

Master of Science Thesis

Entropy-Corrected Full Potential Flow Solver in a Finite Element Framework

Enhanced Representation of Shock Waves in Transonic Flow

Vincent Hoogeboom

December 11, 2024

Entropy-Corrected Full Potential Flow Solver in a Finite Element Framework

Enhanced Representation of Shock Waves in Transonic Flow

Master of Science Thesis

For obtaining the degree of Master of Science in Aerospace Engineering
at Delft University of Technology

Vincent Hoogeboom

December 11, 2024



Delft University of Technology

Copyright © Aerospace Engineering, Delft University of Technology
All rights reserved.

DELFT UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF AERODYNAMICS

The undersigned hereby certify that they have read and recommend to the Faculty of Aerospace Engineering for acceptance the thesis entitled “**Entropy-Corrected Full Potential Flow Solver in a Finite Element Framework**” by **Vincent Hoogeboom** in fulfillment of the requirements for the degree of **Master of Science**.

Dated: December 11, 2024

Supervisors:

Reader 1

Reader 2

Reader 3

Reader 4

Preface

This thesis marks the end of my studies in aerospace engineering. It has been a long and rewarding journey and I am proud to have reached its end.

I would like to thank Roeland de Breuker for his guidance throughout my thesis. Although I feel like I have tried to do things mostly by my self, the support and occasional reassurance that I was on the right path were very much appreciated. Initially, I wanted to start my research in aeroelasticity. However, it concluded in a more aerodynamic-focused thesis, meaning that my interests in aeroelasticity were left unexplored. But who knows what the future will bring.

Additionally, I would like to thank Adrien Crovato for his help throughout the process. His model was easy to work with - which is definitely not a guarantee when talking about computer models - and formed the basis of the work in this thesis.

Vincent Hoogeboom

Summary

This thesis focuses on the development and application of an entropy-corrected full potential flow solver (*EC-Flow*) within a finite element framework, designed to improve the representation of shock waves in transonic flow regimes. By addressing the limitations of traditional isentropic full potential flow solvers, this research aims to enhance the accuracy of aerodynamic predictions at low computational cost, making it suitable for early design phases in aeroelastic analysis.

The work begins with an overview of the theoretical background of transonic aerodynamics and the hierarchy of numerical solutions to the Navier-Stokes equations. It identifies the full potential flow model as a computationally efficient alternative, enhanced with an entropy correction to address its inherent weaknesses in capturing entropy effects across shock waves.

Key components of the methodology include the numerical implementation of the entropy correction within the *Flow* solver to obtain the *EC-Flow* solver, the use of gradient-based and normal Mach number methods for shock detection, and an iterative downstream approach to apply the correction consistently across the domain that is influenced by the shock. The entropy correction is formulated using the Rankine-Hugoniot relations to account for the entropy change across shocks.

The results are demonstrated for two airfoil cases: the symmetric NACA0012 airfoil and the non-symmetric, supercritical RAE2822 airfoil. Convergence analyses show that *EC-Flow* achieves smoother and more stable solutions compared to the uncorrected solver, with a notable reduction in shock strength and improved shock localization. Comparisons with results from the Euler equations validate the effectiveness of the entropy correction in improving solution accuracy without significant computational overhead. Besides improving results, there is also a significant improvement in computational performance.

The analysis concludes with a discussion of the limitations and applicability of the solver, particularly its performance for varying angles of attack. While the entropy correction becomes more pronounced at more extreme angles of attack, the results in these cases should be taken with care, as *EC-Flow* also begins to diverge from the Euler solutions. Recommendations for future work include enhancing the shock detection algorithm, further improving computational efficiency, and extending the solver to three-dimensional transonic flow fields.

Table of Contents

Preface	v
Summary	vii
List of Figures	xiii
List of Tables	xv
Nomenclature	xvii
1 Introduction	1
2 Literature Review	3
2.1 Transonic Flow Regimes	4
2.2 Hierarchy of Different Solutions to the Navier-Stokes Equations	5
2.3 Shock Wave Detection	6
3 Theoretical Background and Research Questions	9
3.1 Steady Isentropic Full Potential Flow Model	9
3.2 Entropy Correction	10
3.3 Research Questions	11
4 Methodology	13

4.1	<i>Flow</i> Numerical Scheme and Implementation	13
4.2	Shock Wave Detection	14
4.2.1	Gradient-Based Detection of Shock Waves	15
4.2.2	The Normal Mach Number	16
4.2.3	Extra Criterion Based on Velocity Decrease	16
4.2.4	Iterative Downstream Implementation	17
5	Results	19
5.1	Domain Convergence	21
5.2	Airfoil Mesh Size Convergence	23
5.3	Far Field Mesh Size Convergence	24
5.4	Remarks on Convergence	25
5.5	Shock Location Determination	27
5.6	Shock Domain of Influence	30
5.7	Pressure Distributions	32
5.8	Polar Data	34
5.8.1	Lift Polar	35
5.8.2	Drag polar	36
5.8.3	Moment Polar	38
5.9	Computational Performance	39
5.10	Sensitivity Analysis	40
5.10.1	Iteration Updates	41
5.10.2	Symmetry Analysis for NACA0012 Airfoil	42
6	Discussion	45
6.1	Effects of the Entropy Correction on Solution Quality	45
6.2	Comparison with State-of-the-Art Models	46
6.3	Numerical Convergence and Stability	46

Table of Contents	xi
7 Conclusion	47
8 Recommendations	49
8.1 Current State of Three-Dimensional Flow Analysis	49
8.2 Smoothing of Convergence Rates	51
8.2.1 Threshold-Based Updates	51
8.2.2 Conditional Field Updates	52
8.2.3 General Remarks	52
8.3 Add Thread Splitting for Entropy Field Update	52
8.4 Adaptive Mesh Refinement	53
Bibliography	55
A NACA0012 Case with Experimental Data	57

List of Figures

2.1	Overview of how local supersonic flow exists on an aircraft wing (Anderson (2017)).	4
2.2	An overview of the different paths towards aerodynamic problem solving (Bendiksen (2011)).	5
5.1	The NACA0012 airfoil.	19
5.2	The RAE2822 airfoil.	20
5.3	Convergence history of the NACA0012 airfoil at $M_\infty = 0.752$ and $\alpha = 1.49^\circ$.	26
5.4	Convergence history of the RAE2822 airfoil at $M_\infty = 0.725$ and $\alpha = 2.4^\circ$.	27
5.5	Overview of the shock wave locations for the NACA0012 airfoil at $\alpha = 1.49^\circ$ and $M = 0.752$.	28
5.6	Zoomed in plot of the NACA0012 airfoil shock locations at $\alpha = 1.49^\circ$ and $M = 0.752$. All elements are shown, overlaid with the elements tagged as shock locations.	28
5.7	Overview of the shock wave locations for the RAE2822 airfoil at $\alpha = 2.4^\circ$ and $M = 0.725$.	29
5.8	Zoomed in plot of the RAE2822 airfoil shock locations at $\alpha = 2.4^\circ$ and $M = 0.725$.	30
5.9	Entropy correction field for the NACA0012 airfoil at $\alpha = 1.49$ and $M_\infty = 0.752$.	31
5.10	Entropy correction domain of influence for the RAE2822 airfoil at $\alpha = 2.4$ and $M_\infty = 0.725$.	32
5.11	Pressure coefficient plot for the NACA0012 airfoil at $\alpha = 1.49$ and $M_\infty = 0.752$.	33
5.12	Pressure coefficient plot for the RAE2822 airfoil at $\alpha = 2.4$ and $M_\infty = 0.725$.	34
5.13	Lift polar plot for the NACA0012 airfoil at $M_\infty = 0.752$.	35
5.14	Lift polar plot for the RAE2822 airfoil at $M_\infty = 0.725$.	36
5.15	Drag polar plot for the NACA0012 airfoil at $M_\infty = 0.752$.	37

5.16	Drag polar plot for the RAE2822 airfoil at $M_\infty = 0.725$	37
5.17	Moment polar plot for the NACA0012 airfoil at $M_\infty = 0.752$	38
5.18	Moment polar plot for the RAE2822 airfoil at $M_\infty = 0.725$	39
5.19	Pressure coefficient plot for the NACA0012 airfoil at $M_\infty = 0.8$ and $\alpha = 0$	42
8.1	Distance based approach in 2D.	50
A.1	NACA0012 airfoil at $M_\infty = 0.752$ and $\alpha = 1.49$ with experimental data.	57

List of Tables

5.1	Domain size convergence for the NACA0012 airfoil.	22
5.2	Domain size convergence for the RAE2822 airfoil.	22
5.3	Airfoil mesh size convergence for the NACA0012 airfoil, with far-field mesh size 0.05.	23
5.4	Airfoil mesh size convergence for the RAE2822 airfoil, with far-field mesh size 0.05.	24
5.5	Farfield mesh size convergence for the NACA0012 airfoil, with airfoil mesh size 0.01.	24
5.6	Farfield mesh size convergence for the RAE2822 airfoil, with airfoil mesh size 0.01.	25
5.7	Load coefficients for the NACA0012 airfoil at $M_\infty = 0.752$ and $\alpha = 1.49$	33
5.8	Load coefficients for the RAE2822 airfoil at $M_\infty = 0.725$ and $\alpha = 2.4$	33
5.9	Computational performance for the NACA0012 case.	40
5.10	Computational performance for the RAE2822 case.	40
5.11	The sensitivity of the NACA0012 airfoil case to the iteration update of the entropy field.	41
5.12	The sensitivity of the RAE2822 airfoil case to the iteration update of the entropy field.	42
5.13	Load coefficients for the NACA0012 airfoil at $M_\infty = 0.8$ and $\alpha = 0$	43
8.1	Load coefficients for the AGARD445 wing at $M_\infty = 0.9$ and $\alpha = 2$	51

Nomenclature

Abbreviations

AMR	Adaptive Mesh Refinement
DNS	Direct Numerical Simulation
NS	Navier-Stokes
RANS	Reynolds-Averaged Navier Stokes

Latin Symbols

\mathbf{n}	Shock normal vector
\mathbf{V}	Velocity vector
a	Speed of sound
d	Distance to velocity vector line
J	Jacobian matrix
M	Mach number
R	Residual vector
R_g	Universal gas constant
s	Entropy

Greek Symbols

Δ	Difference or change in a quantity
$\delta_1\rho$	First-order directional derivative of the density
$\delta_2\rho$	Second-order directional derivative of the density
γ	Specific heat ratio
∇	Gradient operator
ϕ	Velocity potential
ρ_c	Entropy corrected density
ρ_i	Isentropic density

Subscripts

∞	Condition at free-stream
m	Evaluation at iteration m

n Component normal to the surface

Miscellaneous

$\|\cdot\|$ Magnitude of a vector

Chapter 1

Introduction

”Make everything as simple as possible, but not simpler.” These words, famously attributed to Albert Einstein, serve as a reminder that progress often lies in the balance between simplicity and complexity. This idea resonates deeply with the principles of this thesis, which will explore the connection between reality and mathematics to tackle the challenges of transonic flow. Just as Einstein’s wisdom encourages us to avoid over-simplification, the methods and insights presented here strive to capture the essential complexity of the problem.

In today’s aviation industry, there is an endless surge for more sustainable aircraft. A key factor in achieving this goal is to minimise the structural mass of aircraft, leading to more efficient and lightweight designs that reduce energy consumption during flight. A result of this continuous search for more lightweight aircraft structures is that they become increasingly flexible and elastic in response to the aerodynamic forces encountered in flight.

