids reduces the accuracy further, and the separation between the solid green and dashed green lines shows that convergence is slow.



Cosine grids, CPs on panels
 Cosine grids, CPs on surface
 Linear grids, CPs on panels


Solid Runs with approx. 2,400 panels
Dashed Runs with approx. 400 panels
Shaded Assumed 1% error band for total C_L & C_D

Figure C.4: Comparison of edge suction force contributions with different panel configurations

C.2 Effect of configuration on sample point pressures

Pressures sampled at points A and B on figure 4.1 are shown in figure C.5. Point A near the leading edge is more accurately sampled by the cosine grids, which concentrate panels in this region. At point B the pressures are sampled with similar accuracy by all configurations.

For both points A and B, the lower aspect ratios with lots of panels across the chord produce better results. The distance between the dashed and dotted lines on both graphs show that convergence is very slow in all cases.



Cosine grids, bFlat=1	<i>Solid</i>	Runs with approx. 2,400 panels
Cosine grids, bFlat=0	<i>Dashed</i>	Runs with approx. 400 panels
25/75 grids, bFlat=1	<i>Shaded</i>	Assumed 1% error band for total C_L & C_D

Figure C.5: Comparison of sample point pressures with different panel configurations

C.3 Effect of configuration on lift and drag coefficients

Lift and drag coefficients can be calculated for a vortex lattice solution by one of three methods:

- *Trefftz plane:* Total lift and drag forces are computed via a line integration along the wake in the Trefftz plane, distant from the sail body.
- *Pressure integration:* Doublet gradients are estimated at collocation points by taking a finite difference with neighbouring points; surface velocities are then calculated and assumed constant for each panel, from which constant panel pressures are determined.
- *Kutta-Joukowski method:* Each vortex line is assumed to produce a force equal to the cross-product of the line's vortex vector with the local fluid velocity (interpolated from velocities at collocation points). These forces are summed to determine total lift and drag coefficients.

Figure C.6 compares the convergence of lift and drag coefficients with increasing panel density for the above three calculation methods. It is seen that

the Trefftz plane method produces accurate results with very fast convergence, whereas the other two methods are slower to converge and appear to be converging to an inaccurate solution. Note that both these latter two methods require the addition of the edge suction forces to produce the lift and drag coefficients, so that any error or convergence issues in the edge suction forces will affect the solutions shown.




Figure C.6: Lift and drag coefficient convergence calculated via three methods, for cosine ($N \times N$) grids with collocation points on panels.

Figure C.7 (Top) shows the same results as the previous figure, but now includes a variation with aspect ratio (all results in figure C.6 were for an $(N \times N)$ configuration). For this case of cosine grids with collocation points on panels, the Trefftz plane calculations are largely unaffected by aspect ratio, whereas the other results are most accurate when the number of chordwise panels is greatest.

The middle two graphs in figure C.7 show corresponding results for cosine grids with collocation points on the curved surface, and the bottom graphs show results for linear grids with collocation points on panels. Whilst the Trefftz plane calculations remain by far the most accurate in all cases, both of these configurations produce poor results in comparison to the cosine on-panel grids. Although edge suction forces can be blamed for some of the errors seen in C_D ,

they do not have a significant effect on CL, where it is seen that the linear grids are significantly less accurate than the cosine grids.

C.4 Variation of computation time with number of panels

Figure C.8 shows how computation time increases with the number of body panels in the configuration. A fit line of order N^2 (where N here is the total number of panels) shows that the total run time is approximately proportional to the number of panels multiplied by the number of collocation points in the solution. For very low panel numbers the initialisation and general program control code becomes significant, leading to results above the low end of the trend line.

C.5 Conclusions

Analysis of results from 225 vortex lattice simulations has shown that an optimal configuration for obtaining accurate and quickly converging lift coefficients, sample point pressures and edge suction forces is as follows:

- A cosine-cosine panel spacing
- An $(N \times N)$ or $(N \times 2N)$ panel density (the former is chosen here for consistency with the CSM models)
- Collocation points projected on to the skewed quadrilateral vortex panels

This configuration is employed in section 4.5 where vortex lattice results are compared against those from the CSM.