This trend led to the emergence of the relatively new field of aeroelasticity, which is a multidisciplinary area between structural mechanics and aerodynamics. Aeroelasticity not only considers these fields individually but, more importantly, addresses the coupling between them. The increased flexibility of modern aircraft makes aeroelasticity a critical consideration during the early design phases when significant structural and aerodynamic changes can occur. In this context, numerical simulations must be efficient and require low computational costs to support rapid design iterations.

To satisfy the need for computational efficiency, this research focuses on analysis of existing aerodynamic frameworks and exploring potential improvements. Although the emphasis is placed on aerodynamics, the goal is to develop a framework that can be expanded to include aeroelastic analysis. This research specifically targets *transonic* and *high subsonic*, as these are the operating conditions for most commercial aircraft. It should be noted that this will be done in a *steady* framework, meaning that time variations are not included in this research. A detailed discussion of these flight regimes and their challenges can be found in chapter 2.

In aerodynamics, the conservation of mass, momentum, and energy in a continuum form the foundation for solving aerodynamic problems (Bendiksen (2011)). Although there are some exact solutions for specific case problems, to date there is no general solution to these equations (Anderson (2017)). Consequently, numerical methods are used to simulate flow patterns. These methods vary in fidelity, with higher-fidelity approaches, such as Direct Numerical Simulation (DNS) and Reynolds-Averaged Navier-Stokes (RANS), offering greater accuracy, but at a higher computational cost. Through several possible simplifying assumptions, one ends up on the other side of the fidelity spectrum, with solutions like linearised potential flows (Bendiksen (2011)). Depending on the requirements and design phase of a project, an appropriate solution method is chosen.

Balancing the two requirements of low computational cost and accurate modelling of transonic flow problems presents a complex problem. The transonic flow itself is particularly challenging, as it contains mixed subsonic-supersonic flow. This often means it includes shock wave formation within the domain, which introduces non-linear mathematical behaviour and loss of feasibility of multiple simplifying assumptions. The full potential flow model, when combined with an entropy correction across shock waves, offers a promising solution. The full potential flow assumes inviscid, isentropic and irrotational flow. After the entropy correction is implemented across the shock wave, an approximate solution is obtained which tries to lift the isentropic assumption of the full potential flow equations. A more detailed explanation on transonic flow, the hierarchy of the assumptions that can be made, and an explanation on shock wave detection can be found in chapter 2.

Although the full potential flow model poses possible low-cost solutions of transonic flow regimes, it seems that it is not used in that regard in current literature. This is mainly because of its low accuracy in capturing shock effects, which is an effect that is intrinsically connected to transonic aerodynamics, and thus very important to model accurately. However, as it seems from the current literature, a full potential flow solver with a corresponding entropy correction along the shock wave promises good solution accuracy with comparatively low computational costs with regards to the Euler equations and RANS solutions Whitlow (1988). Therefore, this thesis aims to contribute to the aeroelastic analysis of aircraft in early design phases by developing a transonic full potential flow solver with an entropy correction across shock waves, designed to improve shock wave prediction and analysis. More information on the theoretical background of the entropy correction can be found in chapter 3, with its implementation discussed in chapter 4.

The results that the entropy corrected solver, called *EC-Flow*, produces compared to the standard model and Euler equations, can be found in chapter 5. These results are discussed in chapter 6, after which a conclusion is drawn in chapter 7. Finally, chapter 8 discusses the current developments of the entropy correction in three dimensions, and some recommendations for future work.

Chapter 2

Literature Review

To establish an understanding of transonic aerodynamics, it is important to analyse the governing equations. In aerodynamics, the continuity, momentum (also Navier-Stokes (NS)) and energy equations form the governing equations for flow problems which are modelled as a continuum. Although they form a solid model for a variety of fluid flow problems, there is no general analytical solution, making numerical methods necessary. In addition to this, often simplifying assumptions are made to allow for less resource-intensive analyses ([Anderson \(2017\)](#)).

Given the variety of simplifying assumptions, [section 2.2](#) provides an overview of the different approaches to solving aerodynamic problems. As mentioned in [chapter 1](#), the emphasis of this study lies on low-computational cost solutions within the context of transonic flight. Consequently, the current state of literature within this scope is briefly analysed to conclude if there are possible gaps that can be exploited to improve the *aeroelastic* design process. Again, it should be noted that although this section focuses on aerodynamic models, it is the context of *aeroelasticity* in which these models are used that is of interest.

Before discussing the different solution methods, [section 2.1](#) introduces the characteristics of transonic flow. Understanding these complex flow behaviours is critical to identify the requirements of an effective numerical solver.

From the literature review, it becomes evident that an entropy-corrected full potential flow model is a promising candidate for further exploration. The model has been used a few times in previous research, but only in finite-difference models and in application to two-dimensional airfoils. Hence, its intended novelty is to implement it in a finite-element solver and to apply the model on three-dimensional wings.

However, a significant challenge arises from the entropy correction: the need to numerically determine the location of the shock. Determining the location of the shock is not a trivial task, and the numerical methods to find the location of the shock are discussed in [section 2.3](#).

2.1 Transonic Flow Regimes

In chapter 1, it was explained that this text is mainly concerned with low-computational-cost solutions within the framework of transonic flow. To understand what transonic flow entails and what is needed to solve problems within this context, this section provides a small analysis on the characteristics of this flow regime.

Transonic flow owes its name to the fact that it is concerned with flow around an aerodynamic body that includes both subsonic and supersonic regions. Although free-stream velocities may be below sonic speeds, it is possible that through flow acceleration on, for example, an aircraft wing, local supersonic speeds are encountered. This is also visualised in Figure 2.1, where the supersonic region is denoted by a Mach number larger than 1 ($M > 1$). Hence, these problems are related to mixed subsonic-supersonic flow.




Figure 2.1: Overview of how local supersonic flow exists on an aircraft wing (Anderson (2017)).

As the local supersonic flow decelerates through the speed of sound again, often a shock wave occurs. This is characterised by the solid line at the end of the supersonic region. These shock waves cause abrupt changes in pressure, density, and temperature.

Mixed subsonic-supersonic flow has inherently non-linear behaviour, especially in unsteady aeroelastic problems such as flutter (Bendiksen (2011)). Now, since this regime is highly non-linear, non-linear equations are required to model this phenomenon.

One of the primary challenges in modelling transonic flow is the sensitivity of shock waves to small variations in flow conditions, such as Mach number or angle of attack. Even minor changes in these parameters can result in significant shifts in the shock location and strength, which in turn affect aerodynamic performance. Shock waves cause a sharp increase in wave drag, reducing the aerodynamic efficiency of an aircraft.

Now, in aeroelasticity, accurately modelling flow phenomena has a great influence on the solution quality for aeroelastic parameters. Crawley et al. (1995) notes that the location of the shock and its motion are central to understanding the behaviour in transonic regimes. In this work, steady flow is analysed, meaning that the motion need not be tracked. Therefore, the focus lies on the location of the shock. To fully quantify the implications of the possible improvements in this work, the resulting aerodynamic model should be used in a concurrent aeroelastic analysis. This is outside the scope of the current work. However, the model that is applied will be such that subsequent aeroelastic analysis can be performed.

2.2 Hierarchy of Different Solutions to the Navier-Stokes Equations

In section 2.1, the main characteristics of transonic flow were explained, to which this section provides a guide on the different ways to aerodynamic problem solving. Here, it is explained how the Navier-Stokes equations are transformed, through simplifying assumptions, into problems that have reduced complexity. The differences between these solutions are mainly in the fidelity with which one requires to solve the equations. For example, early design phases generally focus on low-fidelity solutions because of low-cost requirements and rapid change of design. Later design phases, however, require higher quality solutions, therefore opting for the higher-fidelity solutions.

However, it should be kept in mind that theoretical transonic aerodynamics is often considered challenging and sophisticated. Hence, a good trade-off between the low-cost requirement, whilst keeping the complexity and therefore the fidelity of the solution, is in place.

In Figure 2.2, one can see a rough schematic of the solution hierarchy to solve the Navier-Stokes equations (Bendiksen (2011)). This overview does not contain all possible methods, only roughly the ones that are more useful in aeroelastic sense. The Navier-Stokes equations are at the top of the hierarchy, and through using several assumptions (described in the round-edged boxes) one ends up with the lowest fidelity solution in the linearized potential equation.

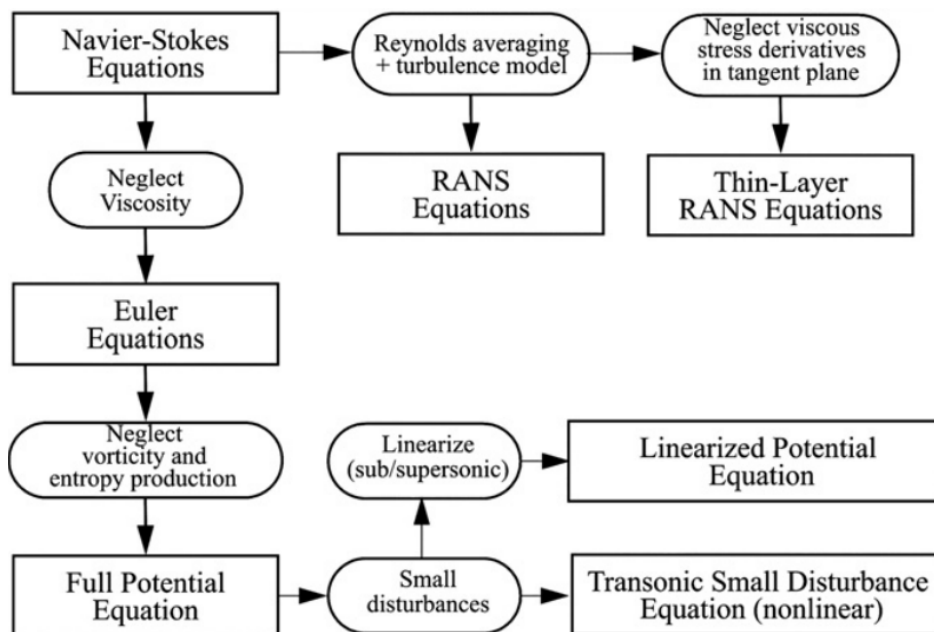


Figure 2.2: An overview of the different paths towards aerodynamic problem solving (Bendiksen (2011)).

The Reynolds Averaged Navier-Stokes (RANS) equations that are depicted are considered

some of the highest fidelity options, though being computationally expensive, especially in changing environments like design optimisation. As this research is focused on less expensive solutions, this leaves one with the other branch of options that are shown in Figure 2.2.

As mentioned in section 2.1, transonic flow exhibits inherently non-linear effects. Therefore, linearized equations such as the Linearized Potential Equation break down, making these methods insufficient in describing the physics of the flow. Not to mention that classical mathematical models for linear theory have a singularity at Mach numbers equal to 1, which is insufficient for transonic regimes.

Now, Bendiksen (2011) mentions that the Full Potential Equation estimates shock waves being too strong and placed too far aft on the airfoil with regard to real solutions, since the solution only accommodates isentropic shocks. Despite these concerns, it is not considered that the FPE can be used alongside an entropy correction or a vorticity correction, as discussed in several research papers (Whitlow (1988), Klopfer and Nixon (1984), Whitlow et al. (1987), Hafez and Lovell (1988), and somewhat more recently in Parrinello and Mantegazza (2012)). The general consensus in these reports is that the entropy correction provides a significant improvement in the results, with the results of the model comparing better to the solutions provided by the Euler equations. Furthermore, Whitlow (1988) shows that the computational cost is around 30-40% higher than the model without the correction, but states that the enhanced shock resolution justifies the extra computational costs. Finally, because of the (partial) inclusion of entropy in the fluid domain, the solver is able to better incorporate strong shocks into the flow field.