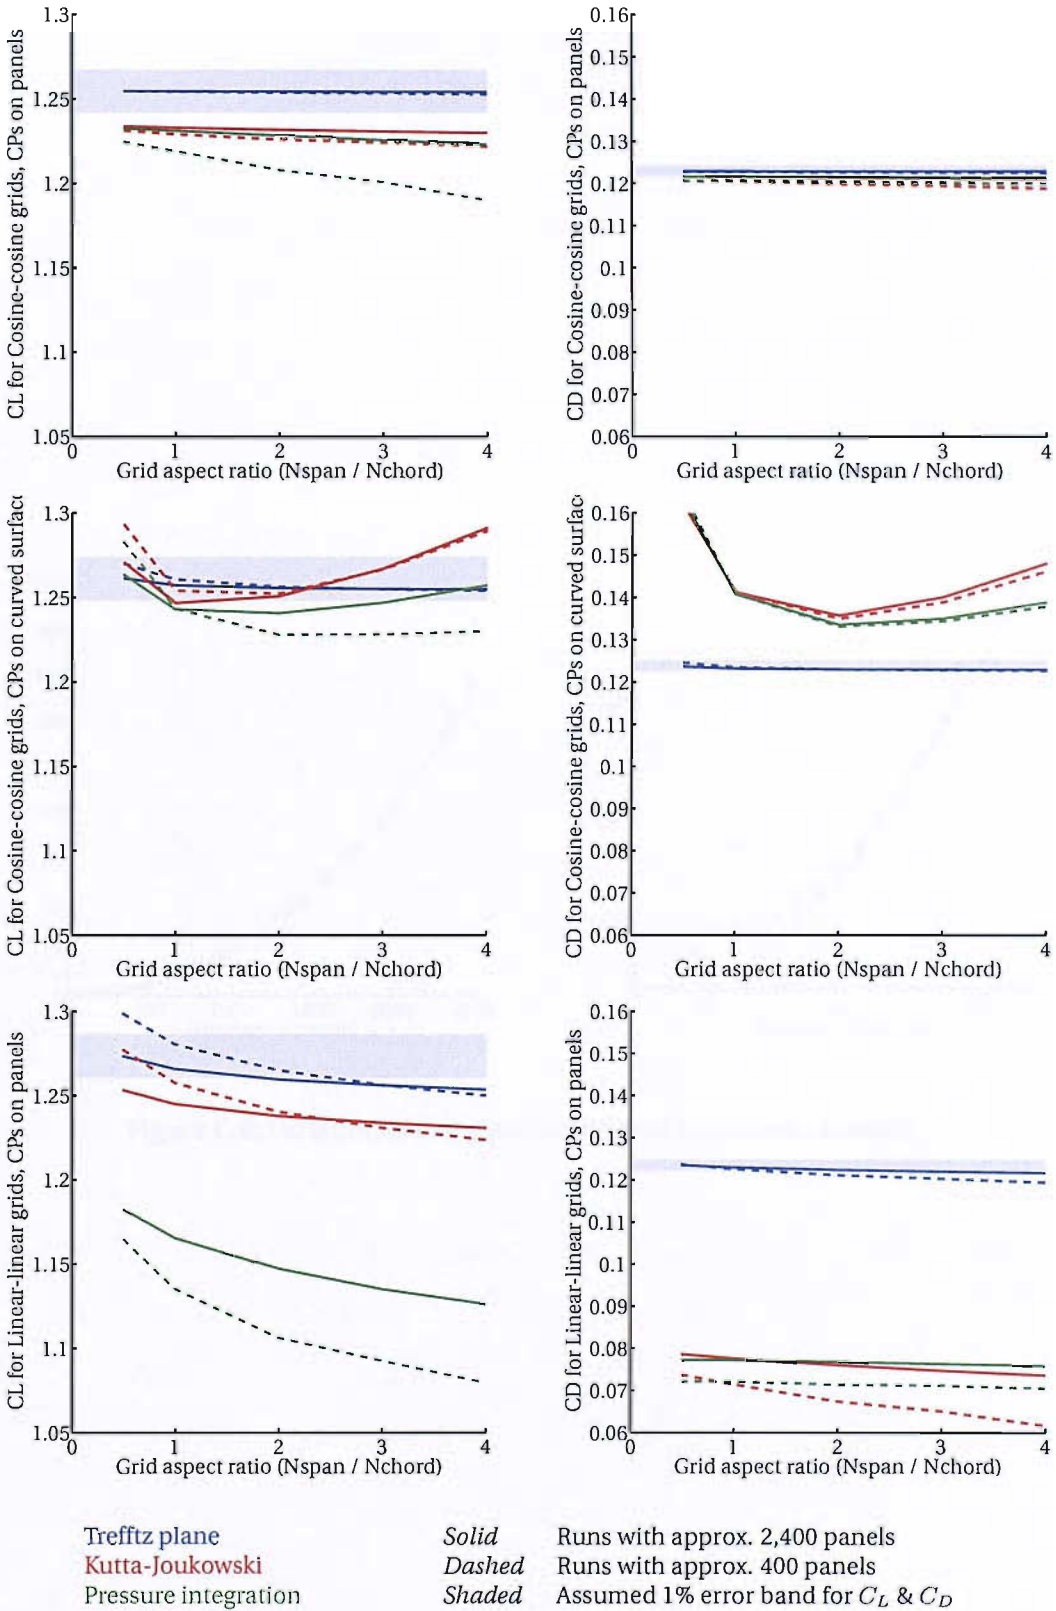


Figure C.7: Lift and drag coefficients calculated via three alternative methods.




Figure C.8: Variation of computation time with number of panels

Appendix D

Alternative wake relaxation schemes

D.1 Studies on a two dimensional cross-section

The problems encountered in the three dimensional wake roll-up can conveniently be investigated in two dimensions by considering a cross-section of the wake downstream of the lifting body. The cross-section's cutting-plane is considered to be within a moving frame of reference that travels downstream with the free-stream velocity. As such, it approximately models the non-streamwise velocity components of a column of particles that leave the trailing edge at the same time. This two dimensional representation assumes the wake upstream and downstream of the cutting plane to be a uniform extrusion of the cross-section, which is reasonable in terms of the influence that it has on the two dimensional wake section under investigation.

The two dimensional cross-sectional wake is allowed to relax in a series of time-steps, where each step in two dimensions corresponds to a displacement downstream in three dimensions. As the CSM's three dimensional wake relaxation process operates by projecting displacements from upstream points to new positions downstream, this two dimensional scheme is a good model for the three dimensional one: the same points are projected to form new wake shapes, just without adding the approximately constant streamwise velocity components.

D.1.1 Basic two dimensional wake relaxation scheme

The two dimensional wake cross-sections considered in this appendix model a three dimensional wake shed from the trailing edge of a flat inclined plate. The wake cross-sectional geometry $\mathbf{x}(\xi_1, t)$ at time t is modelled by a B-spline curve

with parametric coordinate ξ_1 , where $\xi_1 = 0$ corresponds to one end of the line, and $\xi_1 = 1$ to the other. The initial wake shape for a cross-section that starts at the plate's trailing edge is a flat line.

The doublet distribution in a steady-state three dimensional solution is constant with the streamwise parametric coordinate ξ_2 , and as such, we model the two dimensional cross-section's doublet distribution $\mu(\xi_1)$ as being constant with time. This is in fact an approximation, because through the relaxation process of the three dimensional wake towards the steady-state solution, the wake's doublet distribution does change as the relaxing wake shape affects the lifting body's potential flow solution; however, the change is not significant enough to influence the basic shapes and problems being investigated with the two dimensional scheme. The doublet distribution used within the two dimensional analysis is taken from the trailing edge of a three dimensional solution with a fixed wake geometry.

The process for relaxing the two dimensional wake is very simple, and analogous to the three dimensional scheme: in each iteration, induced velocities are calculated at specific points within the wake curve, and by assuming a time-step Δt , displacements for these points are determined. A new wake geometry is then fitted to the revised point positions to complete the iteration. Note that a free-stream velocity is not included in the two dimensional velocity calculations because it is mostly out of plane, and any in-plane components would in any case only shift the wake by a uniform displacement. However, the out-of-plane free stream velocity Q^∞ is used to set the time-step Δt according to

$$\Delta t = \frac{s}{Q^\infty} \quad (\text{D.1})$$

where s is the typical distance between panels in the three dimensional case.

D.1.2 Limitations of a two dimensional model

One significant shortcoming of using a two dimensional model to test the viability of a three dimensional scheme is that in two dimensions, there is no concept of iterating towards convergence: each time-step represents a real development of

a wake shape as it propagates downstream, rather than a numerical development towards an accurate solution. The two dimensional model is therefore ideal for investigating behaviour in the first relaxation step, but less useful in predicting whether or not an equivalent three dimensional scheme will be convergent.

D.1.3 Example of basic two dimensional wake relaxation

Figure D.1 shows a two dimensional wake after five time-steps, representing a cross-section through the three dimensional wake shed from a flat plate at a distance of five wake panels downstream from the plate's trailing edge. The model has a relatively large number of segments (the two dimensional equivalent of panels), with a concentration towards the edges of the wake, intended to capture the geometry of the wake roll-up.

We are generally interested in the behaviour at the edges of the wake, and figure D.2 shows two developments of a wake edge through a number of iterations. The wakes' geometries are represented by piecewise linear and piecewise cubic B-spline curves, where the development of the piecewise linear wake is very similar to that which would be obtained from a vortex lattice code. It is seen that both wakes eventually become unstable, principally because there is insufficient resolution in the representation of the geometry to form the rolled-up shape that the wake is moving towards.