Building on the previous analysis, this research continues with the full potential flow equations, enhanced by an entropy correction. The primary contribution of this work to the current research is the implementation of the entropy correction within a finite element solver, with the additional aim of extending its application to three-dimensional flow fields.

The choice of a finite element framework for solving full potential flow equations enables the selection of a suitable software platform to implement the entropy correction. Taking into account the existing collaboration between Delft University of Technology and the University of Liège, the open-source solver *Flow* — developed by the University of Liège and detailed in Crovato (2020) — was chosen. This software, which employs the finite element method to solve the full potential flow equations, has been validated against benchmark cases, making it a reliable foundation for incorporating and testing the entropy correction. To distinguish between the original solver and the modified version with entropy correction, the latter is referred to as *EC-Flow* throughout this research.

2.3 Shock Wave Detection

Implementing the entropy correction introduces the need for shock detection. Shock detection includes determining the location of the shock, its extent from the airfoil boundary into the fluid domain, and its strength. These characteristics are subsequently used to calculate the

entropy correction and enforce this correction on parts of the domain influenced by the shock.

Both [Ma et al. \(1996\)](#) and [Wu et al. \(2013\)](#) break down some general methods to detect a shock wave. The methods consist of (by increasing resolution quality): a threshold based on the Mach number component normal to the shock, a threshold based on the density/pressure gradients, and a method of characteristics. The normal Mach number uses the fact that the component of the Mach number normal to the shock wave crosses 1 over said shock wave. This can then be exploited to detect shock waves within the flow field. Furthermore, density/pressure gradient methods use the fact that density and pressure rapidly change over a shock wave. Finally, the method of characteristics determines the characteristic vector at each grid vertex and, based on a shock-crossing condition, determines if a shock is crossed.

The normal Mach number and gradient-based methods both yield acceptable results, although they may occasionally produce false positives. In contrast, the method of characteristics provides high accuracy but is computationally expensive. According to [Ma et al. \(1996\)](#), a combination of the gradient-based method and a normal Mach number threshold offers a reliable compromise between accuracy and efficiency.

In this research, a combination of the first two methods is used for shock detection. The specific thresholds used for detection will be determined by trial and error, as there are no established guidelines. This approach not only balances accuracy and efficiency but also automatically calculates the Mach number normal to the shock. This normal component is needed to determine the size of the entropy correction, as detailed in chapter 3.

Chapter 3

Theoretical Background and Research Questions

This chapter outlines the mathematical model for steady, nonisentropic full potential flow. First, the isentropic full potential flow model is introduced in section 3.1, along with the simplifying assumptions required to derive the model and its impact on the precision of the solution. Subsequently, the implementation of the entropy correction is detailed in section 3.2, including the formulation of the correction and the additional numerical challenges it poses.

3.1 Steady Isentropic Full Potential Flow Model

In chapter 2, it was concluded that the steady, non-isentropic, full potential flow model will be used in this research. Here, the isentropic full potential flow model is described in more detail. The general full potential equation for steady, three-dimensional flow is given by Equation 3.1. Here, ρ_i is the isentropic density, and ϕ is the velocity potential. These parameters are duly differentiated with respect to spatial coordinates x , y , and z through the gradient operator ∇ . Compared to a two-dimensional flow, the components with regard to z can be ignored.

$$\nabla \cdot (\rho_i \nabla \phi) = 0 \tag{3.1}$$

Here, the density is determined from the Bernoulli relation, which is given in Equation 3.2. Note that this equation is based on the assumption that the incoming flow is uniform. Here, γ is the specific heat ratio for air, and M_∞ is the Mach number in free-stream conditions (far from the analysed body). Furthermore, ϕ_i denotes the derivative of ϕ with respect to the spatial coordinate i , for $i = x, y, z$.

$$\rho_i = \left[1 + \frac{\gamma - 1}{2} (M_\infty^2 - \phi_x^2 - \phi_y^2 - \phi_z^2) \right]^{\frac{1}{\gamma - 1}} \quad (3.2)$$

Important to note is that the full potential flow model is not able to determine loads, as it is irrotational. In *Flow*, this problem is overcome by enforcing the Kutta condition by extending a flat wake sheet from the trailing edge of the airfoil/wing. This wake sheet enables the generation of aerodynamic loads by introducing a necessary discontinuity between the flows on the upper and lower surfaces of the airfoil.

The Kutta condition is a fundamental concept in aerodynamics but is especially highlighted here because it provides another important feature that will be indirectly used by *EC-Flow*. The flat wake sheet serves as a sort of discontinuity in *Flow*, separating flows on the upper and lower sides of the airfoil. In *EC-Flow*, the entropy correction is applied to streamlines behind a shock wave, achieved through downstream iteration from the location of the shock. The wake sheet acts as a natural boundary, ensuring that the downstream iteration remains confined to the respective flow region and does not interfere with the flow on the opposite side of the airfoil — an advantageous and intended side effect.

For a comprehensive discussion of the finite element formulation used in *Flow* and the associated boundary conditions, the PhD dissertation by A. Crovato ([Crovato \(2020\)](#)) offers an in-depth analysis. Given the detailed treatment provided in that work, it was deemed unnecessary to replicate the analysis here.

3.2 Entropy Correction

This section describes the mathematical formulation of the entropy correction and how it is used in aerodynamic analysis. The analysis in this section will focus on a combination of [Klopfer and Nixon \(1984\)](#), [Hafez and Lovell \(1988\)](#), [Whitlow et al. \(1987\)](#), and [Whitlow \(1988\)](#). In chapter 4, a better description of the numerical model will be given to complement the mathematical model in this chapter.

The entropy change across a shock wave can be derived from the Rankine-Hugoniot jump relations, which relate the physical states on both sides of the shock wave. The change in entropy over a shock wave, Δs , can be expressed as a function of the component of the upstream Mach number normal to the shock wave (M_n), the specific heat ratio γ , and the universal gas constant R_g as by Equation 3.3.

$$\frac{\Delta s}{R_g} = \frac{1}{\gamma - 1} \left[\ln \left(\frac{2\gamma}{\gamma + 1} M_n^2 - \frac{\gamma - 1}{\gamma + 1} \right) - \gamma \ln \left(\frac{(\gamma + 1) M_n^2}{(\gamma - 1) M_n^2 + 2} \right) \right] \quad (3.3)$$

This equation captures the irreversible nature of the shock wave and quantifies the increase

in entropy experienced by the fluid particles travelling through a shock wave. M_n is the component of the Mach number normal to the shock, and its value must be tracked across the shock front due to its variation. Here, γ represents the ratio of specific heats, which is approximately 1.4 for air under normal conditions. With the entropy change determined, the isentropic density ρ_i determined behind a shock wave can be corrected for using this value, as shown in Equation 3.4.

$$\rho_c = \rho_i e^{\frac{-\Delta s}{R_g}} \quad (3.4)$$

In this equation, ρ_i represents the isentropic density behind the shock wave, as determined by the full potential flow equations. This value is adjusted to obtain the corrected density ρ_c . For example, for an M_n of 1.3 (a strong shock), the correction value $e^{-\frac{\Delta s}{R_g}}$ is roughly equal to ~ 0.98 . These corrections are small, highlighting the importance of accurate shock resolution.

The entropy correction thus reduces the density magnitude behind the shock, thereby decreasing the density jump across the shock wave and weakening its strength. As discussed in chapter 2, the full potential flow equations typically overestimate the shock strength and position the shock too far aft. The entropy correction addresses these inaccuracies by indirectly reducing the predicted shock strength.

While the entropy correction is straightforward in its formulation, its practical implementation presents certain challenges. Significant challenges are the shock detection and the determination of the upstream component of the Mach number normal to the shock wave. In transonic and supersonic flows, the location and strength of the shocks can vary depending on the geometry and flow conditions. To correctly apply the entropy correction, the shock location must be identified, and the upstream Mach number must be tracked for the complete shock front.

The determination of the Mach number component normal to the shock wave inherently enables the analysis of oblique shock waves. While the curvature of the shock front is expected to have only a minor influence on the cases analyzed in this research, it remains an important consideration.

3.3 Research Questions

The previous analysis concludes with the following research questions and sub questions that are to be answered in the upcoming research.

1. What are the effects of the incorporation of an entropy correction in a full potential flow solver for aerodynamic problem solving?

- How does one implement the entropy correction in an existing, finite element, full potential flow solver?
 - How does the entropy correction impact the solution quality with respect to standard full potential flow solvers?
2. How does the full potential flow solver with entropy correction compare to state-of-the-art models?
 - What are the differences in computational efficiency compared to the standard full potential equation and the Euler equations?
 3. What are the limitations of the full potential flow solver with entropy correction for aerodynamic problem solving in aeroelastic context?
 - What ranges of angle of attack is this method most suited for?

Chapter 4

Methodology

This chapter outlines the numerical application of the entropy correction model, referred to as *EC-Flow*, within the framework of the standard *Flow* solver. To establish the foundation for this integration, section 4.1 provides an overview of the numerical scheme and the details of the implementation of *Flow*. This overview serves as the basis for the subsequent description of the numerical integration of the entropy correction.

The numerical methods for detecting shock waves are detailed in section 4.2. Shock detection presents a significant challenge, especially when computational efficiency must be preserved. This section outlines the complexities involved in identifying shocks and describes the approach adopted to address these challenges within the context of this study.

Finally, subsection 4.2.4 discusses the application of the entropy correction to the flow field downstream of the shock wave. This process presents the greatest challenge, as different correction values need to be applied to different streamlines. Since the finite-element mesh typically does not align perfectly with the streamlines, numerical inconsistencies can creep into the solver.

4.1 *Flow* Numerical Scheme and Implementation

This section provides a concise overview of the numerical schemes employed in *Flow*, with a focus on the quasi-Newton algorithm, where the entropy correction is applied. Detailed information on the implementation of the finite element method is beyond the scope of this work and is therefore not included here. For a comprehensive discussion on the implementation of the finite element method, readers are referred to [Crovato \(2020\)](#).

The finite element discretized weak form of the full potential flow equation can be rewritten as a set of equations, $R = 0$, with R the residual vector. The weak form is derived by

integrating the governing equations against a set of test functions and is central to the finite element formulation. Solving the full potential flow equation for transonic flows, particularly using the finite element method, requires a robust numerical scheme to address the non-linearities inherent in compressible flow. Consequently, the system of non-linear equations must be solved iteratively.

In [Crovato \(2020\)](#), two primary iterative techniques are employed to solve these non-linear equations: the Picard iteration and the Quasi-Newton method. As the Picard method is used primarily for debugging and has poor convergence characteristics, the focus of this work lies on the Quasi-Newton algorithm, which offers superior convergence for transonic flow problems.

Ignoring higher-order terms and assuming that a solution is available in iteration m , a better estimate of the solution can be found by Equation 4.1.

$$J_m(\phi_{m+1} - \phi_m) = -R_m \quad (4.1)$$

Here, J_m denotes the Jacobian matrix in iteration m , which is defined by Equation 4.2. The Jacobian matrix represents the matrix of first-order partial derivatives of the residual vector R_m with respect to the solution ϕ_m .

$$J_m = \left. \frac{\partial R}{\partial \phi} \right|_{\phi_m} \quad (4.2)$$

The quasi-Newton solution effectively solves for Equation 4.1. The idea behind Equation 4.1 is to linearize the non-linear problem locally with the current solution ϕ_m . By solving this linearized system, one obtains a correction $\phi_{m+1} - \phi_m$ that reduces the residual R_m . The Jacobian matrix J_m helps determining the appropriate size and direction of this correction.

The entropy correction must be incorporated into the residual vector R_m and the Jacobian matrix J_m . This integration involves modifying the underlying density variables for specific elements in the computational domain. The solver must identify the finite elements where the entropy correction is required and apply the corresponding correction values. To achieve this, all entropy correction values are pre-computed for the entire flow domain before the next iteration begins. These precomputed values are then utilized during the solution updates, as outlined earlier, ensuring consistency and computational efficiency.

4.2 Shock Wave Detection

As discussed in section 2.3, gradient-based methods and the normal component of the Mach number relative to the shock wave are commonly used to determine the location of the shock.

This section elaborates on the methodology used to identify shock waves in the flow field. The approach combines density gradient analysis, detailed in subsection 4.2.1, with Mach number thresholds, described in subsection 4.2.2. Furthermore, subsection 4.2.3 introduces an auxiliary selection criterion based on a velocity parameter. This parameter helps to exclude regions of the flow that are not relevant to the shock detection process.

4.2.1 Gradient-Based Detection of Shock Waves

As discussed in section 2.3, the methodologies presented by [Ma et al. \(1996\)](#) and [Wu et al. \(2013\)](#) are utilized to detect shock waves in the flow field. The detection process leverages the first-order directional derivative of the density, $\delta_1\rho$, and the second-order directional derivative, $\delta_2\rho$. These derivatives are instrumental in identifying regions of rapid density variation, such as those found in shock waves. By analysing how the density changes along the local velocity vector, these methods enable precise localization of shocks in the flow field.

In Equation 4.3, one can find the first-order directional derivative of the density. Here, $\frac{\mathbf{V}}{\|\mathbf{V}\|}$ is the normalised velocity vector and $\nabla\rho$ the density gradient.

$$\delta_1\rho = \frac{\mathbf{V}}{\|\mathbf{V}\|} \cdot \nabla\rho \quad (4.3)$$

This criterion is informative for shock locations, as the maximum of this variable is attained at the shock location. Moreover, if the value is positive or larger than zero, it means that the flow is compressing. This means that the velocity and the density gradient are oriented in the same direction. If it is negative, it means that the flow is expanding. This criterion thus also provides a distinction between expansion and compression waves.

In practical applications, identifying the maximum values of $\delta_1\rho$ provides an initial indication of the locations of the shock, marking the regions of drastic density changes. To find these maximum values, the second order derivative $\delta_2\rho$ is used and is given by Equation 4.4. Determining the maxima of $\delta_1\rho$ is the same as determining the zero values of $\delta_2\rho$.

$$\delta_2\rho = \frac{\mathbf{V}}{\|\mathbf{V}\|} \cdot \nabla(\delta_1\rho) = \frac{\mathbf{V}}{\|\mathbf{V}\|} \cdot \nabla\left(\frac{\mathbf{V}}{\|\mathbf{V}\|} \cdot \nabla\rho\right) \quad (4.4)$$

Although $\delta_2\rho$ is zero at shock waves, it can also reach zero in smooth regions of the flow, resulting in spurious shock wave locations. This requires additional criteria to differentiate between these cases.

In addition to adding additional parameters to the threshold with, the values for $\delta_1\rho$ and $\delta_2\rho$ can also be used solely. Then, one has to threshold the $\delta_1\rho$ value by $\delta_1\rho > \epsilon$, as mentioned

by Wu et al. (2013). This means that this value has to be user-selected. This loses some generality and ease of use of *Flow*, and introduces a new very sensitive parameter. Therefore, it was chosen to add extra parameters instead.

4.2.2 The Normal Mach Number

As mentioned, it is essential to couple the gradient parameters with the *normal Mach number* to ensure that shock detection only occurs in regions where supersonic flow decelerates to subsonic speeds. The Mach number component normal to the shock M_n is defined by Equation 4.5.

$$M_n = \frac{\mathbf{V} \cdot \mathbf{n}}{a} \quad (4.5)$$

Here, \mathbf{V} is the velocity vector, \mathbf{n} is the local surface normal, and a is the speed of sound under local conditions. The local surface normal is determined by Equation 4.6. This is equal to the normalized density gradient.

$$\mathbf{n} = \frac{\nabla \rho}{\|\nabla \rho\|} \quad (4.6)$$

The normal Mach number provides a critical threshold mechanism for shock detection. During analysis, elements are evaluated to determine whether $M_n > 1$, hence only select regions where the flow has supersonic components normal to the local surface. This ensures that the detection mechanism accurately identifies the shock locations, as these regions correspond to supersonic deceleration indicative of a shock wave. However, even if this method is coupled with gradient-based parameters, it is also mildly sensitive to spurious selection of slightly upstream locations of the shock. Therefore, another parameter was used, as discussed in subsection 4.2.3.

4.2.3 Extra Criterion Based on Velocity Decrease

During the shock determination process, an additional selection criterion was found to be helpful in deselecting upstream parts of the domain. The criteria is given by Equation 4.7, and was also referenced by Wu et al. (2013) as a useful tool.

$$\mathbf{V} \cdot \nabla \|\mathbf{V}\| < 0 \quad (4.7)$$

If $\mathbf{V} \cdot \nabla \|V\|$ is smaller than zero, it means that the flow is decelerating in the direction of the flow, which is expected across a shock wave. This criterion helps in deselecting upstream parts as the velocity is increasing there in the direction of the flow.

Combining this method with the other two provides a solid basis for shock wave detection. Although being of major importance, the implementation of the correction that follows from this shock wave location is arguably even more important. This is discussed in subsection 4.2.4.

4.2.4 Iterative Downstream Implementation

Once the shock wave is localized, the entropy correction is calculated using the Mach number component normal to the shock wave. The implementation of the entropy correction involves traversing downstream over streamlines. This is by finding neighbouring elements for which their relative position is most in streamwise direction (in the velocity vector direction). This approach ensures that the entropy change remains consistent along each streamline, as the shock-induced entropy change is constant for streamlines passing through the shock.

The iteration downstream continues until one of the following conditions is met: the edge of the domain is reached, no downstream element can be identified, or a predefined iteration counter is exceeded (to ensure convergence).

It should be noted that, when the solver has reached a downstream element, it also applies the entropy correction to its direct neighbouring elements. This is done to ensure that all elements downstream of the shock are selected and corrected for using the appropriate entropy correction. In practice, when not adding those neighbouring elements, some elements are missed whilst adding the entropy correction. The implications of this are discussed in chapter 5.

Chapter 5

Results

This chapter presents the results of the influence of the entropy correction within the framework of *Flow*. The analysis includes two cases: a NACA0012 airfoil with a moderate strength shock and an RAE2822 airfoil with a strong shock. The NACA0012 airfoil was selected for its widespread use in aerodynamic testing and validation and is analysed at a Mach number of $M_\infty = 0.752$ and an angle of attack of $\alpha = 1.49$. Its generality also facilitates the inclusion of experimental data for the pressure coefficient plots, providing an additional layer of validation. The NACA0012 airfoil is represented in Figure 5.1.

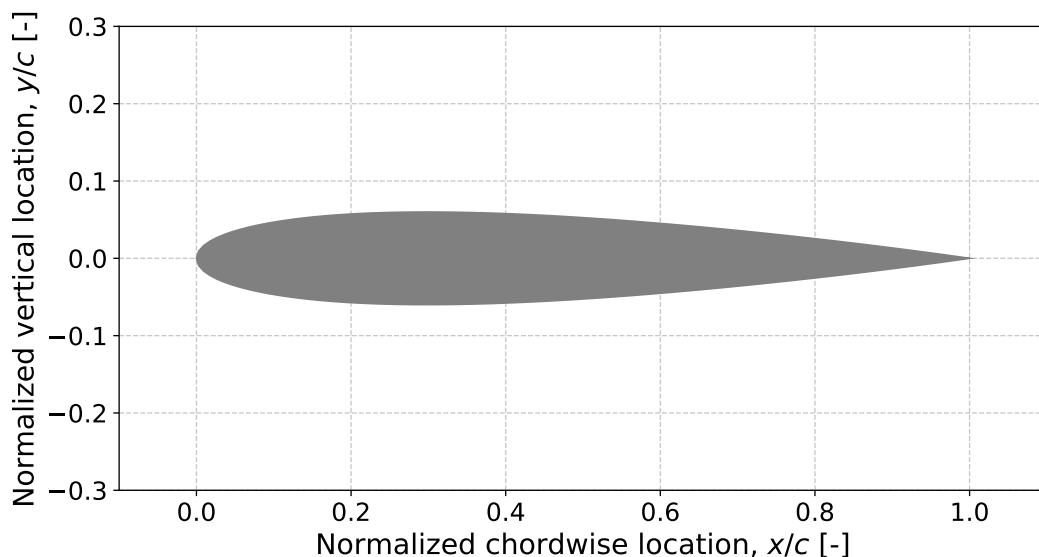


Figure 5.1: The NACA0012 airfoil.

The RAE2822 is a nonsymmetric supercritical design and was chosen as another commonly used validation model. This case is examined with a Mach number of $M_\infty = 0.725$ and an angle of attack of $\alpha = 2.4$. The airfoil can be found in Figure 5.2.




Figure 5.2: The RAE2822 airfoil.

The results are presented systematically, beginning with a convergence analysis of the domain size, airfoil mesh size, and far-field mesh size. Although the standard *Flow* solver demonstrates convergence for these parameters, it is essential to verify whether *EC-Flow*, which incorporates the entropy correction, exhibits similar behaviour. The entropy correction, applied to the flow field as it passes through the shock wave, introduces a potential sensitivity to these parameters. This necessitates a more detailed examination of the convergence characteristics of these parameters, albeit for different reasons.

The size of the domain ensures general convergence and determines the required extent of the computational domain. As the entropy correction is applied throughout the entire fluid stream that interacts with the shock wave, the domain size directly influences the extent of the region where the correction is active, making its convergence analysis particularly significant. The domain size convergence will be analysed in section 5.1. Furthermore, the airfoil mesh size is equally important as it directly impacts the shock wave resolution near the body of the airfoil. Since the entropy correction relies on accurately identifying the location of the shock wave, it is critical to determine whether this parameter significantly affects the solution. It will be analysed in section 5.2. Then, the far-field mesh density influences the resolution in the regions where the entropy correction is applied. Assessing the convergence of this parameter is essential not only for general solver stability but also for identifying whether *EC-Flow* imposes additional requirements on far-field mesh refinement compared to *Flow*. This will be addressed in section 5.3.

The convergence analysis concludes with remarks on the general convergence in section 5.4. From the convergence characteristics of the three mesh defining parameters between the two models, it will become clear that *EC-Flow* seems to exhibit a stronger convergence behaviour than *Flow*. This is analysed in slightly more detail in this section, by addressing the general convergence of the models. By examining the general convergence of the two models, this

section provides a bit more detail on the influence of the entropy correction on the general convergence.

Following the convergence analysis, the results of the shock location determination and the shock domain of influence are presented. These aspects form the numerical implementation of the entropy correction. Especially the domain of influence analysis highlights rather specifically where the corrections are applied, and the values it uses.

Subsequently, the impact of these corrections is demonstrated through pressure coefficient plots and load coefficients for the two cases. To provide a broader analysis, the lift, drag, and moment polar plots are provided for the NACA0012 and RAE2822 airfoils. This generalization not only illustrates the effects of the entropy correction more comprehensively but also reveals the solver's enhanced capability to handle a slightly extended range of angles of attack when the correction is applied.

The foundation of the results section is the comparison between the standard *Flow* model and its entropy-corrected extension, *EC-Flow*. To reiterate, the standard model *Flow* assumes irrotational, and isentropic flow when compared to the Euler solution. In contrast, *EC-Flow* tries to lift the isentropic assumption by applying a correction. Hence, to check the validity, and the influence of this correction, it is best to also compare these results with the Euler solution. For this purpose, Euler solution results will be presented in this chapter. All Euler results were generated using the open-source software of SU2 (Economon et al. (2016)).