Although the wake developments seen in the final steps in figure D.2 are inaccurate and undesirable because they prevent convergence in a wake relaxation, they do not generally produce the near-singular velocities that can cause vortex lattice relaxations to 'explode'. This is because they are not modelled by discrete vortex lines that produce singular velocity fields as they are approached. The exception, where near singular velocities can be seen within




Figure D.1: A two dimensional wake section represented by 42 B-spline segments after five time-steps

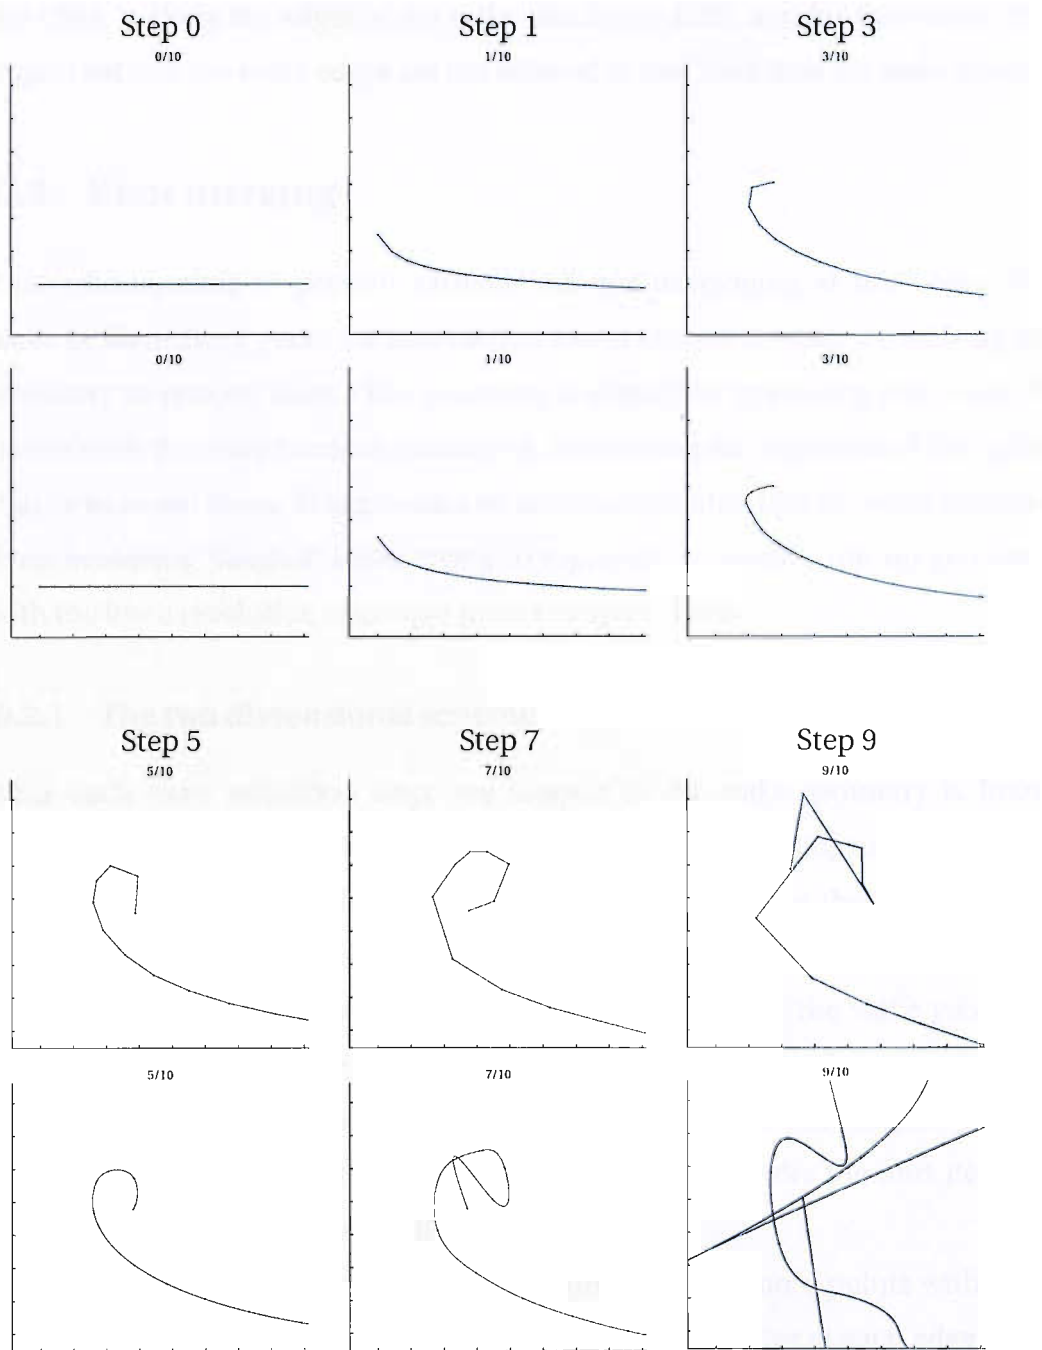


Figure D.2: Wake development with geometry modelled by piecewise linear (first and third rows) and piecewise cubic (second and fourth rows) B-spline curves. The points within the wake that define the two curves initially move along very similar paths, such that the only difference between the geometries is the straight-line or cubic interpolation between them. However, in Step 7, the cubic B-spline geometry fits a poor wake curve through the wake points; the marked difference between the two geometries then causes differences in the induced velocities they produce, and hence divergence in future geometries. Although the piecewise linear B-spline curve remains acceptable for one or two steps more than the piecewise cubic curve, it also eventually fails in its ability to model the roll-up.

the CSM, is along the edges of the wake (see figure 4.25), and for this reason it is important that the wake edges are not allowed to fold back into the wake sheet.

D.2 Knot merging

This scheme aims to prevent extreme roll-ups developing at the edges of a wake by identifying potential instabilities, and if any are present, simplifying the geometry to remove them. The geometry is simplified by moving edge wake fit points such that they become coincident, eliminating the segments of the spline that lie between them. This provides an automatic control that prevents the wake from becoming ‘tangled’ when trying to represent a complex roll-up geometry with too low a resolution of sample points or spline knots.

D.2.1 The two dimensional scheme

After each wake relaxation step, the tangent to the wake geometry is found at each spline knot, and the angle between neighbouring tangents tested. An angle greater than a threshold, for example π , indicates that the geometry has become too tight, and must be simplified. This simplification is achieved by making the two outermost points that define that edge of the wake geometry mutually coincident: they are moved to some point lying in between the two according to the relative vorticity associated with each one. The tangent tests are then repeated, and further points merged if necessary, before the next iteration proceeds. The procedure is detailed in figure D.3.