In conclusion, the convergence analysis uses four key parameters: the lift coefficient C_l , the drag coefficient C_d , the shock location normalized by the chord length x_s/c , and the shock strength. Now, the shock strength is quantified by the jump in pressure coefficient of the coefficient on the airfoil body. This is by no means the most accurate metric for this parameter, but it only serves as an indicator as to how the shock strength changes with mesh fidelity and also between the two models.

5.1 Domain Convergence

Analysing domain convergence is essential to ensure that the solver reaches a stable solution as domain size increases. If the domain boundaries are placed too close to the airfoil, they may interfere with the solution by imposing artificial constraints on the flow field. Conversely, extending domain boundaries too far can unnecessarily increase computational cost without improving solution accuracy significantly.

In addition to these general considerations, the domain size requires particular care for *EC-Flow*, because the correction is applied on the flow field downstream of the shock. Therefore, the size of the domain directly impacts the extent of the application of the correction. This means that analysing the domain size not only demonstrates the solver's convergence, but can also highlight potential differences in domain size requirements for the two solvers. Thus, it will ensure that the entropy correction is applied accurately.

For the domain convergence analysis, an airfoil mesh size of 0.01 and a far-field mesh size of 0.5 were selected. These values were found to produce the best results, particularly for *EC-Flow*, as the refinement they provide in the shock domain of influence is critical for ensuring solution quality. The importance of these mesh sizes and their impact on the solution will be further detailed in section 5.2 and section 5.3, respectively.

In Table 5.1, one can find the convergence parameters for the NACA0012 airfoil case. Note that this airfoil was analysed at a Mach number of $M_\infty = 0.752$ and an angle of attack of $\alpha = 1.49$. The domain size itself is given in chord lengths on the x- and z-axis.

Table 5.1: Domain size convergence for the NACA0012 airfoil.

Domain size	C_l		C_d		Shock loc. x_s/c (shock strength)	
	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>
7 x 6	0.428	0.389	0.0095	0.0080	0.53 (1.07)	0.49 (0.98)
13 x 12	0.394	0.369	0.0076	0.0069	0.50 (1.13)	0.46 (1.00)
21 x 20	0.393	0.366	0.0074	0.0067	0.50 (1.15)	0.45 (1.01)

The results indicate that both models exhibit similar convergence behaviour, with a domain size of approximately 13×12 chord lengths being sufficient to achieve stable solutions. Importantly, this domain size provides a balance between capturing the essential flow physics and avoiding unnecessary computational costs. Furthermore, this domain provides better resolution of the shock wave, as shown by the shock location and shock strength

A first implication on what the entropy correction does to the solution can also be seen; it seems to reduce the shock strength, and place the shock more towards the leading edge. These changes also influence the C_l and C_d , as these values experience a drop.

For the non-symmetric RAE2822 airfoil, the convergence parameters are presented in Table 5.2. The dashes for *Flow* at the domain size of 7x6 means that the solution did not converge.

Table 5.2: Domain size convergence for the RAE2822 airfoil.

Domain size	C_l		C_d		Shock loc. x_s/c (shock strength)	
	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>
7 x 6	-	1.104	-	0.027	-	0.73 (1.45)
13 x 12	1.100	0.959	0.022	0.012	0.74 (1.56)	0.64 (1.20)
21 x 20	1.076	0.955	0.019	0.012	0.72 (1.55)	0.64 (1.18)

As observed with the NACA0012 case, both *Flow* and *EC-Flow* exhibit similar convergence behaviour, requiring a domain size of approximately 13×12 chord lengths for stable solutions. This size represents again the balance between capturing the essential flow dynamics and minimizing computational cost.

However, a notable difference is the slightly stronger convergence exhibited by *EC-Flow*, particularly for the lift and drag coefficients. The parameters for *EC-Flow* demonstrate less variation when transitioning from the 13×12 domain to the 21×20 domain. This suggests

that *EC-Flow* achieves a more stable and smoother solution.

The stronger convergence of *EC-Flow* can most likely be attributed to its prediction of a weaker shock. This results in a smoother flow field with reduced solution jumps in the shock region, improving overall stability and reducing sensitivity to domain size.

Further insight into the convergence behaviour is provided in section 5.4, where the general iteration requirements and solution quality are analyzed for both solvers. This section highlights the number of iterations required for convergence and provides a comparative assessment of the stability and accuracy of *Flow* versus *EC-Flow*.

5.2 Airfoil Mesh Size Convergence

The size of the mesh at the airfoil plays an important role in the fidelity of the solution on the airfoil boundary. As mentioned, the shock location needs to be determined numerically to be able to apply the entropy correction on the flow field. Since the shock wave is strongest near the airfoil boundary, the numerical determination of its location is probably dependent on the quality and density of the mesh in this region.

In Table 5.3, the convergence values for the NACA0012 airfoil mesh size are shown.

Table 5.3: Airfoil mesh size convergence for the NACA0012 airfoil, with far-field mesh size 0.05.

Airfoil mesh size (No. of elements)	C_l		C_d		Shock loc. x_s/c (shock strength)	
	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>
0.02 (5700)	0.391	0.364	0.0075	0.0067	0.47 (1.08)	0.44 (0.96)
0.015 (9700)	0.393	0.369	0.0076	0.0068	0.48 (1.09)	0.45 (0.99)
0.01 (10800)	0.394	0.369	0.0076	0.0069	0.50 (1.13)	0.46 (1.01)

Both versions of *Flow* exhibit the same convergence pattern. Already at an airfoil mesh size of 0.015, both solvers seem to produce comparatively converged results. However, especially for shock resolution, in terms of strength and location, it is advised to use an airfoil mesh size of 0.01. This only comes at a cost of slightly more fluid domain elements (10800 compared to 9700) and produces a better shock prediction. Moreover, the coarser meshes introduce some slight shock smearing due to a slight lack of mesh resolution. Again, especially for *EC-Flow*, the refinement of the shock is of importance.

Table 5.4 shows the values for the RAE2822 airfoil. This table shows again that both solvers exhibit the same general behaviour.

Table 5.4: Airfoil mesh size convergence for the RAE2822 airfoil, with far-field mesh size 0.05.

Airfoil mesh size (No. of elements)	C_l		C_d		Shock loc. x_s/c (shock strength)	
	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>
0.02 (5700)	1.041	0.942	0.015	0.013	0.70 (1.48)	0.62 (1.18)
0.015 (9800)	1.058	0.957	0.019	0.013	0.72 (1.52)	0.63 (1.19)
0.01 (11100)	1.100	0.959	0.022	0.012	0.74 (1.56)	0.64 (1.20)

5.3 Far Field Mesh Size Convergence

The far-field mesh size directly influences the fidelity of the shock domain of influence. Accurate resolution in this region is crucial to ensure that the streamlines that pass a shock wave are refined enough, and therefore correctly analysed. A poorly resolved far-field mesh can lead to inaccuracies in capturing the flow behaviour downstream of the shock, potentially compromising the effectiveness of the correction. If needed, to get a better feel of how the entropy correction is applied in the flow field, it is referred to section 5.6. There, figures illustrate the domain of influence of the entropy correction on the streamlines of the flow field, and its values. This can create a better feeling for how the far field mesh size can influence the solution behaviour.

The values for the far field mesh size convergence of the NACA0012 airfoil convergence are shown in Table 5.5.

Table 5.5: Farfield mesh size convergence for the NACA0012 airfoil, with airfoil mesh size 0.01.

Farfield mesh size (No. of elements)	C_l		C_d		Shock loc. x_s/c (shock strength)	
	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>
1.5 (6400)	0.356	0.347	0.0069	0.0061	0.48 (1.02)	0.42 (0.93)
1 (7400)	0.387	0.354	0.0071	0.0062	0.49 (1.07)	0.43 (0.95)
0.5 (10800)	0.394	0.369	0.0076	0.0069	0.50 (1.13)	0.46 (1.00)
0.35 (18300)	0.400	0.373	0.0083	0.0075	0.51 (1.17)	0.46 (1.02)

Although both models seem to convergence, it seems like *EC-Flow* requires a far-field mesh size of 0.5 for accurate shock resolution. Its values, especially in shock location and lift coefficient, seem to change significantly between far-field mesh sizes of 1 and 0.5

The far-field mesh size convergence values for the RAE2822 airfoil are presented in Table 5.6. Both models again exhibit the same general behaviour, but here the difference between a far-field mesh size of 1 and 0.5 is not particularly evident for *EC-Flow*.

For *EC-Flow*, it is advised to use a far-field mesh of 0.5, to provide the best shock resolution, and to provide the best overall convergence results, as analysed in section 5.4.

Table 5.6: Farfield mesh size convergence for the RAE2822 airfoil, with airfoil mesh size 0.01.

Farfield mesh size (No. of elements)	C_l		C_d		Shock loc. x_s/c (shock strength)	
	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>	<i>Flow</i>	<i>EC-Flow</i>
1.5 (6800)	1.048	0.956	0.017	0.012	0.70 (1.51)	0.62 (1.18)
1 (7900)	1.079	0.959	0.020	0.012	0.72 (1.54)	0.62 (1.19)
0.5 (11100)	1.100	0.959	0.022	0.012	0.74 (1.56)	0.64 (1.20)
0.35 (18200)	1.166	0.976	0.029	0.013	0.76 (1.59)	0.66 (1.22)

5.4 Remarks on Convergence

The previous analyses indicate that, in general, using a domain of roughly 13x12 chord lengths, an airfoil mesh size of 0.01, and a far-field mesh size of 0.5 are best for usage of *EC-Flow*.

It also became evident that, especially for the RAE2822 case, the solution seems to stabilise more effectively for *EC-Flow* compared to *Flow*. This is probably attributed to the lower predicted shock strength, allowing for smoother flow. It does raise the question of how this influences the overall convergence of the solver. To investigate this further, the general convergence rates of the cases are analysed here, or the relative residual values. This provides a little more insight into the solvers' efficiency and stability in achieving a stabilised solution. The relative residuals convergence for the NACA0012 and RAE2822 airfoils can be found in Figure 5.3 and Figure 5.4, respectively.

The convergence plot for the NACA0012 airfoil highlights several features of the models and illustrates notable differences in their behaviour. The first observation is that *EC-Flow* converges faster than *Flow* in iterations (17 compared to 22). It should be noted that *EC-Flow* does require extra computations for the entropy correction, hence this does not necessarily mean that it has faster computational performance. Its computational performance will be analysed in section 5.9.

Moreover, the first two times the relative residuals converge to a value below 10^{-2} , the viscosity damping parameter is altered (which was already present in the standard solver), which can best be seen in the standard model (the two increases in relative residual values at iterations 10 and 13). This is also apparent for *EC-Flow*, at iterations 9 and 12.

For *EC-Flow*, the entropy correction field is updated at every sixth iteration. This was determined by trial and error, and the sensitivity of this parameter is analysed in subsection 5.10.1. The updates of the entropy field correction in *EC-Flow* is specially seen here at iteration 6, the first update. After this, the two solver exhibit different convergence values.

Furthermore, for the RAE2822, the influence of the entropy correction field update is more clear, as seen in Figure 5.4.

First of all, *EC-Flow* converges significantly faster in iterations than *Flow*. Again, although converging faster in iteration count, *EC-Flow* needs to determine the entropy correction field at every update it places. Therefore, this does not necessarily mean that it converges faster

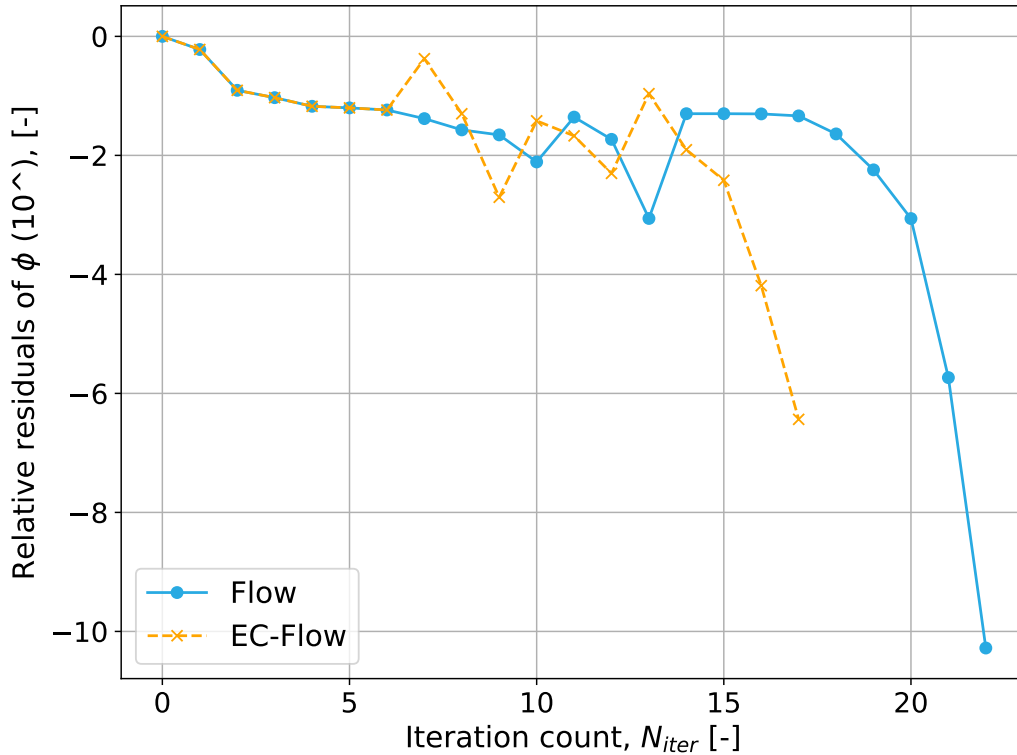


Figure 5.3: Convergence history of the NACA0012 airfoil at $M_\infty = 0.752$ and $\alpha = 1.49$.

in actual computational time. The actual computational time for the two cases is analysed in section 5.9.

In Figure 5.4, the spikes at iterations 15 and (partially) 18 can be attributed to the update of the viscosity damping parameter. However, the spikes at 6, 12, and (partially) 18, are attributed to the updates of the entropy correction field. These oscillations pose a negative behaviour on the convergence, causing some leakage of performance. Smoothing of the residuals, resulting in a better performance of the correction updates, is a recommendation for further research. Especially as it seems that the convergence rate looks fast already at earlier iteration counts, and the entropy field updates cause a bounce-back in residual values. This compares to the rather smooth behaviour of the original *Flow* model.

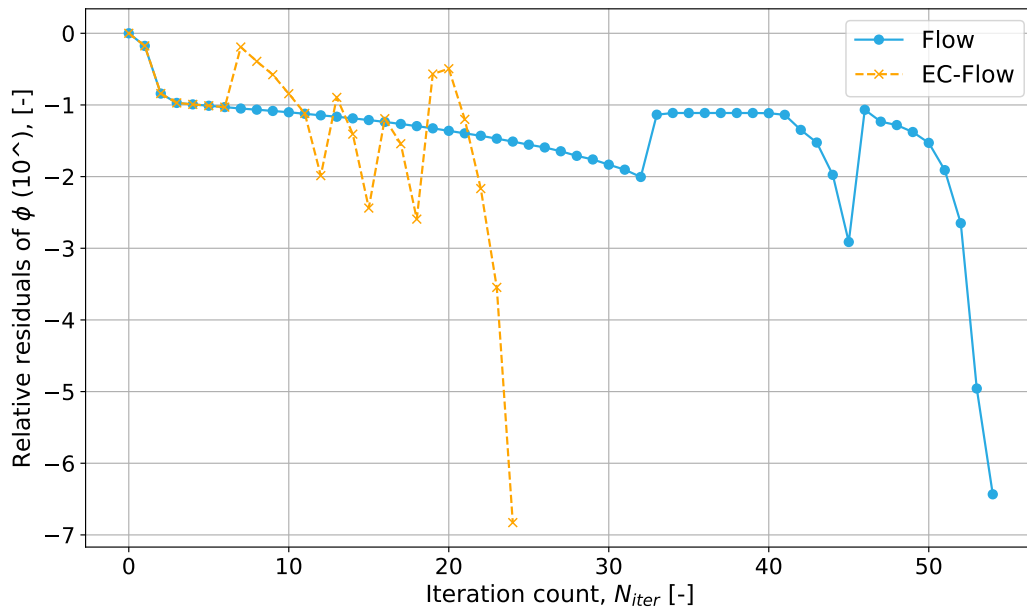


Figure 5.4: Convergence history of the RAE2822 airfoil at $M_{\infty} = 0.725$ and $\alpha = 2.4$.

5.5 Shock Location Determination

Following the detailed and somewhat tedious convergence analysis, the shock location determination is now analysed. An overview of the shock location determined for the NACA0012 airfoil is presented in Figure 5.5. Overall, the selection of shock elements in this simulation performs well, accurately identifying most of the elements within the shock region on the upper surface of the airfoil. However, near the top of the shock wave, the resolution becomes coarser and some elements appear to be missed.

These omissions are better showcased in Figure 5.6. In this plot, the shock location is examined more closely, with the centres of gravity of all elements displayed alongside the elements identified as part of the shock.

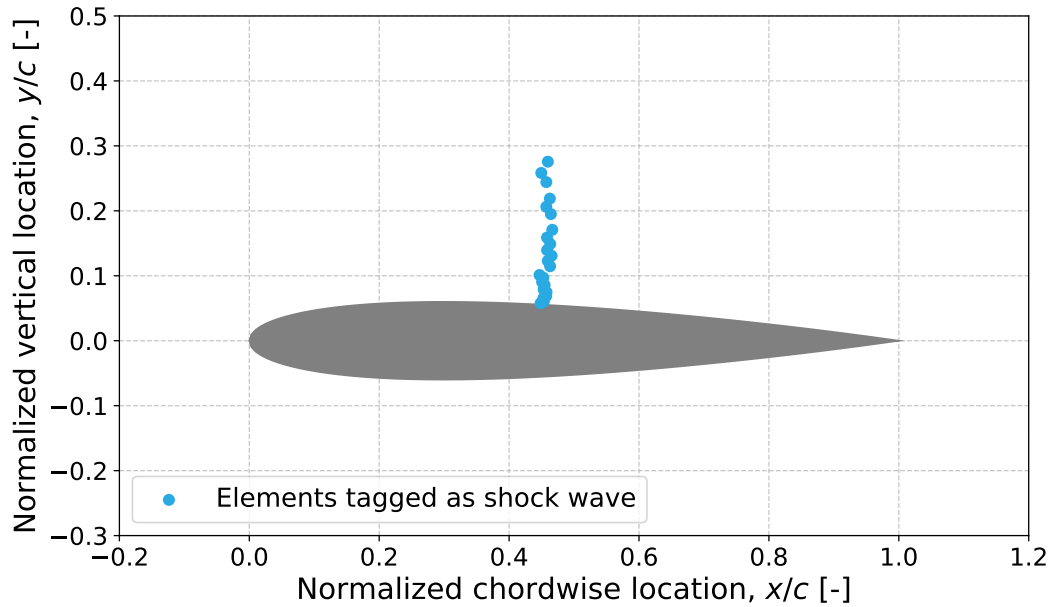


Figure 5.5: Overview of the shock wave locations for the NACA0012 airfoil at $\alpha = 1.49^\circ$ and $M = 0.752$.

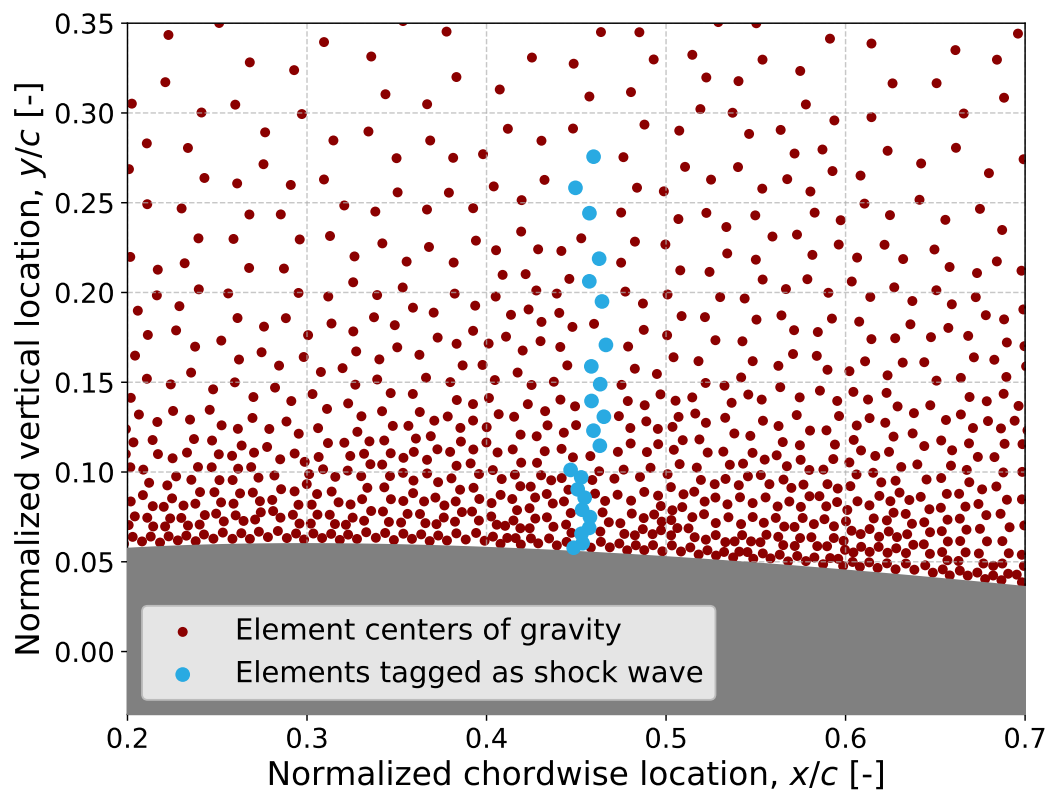


Figure 5.6: Zoomed in plot of the NACA0012 airfoil shock locations at $\alpha = 1.49^\circ$ and $M = 0.752$. All elements are shown, overlaid with the elements tagged as shock locations.

There indeed seem to be some omissions, one at the bottom, one in the middle, and one at the top of the shock. These omissions, while noticeable, are not particularly of great influence on the implementation of the entropy correction. This is because neighbouring elements are also selected when moving downstream of a shock.

For the RAE2822 case, the overview of the shock locations can be found in Figure 5.7. This shows a correct shock determination at the base of the shock, whilst at the top of the shock, an omission seems apparent. This is best analysed via the zoomed in plot, which is given by Figure 5.8.