In this way, as the relaxation develops, progressively more points within the wake become combined with the point at the wake’s edge (or at each edge of the wake). This means that a smaller range of the parametric coordinate ξ_1 now spans the entire wake, such that:

$$\xi_1 \leq \xi_1^L \quad : \quad \text{Left-hand wake edge} \quad (\text{D.2})$$

$$\xi_1^L < \xi_1 < \xi_1^R \quad : \quad \text{Wake span} \quad (\text{D.3})$$

$$\xi_1^R \leq \xi_1 \quad : \quad \text{Right-hand wake edge} \quad (\text{D.4})$$

RelaxAndMerge2D

Initialisation

- ▷ Define parametric coordinates of wake sample/constraint points (all knots plus mid-points of end panels)

Main loop

For each wake relaxation step

- ▷ Calculate velocity at each wake sample point
- ▷ Determine new geometric position of wake points
- ▷ Define additional gradient constraints

Repeat

- ▷ Build square system of equations from constraints
- ▷ Solve to determine new spline geometry

For each spline edge

- ▷ Test angle between tangents at adjacent wake points
- ▷ Test maximum velocities induced by concentrated vortices

If any point on this edge fails

- ▷ Merge edge two spline knots (weighted by vorticity)
- ▷ Combine edge two wake sample points

End If

Next spline edge

Until no more points need merging

Next wake relaxation step

Figure D.3: An outline algorithm for a two dimensional wake relaxation strategy with wake point merging

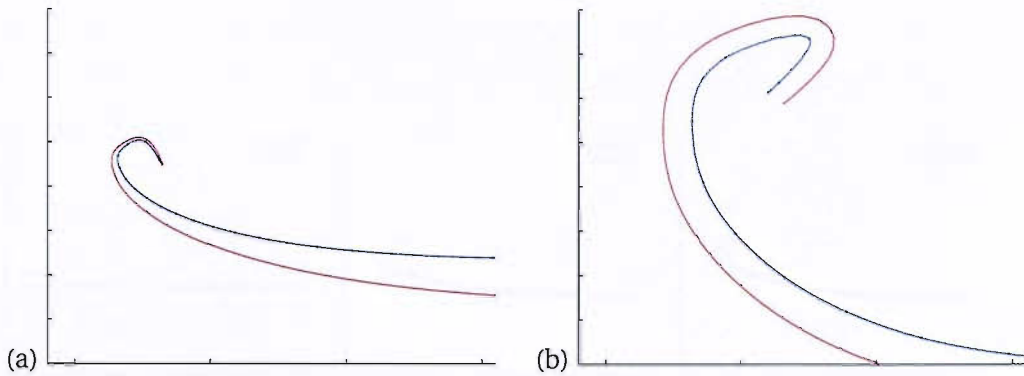


Figure D.4: The effect of wake truncation on doublet distribution. The blue line shows a portion of the wake geometry, and the perpendicular distance of the red line away from the blue line represents the doublet strength at that point. In (a), the wake is not truncated, and the doublet strength falls gradually to zero towards the wake edge. In (b), several wake iterations later, the wake roll-up has grown, and several geometric and doublet strength segments have been ‘absorbed’ into the concentrated edge vortex. The magnitude of the edge vortex is represented by the end value of the doublet strength.

where ξ_1^L and ξ_1^R mark the extent of ξ_1 on the remaining geometric wake. Recalling that the doublet distribution $\mu(\xi_1)$ is a function of the same parametric coordinate as the geometry, it is clear that within the two wake edges, now compressed to points, there will be a range of doublet distribution: in the left edge $\mu(0) \rightarrow \mu(\xi_1^L)$, and in the right edge $\mu(\xi_1^R) \rightarrow \mu(1)$. These steps in doublet strength over a geometric point are in fact concentrated vortices, so the scheme can be thought of as replacing the most tightly rolled-up portions of the wake with a concentrated vortex line attached to the wake edge.

Figure D.4 shows how, with the wake truncated in this way, the doublet distribution over the wake does not fall gradually to zero towards the wake edge, but ends in a step.

D.2.2 Results in two dimensions

Figure D.5 shows a selection of relaxation steps created using this scheme, detailing one edge of the wake, and figure D.6 shows the complete wake at the end of the series.

One or more merge operations is required after every few steps, and it is seen that although the wake sometimes becomes misshaped, it tends to revert

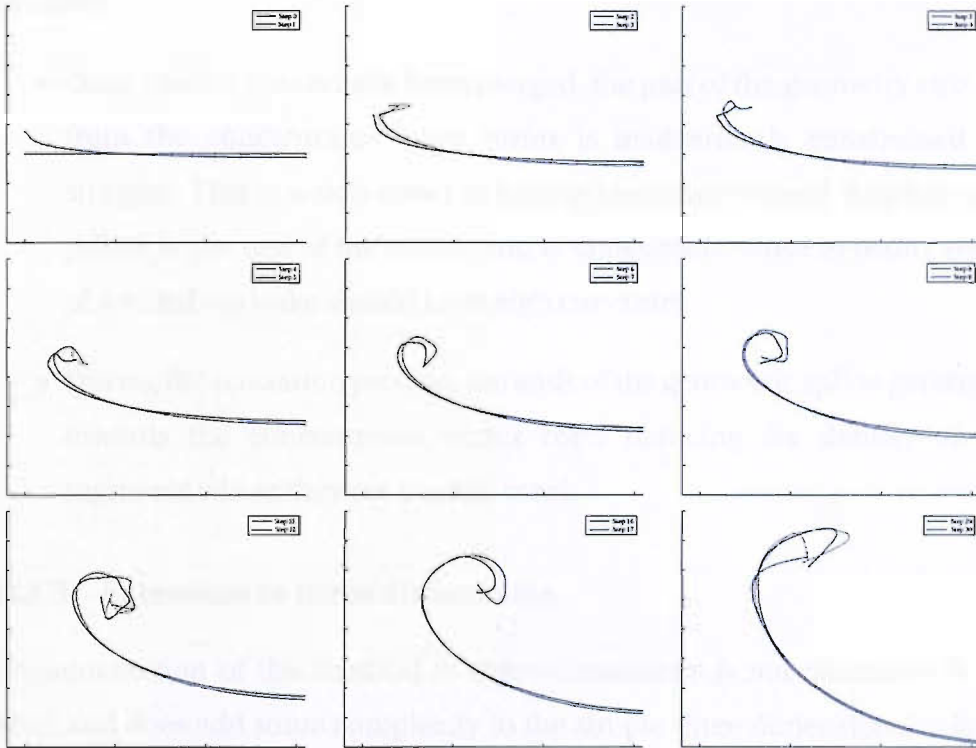


Figure D.5: A selection of steps from a wake relaxation with knot merging. The black line in each figure is the starting point for each relaxation step and the blue line is the relaxed wake at the end of that time step. In some of the plots, one or more grey lines indicate the relaxed shapes before point merging has taken place, with one grey line shown prior to each merge operation. In these plots, the sequence is thus black \rightarrow grey (following relaxation) \rightarrow blue (following merge).




Figure D.6: A wake relaxed using the knot-merging scheme, at step 29

to a smooth curve after a few more steps. Overall, the method is stable in two dimensions, but is thought to produce a somewhat inaccurate wake shape because:

- Once control points have been merged, the part of the geometry emanating from the concentrated edge vortex is inadvertently constrained to be straight. This is a side-effect of having identically valued B-spline control points at the end of the spline, and is undesirable, since in reality the edge of a rolled-up wake should have high curvature.
- During the relaxation process, the ends of the geometric spline get stretched towards the concentrated vortex core, reducing the density of spline segments where they are needed most.