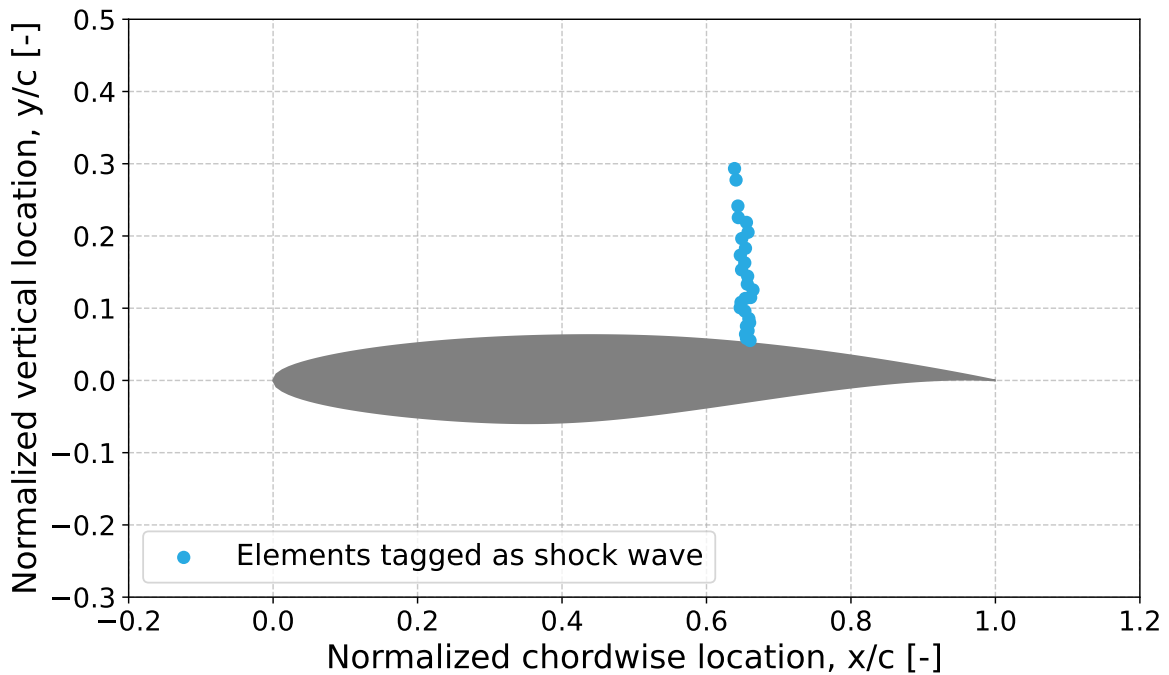


Figure 5.7: Overview of the shock wave locations for the RAE2822 airfoil at $\alpha = 2.4^\circ$ and $M = 0.725$.

It again shows that at the base of the shock, where the shock is stronger, the shock selection appears to be correct and in a straight manner. At the top, the small gap in the shock wave selection can be attributed to the absence of an element there.

Though the solver exhibits coarse selection at the top of the shock for both cases, this was found to have negligible effects on the overall solution quality. Mostly because the shock is weaker here anyway, hence the correction value is also smaller (close to 1). Furthermore, as mentioned before, the correction is also applied to neighbouring elements in the downstream domain, removing the strict necessity of selecting all shock wave elements. This is better showcased in section 5.6, where the domain of influence of the shock is analysed, together with the used correction values.




Figure 5.8: Zoomed in plot of the RAE2822 airfoil shock locations at $\alpha = 2.4^\circ$ and $M = 0.725$.

5.6 Shock Domain of Influence

Though the shock location determination showed some coarse selection at the top of the shock, its application towards the downstream flow field is considered to be more important. This section analyses the shock domain of influence regarding the entropy correction that is applied.

For the NACA0012 case, the shock domain of influence can be found in Figure 5.9. In this plot, the elements that are corrected for are shown as non-red, with values conforming to the colour-coding on the right. As can be seen, the corrections vary from values of 0.982-1. The colour-coding can also be seen as the different streamlines in the influenced domain of the shock wave.

Figure 5.9 immediately shows the importance of the far-field mesh size. It has a direct influence on the fidelity of the downstream region, and hence determines the resolution of this area.




Figure 5.9: Entropy correction field for the NACA0012 airfoil at $\alpha = 1.49$ and $M_\infty = 0.752$.

Another feature of *Flow*, which was introduced earlier, can be spotted at the wake. In the standard model of *Flow*, there is a wake sheet extended from the trailing edge. This is done to ensure the Kutta condition is correctly applied. In *EC-Flow*, this means that the entropy correction field cannot interfere with flow on the bottom side of the airfoil, which is desired.

For the RAE2822 case, the shock domain of influence can be found in Figure 5.9. This case was presented as a strong shock case, which can be seen by larger correction values and the slightly larger shock, which are now in the range of roughly 0.965-1, as compared to the NACA0012 airfoil.

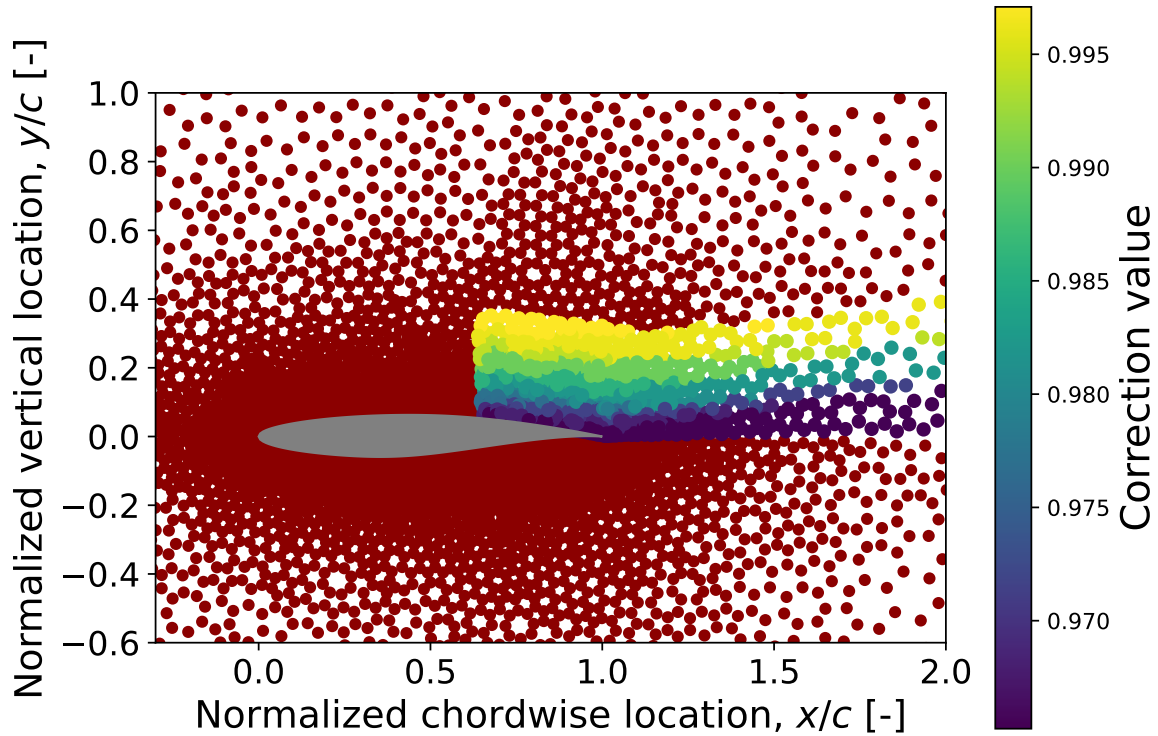


Figure 5.10: Entropy correction domain of influence for the RAE2822 airfoil at $\alpha = 2.4$ and $M_\infty = 0.725$.

5.7 Pressure Distributions

In this section, the results of the previous sections are combined towards the two different cases that have been analysed, determining the pressure coefficient plots and the load coefficients.

The pressure coefficient plot for the NACA0012 airfoil can be found in Figure 5.11, with its corresponding load coefficients in Table 5.7. Three solutions are compared: *Flow*, *EC-Flow*, and the Euler solution.

From Figure 5.11, it becomes clear that *EC-Flow* provides a significant correction to the standard *Flow* model. The shock location, as well as the shock strength, is better compared to the standard *Flow* model. Although the shock domain of influence remains a numerical approximation correction, the solved flow still seems to exhibit smooth behaviour after the shock.

Table 5.7 showcases the same behaviour as found in the pressure coefficient plot. The solution of *EC-Flow* moves towards the Euler solution, especially in the drag coefficient. Still, the lift coefficient and moment coefficient are somewhat off.

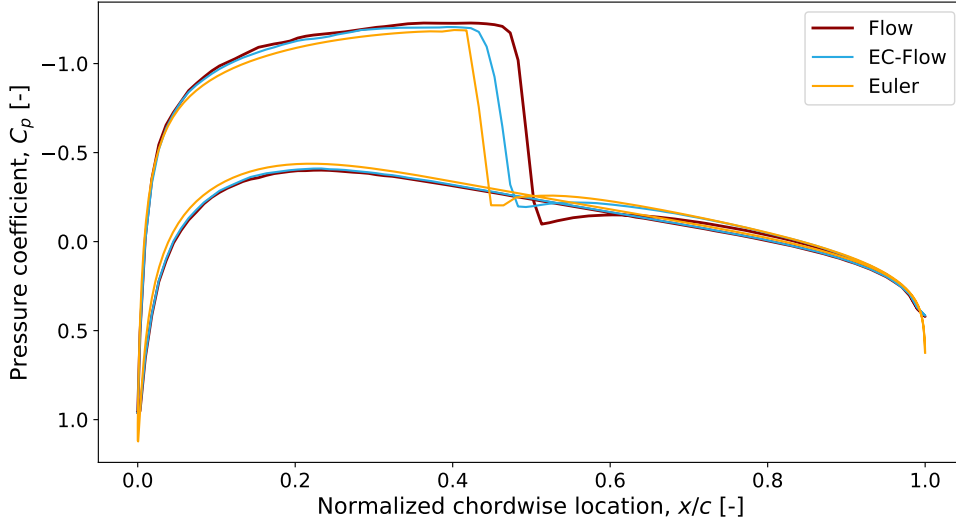


Figure 5.11: Pressure coefficient plot for the NACA0012 airfoil at $\alpha = 1.49$ and $M_\infty = 0.752$.

Table 5.7: Load coefficients for the NACA0012 airfoil at $M_\infty = 0.752$ and $\alpha = 1.49$.

NACA0012	<i>Flow</i>	<i>EC-Flow</i>	Euler
C_l	0.397	0.369	0.315
C_d	0.00782	0.00692	0.00654
$C_{m_{0.25c}}$	-0.00733	-0.00558	-0.000824

The RAE2822 case represents a stronger shock wave and, therefore, the results are somewhat more amplified compared to the NACA0012 case. The pressure coefficient plot can be found in Figure 5.12, and the load coefficients in Table 5.8.

Again, a significant change in shock location and strength can be seen, moving towards the Euler solution. In this extreme case, the solution is approaching the Euler solution.

Table 5.8: Load coefficients for the RAE2822 airfoil at $M_\infty = 0.725$ and $\alpha = 2.4$.

RAE2822	<i>Flow</i>	<i>EC-Flow</i>	Euler
C_l	1.1831	0.9591	0.7900
C_d	0.03035	0.01224	0.01179
$C_{m_{0.25c}}$	-0.1985	-0.1337	-0.1195

Table 5.8 follows the analysis of the previous plot, highlighting a significant move towards the Euler solution, especially in the drag coefficient. Here, the move in moment coefficient is somewhat more clear as well, and significant, as it can be better seen that it approximates the moment coefficient better compared to the standard model.




Figure 5.12: Pressure coefficient plot for the RAE2822 airfoil at $\alpha = 2.4$ and $M_\infty = 0.725$.

5.8 Polar Data

The polar data analysis offers a broader perspective on the aerodynamic performance of the studied models by showing the extent of the lift, drag, and moment coefficients, and their variations across a range of angles of attack. Unlike isolated case studies, polar data provides a more general assessment of the solver’s capabilities and limitations, revealing trends that might not be evident from analysing individual cases alone.