D.2.3 Extension to three dimensions

Implementation of this method in three dimensions is not described in detail here, and does add some complexity to the simple three dimensional relaxation scheme laid out in section 5.2. Tables need to be maintained to store the extent of the remaining wake at each streamwise station, and influence calculations need to be modified to cope with doublet sheets concentrated into an edge vortex.

In practice, it was found that this method did not perform well in three dimensions, and convergent wake relaxations were difficult to achieve. As discussed in section D.1.1, success in two dimensions does not necessarily imply convergence in three dimensions, and it emerges that the discrete nature of the wake merging process (whole segments being either merged or not) hinders convergence. If a segment just passes a roll-up criteria in one iteration, but fails in the next, the geometry can change quite significantly and relaxation has to start again.

D.3 Truncation and redistribution

This scheme aims to solve the main problems of the knot merging process described in the previous section. Rather than merging control points to form

a concentrated vortex at the edge of a wake surface, portions of the wake that are considered too tightly curved to be modelled explicitly are replaced by specific vortex-core elements (point vortices in two dimensions, vortex lines in three dimensions) that are not attached to the wake surface, and are treated separately in the relaxation process. In order to give sufficient flexibility to the remaining wake geometry, the spline knots are redistributed over the surface during each iteration, maintaining the original knot spacing.

D.3.1 The two dimensional scheme

At the beginning of each relaxation step, the wake is tested against roll-up criteria: for instance, a minimum curvature, or a maximum rotation of the wake surface of 2π , measured from the centre of the wake, might be allowed. The portion of the wake that falls within these criteria has parametric coordinate ξ_1 where $\xi_1^L \leq \xi_1 \leq \xi_1^R$. Note that with this scheme, ξ_L and ξ_R can be anywhere on the spline, not only on a spline knot.

On each side of the wake, if there is an edge region that falls outside of the above criteria (i.e. $\xi_1^L > 0$ or $\xi_1^R < 1$), that part of the wake is replaced by a point vortex at some weighted centre of the region. The strength of the vortex is equal to the difference in doublet strength between its two ends. If a point vortex already exists on that side of the wake from an earlier truncation, the new and existing vortices are combined at a weighted centre point. Using a weighted centre approach for both of these cases is an approximation that results in inaccuracies in near-field velocity influences; however in practice the relative distances involved are usually small, and errors are therefore small. Note that when part of the wake is replaced by a point vortex element, an equal and opposite vortex element needs to be included at the edge of the remaining wake when velocity influences are determined: the combination of the two opposite vortex elements then effectively translates the vortex concentration from the edge of the remaining wake to the location of the unattached point vortex.

Within the remaining region $\xi_1^L \leq \xi_1 \leq \xi_1^R$, a table is constructed relating the parametric coordinate ξ_1 to the curvilinear distance $\lambda(\xi_1)$ across the wake,

starting from $\lambda(\xi_1^L) = 0$. This distance is then normalised by

$$\bar{\lambda}(\xi_1) = \frac{\lambda(\xi_1)}{\lambda(\xi_1^R)} \quad (\text{D.5})$$

such that $\bar{\lambda}$ becomes a new parametric coordinate spanning the remaining wake, and has a constant scaling with distance. Where the original wake velocity sampling points were at $\xi_1 = \mathbf{V}$, new velocity sample points are specified at $\bar{\lambda} = \mathbf{V}$ using the inverse function $\xi_1 = \bar{\lambda}'(\mathbf{V})$, effectively redistributing the sample points prior to the relaxation step. Velocity projections are found from these points, and a new geometry fitted to the points $\xi_1 = \mathbf{V}$. Point vortices representing truncated wake edges are relaxed at the same time as the wake geometry, determining their velocities using the off-surface influence coefficient evaluations described in section 3.3.4.

By re-mapping the parametric coordinate ξ_1 in each relaxation step, the direct correspondence between the doublet distribution and the geometry is lost: the same coordinate ξ_1 can no longer be used to directly map the two surfaces to each other. We therefore need to maintain a function relating $\bar{\xi}_1$ for the doublet surface to ξ_1 for the geometric surface. This function is formed by the application of $\lambda(\bar{\xi}_1)$ mappings over successive relaxation steps. In order to reduce numerical errors, it is therefore necessary to construct and maintain $\lambda(\bar{\xi}_1)$ at a fine discretisation, or alternatively approximate $\lambda(\bar{\xi}_1)$ by an analytic function.

An example step of the re-distribution algorithm is shown in figure D.7.

D.3.2 Results in two dimensions

Figure D.8 shows a selection of relaxation steps created using this scheme, detailing one edge of the wake, and figure D.9 shows the complete wake at the end of the series.

As seen in the figures, the scheme is extremely stable, producing very well formed wake roll-ups around vortex core elements. Whilst the example illustrates the stability, the final roll-up shown of over 4π would generally be unnecessary, and a more practical limit of 2π would further increase the robustness of the method.

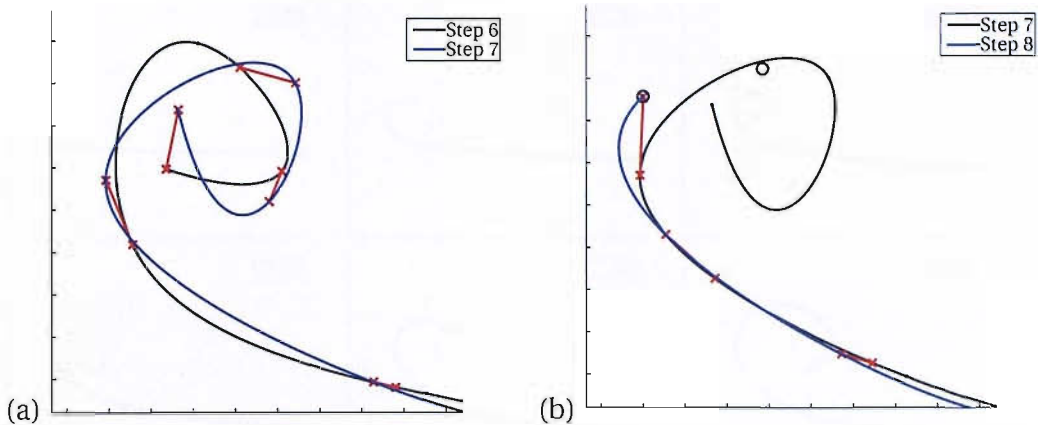


Figure D.7: A sequence of three relaxation steps developed with the truncation and redistribution scheme: In (a) the geometry at the completion of step 6 is shown in black, and the slightly more rolled-up geometry of step 7 in blue. The red lines show how wake sample points have moved from one step to the next. Up to and including step 7, there is no truncation, and the entire wake is modelled explicitly within the spline. In (b), the step 7 geometry is shown again, this time in black, and the new geometry of step 8 is shown in blue. In this case, the geometry has been truncated, such that an internal wake sample point of step 7 has become the extremity of the spline in step 8. The rolled-up region of step 7 between the last red cross and the end black dot has been replaced by a concentrated vortex element shown by the unattached black circle. An equal and opposite vortex element, shown by the other black circle, is attached to the end of the truncated spline.

D.3.3 Extension to three dimensions

Implementation of this method in three dimensions is not described in detail here, and unfortunately adds a great deal of complexity to the simple three dimensional relaxation scheme laid out in section 5.2. The process of maintaining relationships between doublet and geometric parametric coordinates in two dimensions is quite straightforward, and in three dimensions is carried out separately at each streamwise column of velocity sample points. This generates a two dimensional correspondence surface, with linear interpolation between the non-linear relationships at each streamwise station, as shown in figure D.10.

When calculating wake sample point velocities, doublet surface influences must be considered on a panel-by-panel basis. At each integration point, the two dimensional mapping is then used to determine geometric coordinates and any geometric gradients required. Whilst this process is quite simple in the interior of the wake, it is problematic at the wake edge where doublet panels span the geometric boundary, such that part of the doublet panel lies on the surface-

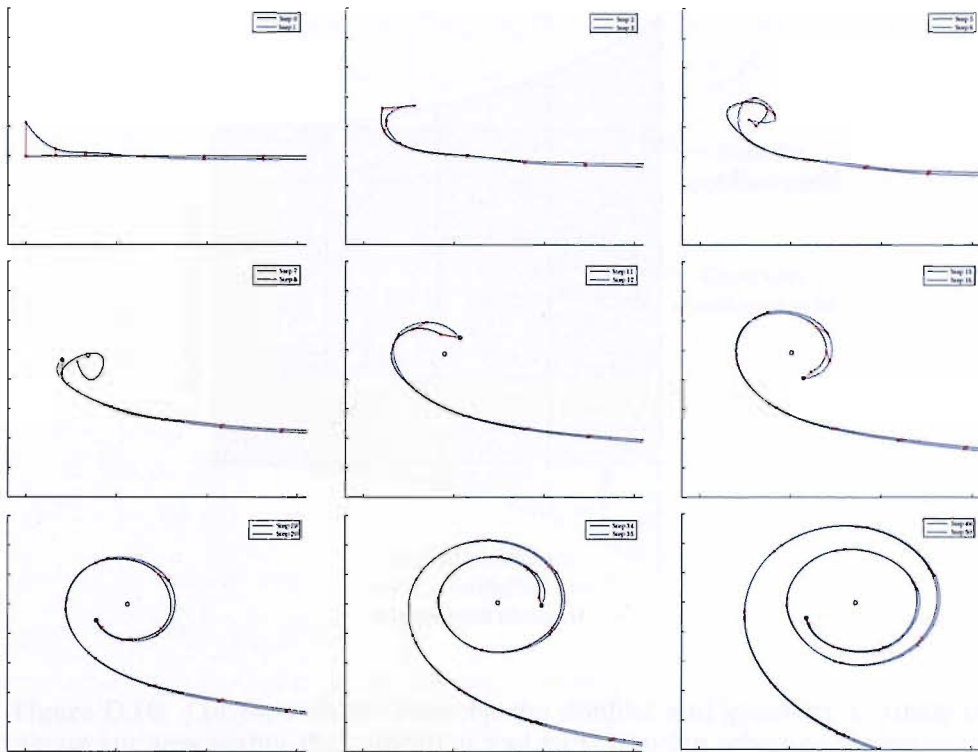


Figure D.8: Selected steps from a wake relaxation developed with the truncation and redistribution scheme. In each image, two steps are shown, the first in black and the second in blue. The wake first becomes truncated in step 8, as shown in the fourth image. The circle in the centre of the roll-up from this step onwards is the concentrated vortex core element.



Figure D.9: The entire wake cross section after 50 relaxation steps of the truncation and redistribution scheme

modelled wake, and part lies on that portion of the wake now modelled as a line vortex element. There are in fact 17 topological conditions that describe how a doublet panel can overlap with the panels of the geometric surface, each requiring separate treatment – some requiring specific sub-divisions of the doublet panels to enable cartesian quadrature.

The truncation and redistribution scheme produces excellent results in two dimensions. However, breaking the direct relationship between the geometry and the singularity surfaces strips the CSM of its simplicity in three dimensions, and

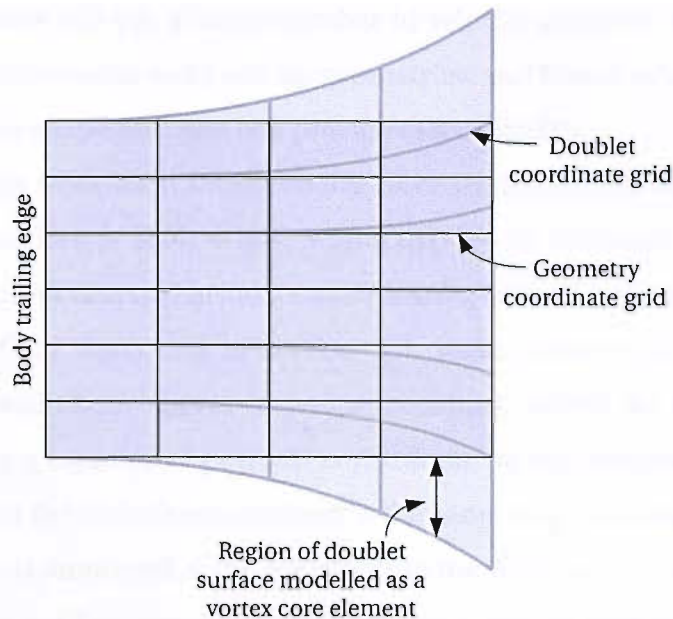


Figure D.10: Correspondence between the doublet and geometry intrinsic coordinate surfaces within the truncation and redistribution scheme. Downstream of the trailing edge of the lifting body, the geometry only spans part of the doublet surface; the vorticity carried by the doublet surface outside of the geometry patch is contained within a vortex core element.

requires numerous extra steps throughout the implementation. Furthermore, whilst in two dimensions it is acceptable to use a relatively large number of spline segments to model the detail of the wake roll-up, computation time in three dimensions requires a much coarser grid. With far fewer panels, the roll-ups shown in figure D.8 are not achievable. Consequently, the scheme was not found to produce a huge improvement over the basic three dimensional CSM relaxation process, and at a cost of much greater complexity.

D.4 Framework shape functions

In common with all higher-order methods, the computational cost of evaluating influence coefficients and wake relaxations is higher for each panel than that for a constant panel method. In order that accuracy for a given computation time can be improved over a low-order method, solutions must therefore comprise fewer panels. If a large density of wake sample points cannot be used to model the

shape of a vortex roll-up, a small number of velocity samples must be capable of producing reasonable wake roll-up geometries, and hence velocity influences. Using a tailored shape function can provide this capability.

Since wakes from most lifting bodies have similar forms, with wake sheets rolling into vortices at their edges, which expand in diameter downstream, a tailored shape function of this form can be employed in the geometric modelling. In the basic CSM described in section 5.2, wake surfaces are approximated as B-spline patches, or curves in two dimensions, where the wake is exactly determined by a number of position constraints. In this modified scheme, the centre region of the wake is constrained in the same way, but a much reduced set of wake points is employed at the edges within the wake roll-up. In order that the geometry does not become under-determined, a ‘roll-up shape-function’ is first fitted to the reduced set of wake points, and the geometry spline is then fitted to this detailed curve. The shape function is designed to approximate standard wake roll-up geometries, but is parameterised by the sparse set of edge wake points in order that it fits the specific situation.

D.4.1 Implementation in two dimensions

The wake geometric spline is divided into a centre region $\xi_1^L \leq \xi_1 \leq \xi_1^R$ which is not expected to roll-up, and edge regions $\xi_1 < \xi_1^L$ and $\xi_1 > \xi_1^R$ which will be controlled by roll-up shape functions. The centre region has widely spaced knots, and a corresponding set of wake sample points which controls its relaxation in the normal way. The edge regions have a higher density of knots, which will allow a more complex roll-up geometry, but have a low density of wake sample points.

During each wake relaxation step, velocities are determined at the sample points, and new positions for them found by forward projection. Specific shape functions (see section D.4.2) are then fitted to the new points within the edge regions, before the whole geometric spline is fitted to the combination of the edge shape functions and the central region’s wake fit points.

An alternative would be to explicitly describe the geometry of the whole wake surface by a parameterised combination of a central B-spline curve and edge

shape functions, removing the need for the extra step of re-fitting a B-spline to the already fitted shape functions. However, this would require more extensive modification of the overall CSM implementation which otherwise is based on all geometries being defined by B-spline curves or surfaces.

D.4.2 Shape functions to model roll-up

Development of an accurate shape function requires further research, which could involve using the truncation and redistribution scheme as a benchmark against which candidate parameterisable models may be tested. For the examples shown here, a simple Archimedean spiral of the form

$$\mathbf{x} = \mathbf{x}_L + \begin{Bmatrix} A\psi \cos(\psi_{0L} + \psi) \\ A\psi \sin(\psi_{0L} + \psi) \end{Bmatrix} \quad (\text{D.6})$$

models the rolled-up part of the wake, where \mathbf{x}_L is the centre of the (left-hand) vortex, A is a constant that controls the diameter of the roll-up, ψ_{0L} is a constant that controls the orientation of the (left-hand) spiral, and ϕ sweeps the range $0 \leq \phi \leq 4\pi$. In each relaxation step, the values of \mathbf{x}_L , A and ψ_{0L} are controlled by: the location of the edge wake fit points; a requirement of continuous curvature where they meet the central spline at $\xi_1 = \xi_1^L$ and $\xi_1 = \xi_1^R$; and the elapsed time since the start of the simulation. A better model might not use simulation time explicitly for the growth of the vortex, but determine it instead from, for example, the divergence of diametrically opposite sample points towards the outside of the vortex.

D.4.3 Results in two dimensions

Figure D.11 shows a selection of relaxation steps created using this scheme, comparing one edge of the wake with that generated using the successful two dimensional truncation and redistribution scheme. It is seen that whilst the two approaches are not perfectly correlated, the spiral shape function, created with only eight velocity evaluations per iteration, is able to approximate the truncation and redistribution scheme's shape which requires around 50 wake sample points.

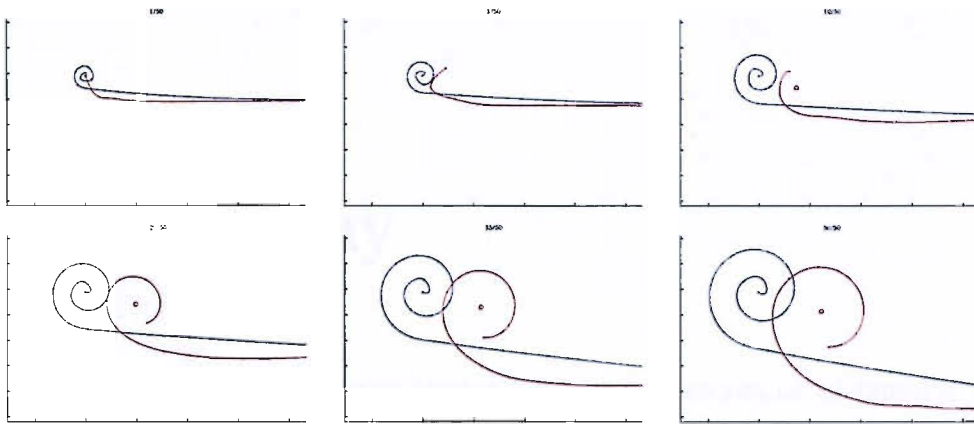


Figure D.11: A selection of steps comparing relaxations with: a spiral tailored shape function (blue), and truncation and redistribution (red). The high density of points in the blue curve shows the number of knots within the geometric spline; the number of wake sample points is much lower.

The Archimedean spiral form was chosen for its simplicity in this example, but is clearly not an accurate representation of a vortex roll-up. In fact, the velocity of the spiral's centre is very dependent on the somewhat arbitrary angle at which the curve approaches the central point, as the doublet strength has a steep gradient here, and has maximum influence on the core itself. In order to obtain reasonable results, this vortex core had to be constrained to move at some average velocity of the points around it. Further research is required to determine how best to sample and weight points within the rolled-up wake in order to estimate an accurate vortex core velocity.

D.4.4 Extension to three dimensions

More detailed analysis of the two dimensional shapes is required before the scheme is implemented in three dimensions, but it is anticipated that the extension from curves to surfaces would be quite straightforward. Shape-functions will be created at each station downstream, before a single surface geometry, with a high density of edge knots, is fitted to the series of the shape-functions and central wake points. An advantage of this method is that once the wake surface is defined, it is treated exactly as a basic CSM surface for influence and velocity computations.

Bibliography

Anderson, R.F. (1936). Determination of the characteristics of tapered wings. Tech. Rep. 572, NACA, Washington D.C.

Bernasconi, D.J. (2007). A dynamic vortex lattice method for unsteady sail aerodynamics. Tech. Rep. SSSU 143, Ship Science, Southampton University, Southampton, England.

Bernasconi, D.J. & Richelsen, P.M. (2008). A higher-order potential flow method for thick bodies, thin surfaces and wakes. *International Journal for Numerical Methods in Engineering*, **73**, 706–727.

Bramesfeld, G. & Maughmer, M.D. (2004). A free-wake, lifting-surface model using distributed vorticity elements. In *22nd Applied Aerodynamics Conference and Exhibit, AIAA Paper 2004-5372, Providence, Rhode Island, 16-19 August 2004*.

Dormand, J.R. & Prince, P.J. (1980). A family of embedded Runge-Kutta formulae. *Journal of Computational Applied Mathematics*, **6**, 19–26.

Engels, H. (1980). *Numerical quadrature and cubature*. Academic Press, London.

Epton, M.A. & Magnus, A.E. (1990). PAN AIR – A computer program for predicting subsonic or supersonic linear potential flows about arbitrary configurations using a higher order panel method. Volume 1 – Theory document (Version 3.0). Tech. Rep. NASA CR-3251, NASA.

Fiddes, S.P. & Gaydon, J.H. (1996). A new vortex lattice method for calculating the flow past yacht sails. *Journal of Wind Engineering and Industrial Aerodynamics*, **63**, 35–59.

Fink, P.T. & Soh, W.K. (1978). A new approach to roll-up calculations of vortex sheets. *Proceedings of the Royal Society of London*, **362**, 195–209.

Friedl, N. (2002). Membrane structures in unsteady potential flow. In *Proceedings of the 5th World Congress on Computational Mechanics*.

Glauert, H. (1959). *The elements of aerofoil and airscrew theory*. Cambridge University Press, 2nd edn.

Guiggiani, M. (1995). Hypersingular boundary integral equations have an additional free term. *Computational Mechanics*, **16**, 245–248.

- Guiggiani, M.** (1998). *Singular integrals in boundary element methods*, chap. 3. WIT Press.
- Guiggiani, M. & Casalini, P.** (1987). Direct computation of cauchy principal value integrals in advanced boundary elements. *International Journal for Numerical Methods in Engineering*, **24**, 1711–1716.
- Guiggiani, M. & Gigante, A.** (1990). A general algorithm for multidimensional cauchy principal value integrals in the boundary element method. *Transactions of the ASME: Journal of Applied Mechanics*, **57**, 906–915.
- Guiggiani, M., Krishnasamy, G., Rizzo, F.J. & Rudolphi, T.J.** (1990). Formulation and numerical treatment of boundary integral equations with hypersingular kernels. *Transactions of the ASME: Journal of Applied Mechanics*, **57**, 906–915.
- Hadamard, J.** (1952). *Lectures on Cauchy's problem in linear differential equations*. Dover, New York.
- Hancock, G.J. & Garner, H.C.** (1975). On the application of subsonic linearised wing theory to second-order forces and moments. Tech. Rep. R&M 3758, Aeronautical Research Council, London.
- Hess, J.L.** (1972). Calculation of potential flow about arbitrary three-dimensional lifting bodies. Final Technical Report MDC J5679-01, McDonnell Douglas, Long Beach, California.
- Hess, J.L.** (1979). A higher order panel method for three-dimensional potential flow. Final technical report, Douglas Aircraft Company, Long Beach, California.
- Hess, J.L. & Friedman, D.M.** (1981). An improved high-order panel method for three-dimensional lifting flow. Tech. Rep. NADC-79277-60, Douglas Aircraft Company.
- Hess, J.L. & Smith, A.M.O.** (1962). Calculation of non-lifting potential flow about arbitrary three-dimensional bodies. Tech. Rep. ES 40622, Douglas-McDonnell.
- Hoeijmakers, H.W.M. & Vaastra, W.** (1983). A higher-order panel method applied to vortex sheet roll-up. *American Institute of Aeronautical Engineering*, **21**, 516–523.
- Horstmann, K.H.** (1987). Ein mehrfach-traglinienverfahren und seine verwendung für entwurf und nachrechnung nichtplanarer flügelanordnungen. Tech. Rep. DFVLR-FB 87-51, Deutsches Zentrum für Luft und Raumfahrt, Braunschweig.
- James, R.M.** (1972). On the remarkable accuracy of the vortex lattice method. *Computational Methods in Applied Mechanical Engineering*, **1**, 59–79.
- Johnson, F.T.** (1980). A general panel method for the analysis and design of arbitrary configurations in incompressible flows. Tech. Rep. NASA CR-3079, NASA.

- Kaden, H.** (1931). Aufwicklung einer unstablen Unstetigkeitsfläche. *Ingenieur Archiv*, **2**, 140–168.
- Kandil, O.A., Chu, L.C. & Tureaud, T.** (1984). A nonlinear hybrid vortex method for wings at large angle of attack. *AIAA Journal*, **22**, 329–336.
- Katz, J. & Plotkin, A.** (2001). *Low Speed Aerodynamics*. Cambridge University Press, 2nd edn.
- Kellogg, O.D.** (1954). *Foundations of potential theory*. Dover, New York.
- Krasny, R.** (1987). Computation of vortex sheet roll-up in the trefftz plane. *Journal of Fluid Mechanics*, **184**, 123–155.
- Labrujere, T.E. & Zandbergen, P.J.** (1973). On the application of a new version of lifting surface theory to non-slender and kinked wings. *Journal of Engineering Mathematics*, **7**, 85–96.
- Lamarre, F. & Paraschivoiu, I.** (1992). Efficient panel method for vortex sheet roll-up. *Journal of Aircraft*, **29**, 28–33.
- Lamb, H.** (1932). *Hydrodynamics*. Cambridge University Press, 6th edn.
- Lan, C.E.** (1974). A quasi-vortex-lattice method in thin wing theory. *Journal of Aircraft*, **11**, 518–527.
- Lee, C.S. & Kerwin, J.E.** (2003). A B-spline higher-order panel method applied to two-dimensional lifting problem. *Journal of Ship Research*, **47**, 290–298.
- Maniar, H.D.** (1995). *A three dimensional higher order method based on B-splines*. Ph.D. thesis, Massachusetts Institute of Technology.
- Maskew, B.** (1982). Prediction of subsonic aerodynamic characteristics: A case for low-order panel methods. *Journal of Aircraft*, **19**, 157–163.
- Moore, D.W.** (1974). A numerical study of the roll-up of a finite vortex sheet. *Journal of Fluid Mechanics*, **63**, 225–235.
- Mracek, C.P., Kim, M.J. & Mook, D.T.** (1992). Three-dimensional potential flows by a vorticity-panel method. *Computers & Fluids*, **21**, 31–42.
- Piegl, L. & Tiller, W.** (1997). *The NURBS book*. Springer-Verlag, New York, 2nd edn.
- Plotkin, A. & Yeh, D.T.** (1986). Vortex panel calculation of wake rollup behind a large aspect ratio wing. *AIAA Journal*, **24**, 1417–1423.
- Pyo, S. & Kinnas, S.A.** (1997). Propeller wake sheet roll-up modeling in three dimensions. *Journal of Ship Research*, **41**, 81–92.
- Rogers, D.F. & Adams, J.A.** (1990). *Mathematical Elements for Computer Graphics*. McGraw-Hill, 2nd edn.
- Rubbert, P.E. & Saaris, G.R.** (1972). Review and evaluation of a three-dimensional lifting potential flow analysis method for arbitrary configurations. In *AIAA Paper 72-188*.

Weise, M.R. (1986). USSAERO version D computer program development using ANSI standard FORTRAN 77 and DI-3000 graphics. Tech. Rep. NASA CR-3980, NASA.

Woodward, F.A. (1968). Analysis and design of wing-body combinations at subsonic and supersonic speeds. *Journal of Aircraft*, **5**, 528–534.

Woodward, F.A. (1973). An improved method for the aerodynamic analysis of wing-body-tail configurations in subsonic and supersonic flow. Tech. Rep. NASA CR-2228, NASA.