This broader approach is particularly valuable in understanding the solver’s performance boundaries. The results presented here compare *Flow*, *EC-Flow*, and the Euler solutions. The results for the Euler solution were generated using the open-source software SU2 ([Economon et al. \(2016\)](#)). By evaluating changes in aerodynamic loads and overall performance trends, this section demonstrates the implications of the entropy correction and contextualizes the solver’s strengths and constraints in predicting aerodynamic behaviour.

Before analysing the results, it should be noted that these plots were generated by analysing the polars until numerical divergence occurred in the solvers. Hence, beyond the results analysed in the subsequent section, the solvers diverge and no physical solution was obtained.

5.8.1 Lift Polar

The lift polar for the NACA0012 and RAE2822 airfoils can be found in Figure 5.13 and Figure 5.14, respectively. Here, the numerical nature of the solver can be seen, as it does not provide as smooth curves as *Flow* and the Euler solution.

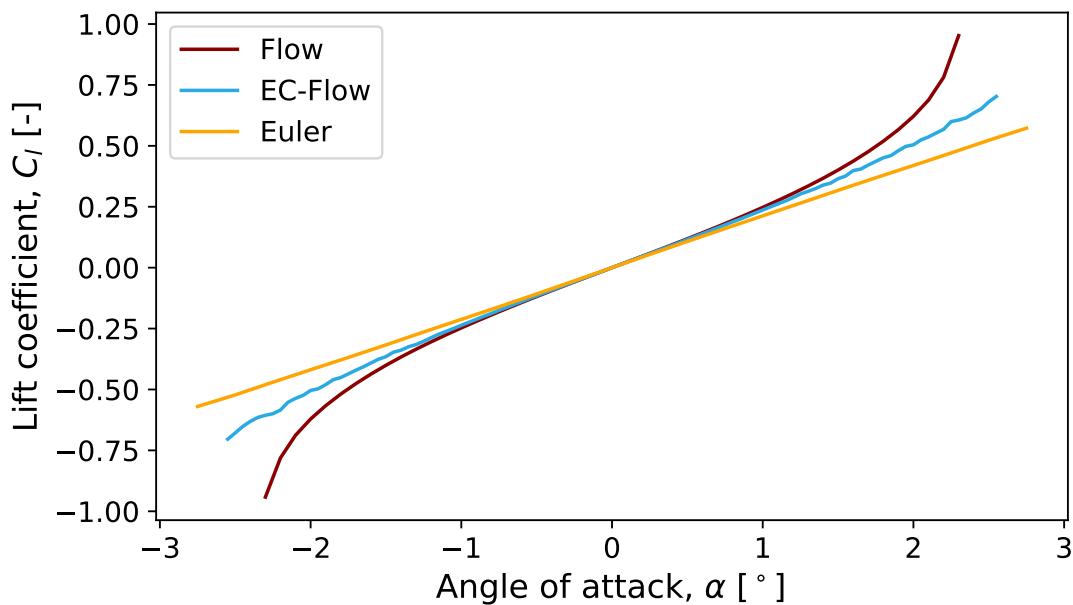


Figure 5.13: Lift polar plot for the NACA0012 airfoil at $M_\infty = 0.752$.

For the NACA0012 airfoil, the effects of the entropy correction are clear. The entropy corrected method more closely follows the Euler solutions, seemingly extending the linear behaviour of the lift polar plot somewhat. Furthermore, it can be seen that *EC-Flow*, though very slightly, extends the extent to which the solver is capable of solving (by around 0.25 degrees angle of attack).

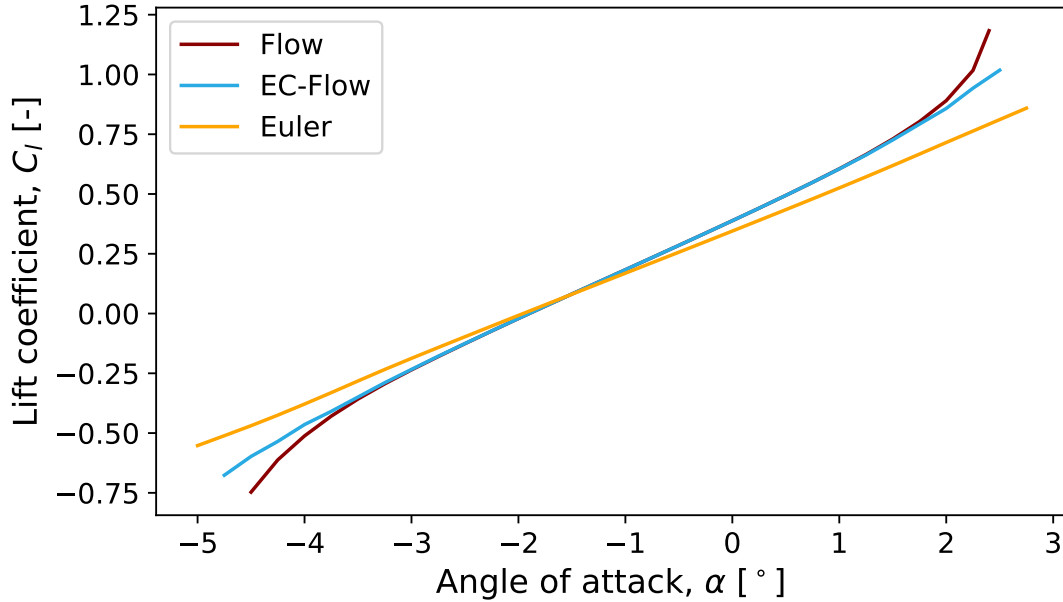


Figure 5.14: Lift polar plot for the RAE2822 airfoil at $M_\infty = 0.725$.

The same results apply to the RAE2822 airfoil. However, here, *EC-Flow* seems to more closely follow the curve of *Flow*. This can be explained by the slightly lower Mach number $M_\infty = 0.725$. This causes the strong shock effects to really start to take effect only at the extreme angles of attack; this is also the reason why the RAE2822 airfoil was analysed at the limit of 2.4 degrees angle of attack. Again, the solver, only slightly, extends the domain by around 0.25 degrees angle of attack.

It was first the intention to analyse the RAE2822 at an angle of attack of 2.55 degrees, as there is also experimental data available for this available. However, this proved to be outside the region of capability for the *Flow* solver.

5.8.2 Drag polar

The drag polar for the NACA0012 airfoil is illustrated in Figure 5.15.

This plot shows some interesting things. For small angles of attack, though hardly visible, the Euler solution predicts a slightly higher values for the drag coefficient. Once the shock starts taking effect on the solution at higher angles of attack, the drag coefficient predicted by the two full potential flow models is higher than the Euler solution. This is because the shock strength, predicted by the full potential solvers, is significantly stronger than the Euler solution. The drag coefficient therefore becomes higher as well.




Figure 5.15: Drag polar plot for the NACA0012 airfoil at $M_\infty = 0.752$.

Then, the drag polar for the RAE2822 airfoil can be found in Figure 5.16.




Figure 5.16: Drag polar plot for the RAE2822 airfoil at $M_\infty = 0.725$.

This plot shows the same behaviour as for the NACA0012 airfoil, though it being with ampli-

fied effects. At angles of attack ranging from -1 to 1, the Euler solution predicts a significantly higher drag coefficient for the airfoil compared to the two full potential solutions. At these angles of attack, the shock in the domain is not that strong, hence it does not influence the full potential flow solvers significantly. Again, once at extreme angles of attack, where the shocks are the predominant force, the two full potential solvers diverge from the Euler solution, and compute higher values than the Euler solution again.

5.8.3 Moment Polar

The moment polars for the NACA0012 and RAE2822 airfoils can be found in Figure 5.17 and Figure 5.18, respectively. For the NACA0012 case, it can be seen that the solver significantly alters the values of $Flow$ at the more extreme angles of attack. Here, the shock wave is most present, and the significant changes in especially the location of the shock seem to have a major influence on the moment coefficient prediction.




Figure 5.17: Moment polar plot for the NACA0012 airfoil at $M_\infty = 0.752$.

Both plots seem to experience the same behaviour, with the entropy correction pushing the solution of $Flow$ towards the Euler solution quality. However, at the more extreme angles of attack, both models still diverge from the Euler solution due to the absence of vorticity.




Figure 5.18: Moment polar plot for the RAE2822 airfoil at $M_\infty = 0.725$.

Although providing a better prediction, still only at extreme angles of attack the entropy change becomes visible, as that is where the shock waves are the dominant force. Furthermore, an extension of around 0.25 degrees in angle of attack is experienced, only a slight increase.

5.9 Computational Performance

To complement the preceding results, this section provides a final analysis of the computational time required by the three models. The results thus far have demonstrated that *EC-Flow* significantly improves the solution quality of *Flow*, achieving accuracy levels closer to those of the Euler solution.

However, computational efficiency remains a critical consideration in assessing the practicality of a solver. If the computational time required by *EC-Flow* approaches that of the Euler solver, the added complexity of *EC-Flow* may not justify its use over the more robust Euler solution. Therefore, while *EC-Flow* shows promise in balancing accuracy and efficiency, its computational performance relative to the other solvers requires careful evaluation.

The computational performance of the solvers for the NACA0012 case is summarized in Table 5.9. Among the solvers, *EC-Flow* actually requires approximately 22% less computational time than *Flow*, while the Euler solution demands more than twice the computational effort.

This decrease in computational time for *EC-Flow* was not readily expected before the start of this research, but can be attributed to the smoother flow prediction (less strong shock) by *EC-Flow*. As highlighted in section 5.4, the general convergence behaviour still exhibits oscillations in the residuals, primarily caused by updates to the entropy correction field. Addressing these oscillations could lead to notable performance improvements, making *EC-Flow* even more competitive computationally.

Table 5.9: Computational performance for the NACA0012 case.

Solver	Wall-Clock Time [s]	CPU Time [s]
Flow	3.74	3.74
EC-Flow	2.91	2.91
Euler	7.87	-

The computational performance for the RAE2822 case is summarized in Table 5.10. As discussed in section 5.4, *EC-Flow* requires fewer iterations to converge compared to *Flow*, demonstrating improved convergence behaviour. This is attributed to the weaker shock that is predicted using the *EC-Flow* model, allowing for smoother flow in the domain. The corresponding computational time for each solver is provided in Table 5.10, reflecting the gains in efficiency discussed previously in the number of iterations.

Table 5.10: Computational performance for the RAE2822 case.

Solver	Wall-Clock Time [s]	CPU Time [s]
Flow	9.0	9.0
EC-Flow	4.3	4.3
Euler	25	-

Notably, *EC-Flow* actually achieves convergence approximately 53% faster in computational time compared to *Flow*. This is a significant improvement again, and shows that the smoother flow predicted by *EC-Flow* allows for better computational performance.

Also for this case, although there is faster convergence, oscillations were observed in the convergence rates, indicating that there is still room for further optimization. Addressing these oscillations could enhance the computational performance of *EC-Flow* even further.

5.10 Sensitivity Analysis

Due to the addition of the entropy correction, some additional parameters have been added to the structure of *Flow*. In this section, the parameters that could influence the solution, or convergence, are analysed.

The entropy correction is currently applied on iterations 10-15, and afterwards for iterations

that are divisible by 6. These numbers have been determined by trial and error, and have been found to perform best. Therefore, it is interesting to see what the influence is of changing this parameter, which will be done in subsection 5.10.1.

Finally, the symmetry-preserving property of *Flow* is revisited. For symmetric airfoils at zero angle of attack, *Flow* maintains symmetric flow solutions. To confirm whether *EC-Flow* also retains this symmetry, a zero angle of attack case for the symmetric NACA0012 airfoil is evaluated in subsection 5.10.2.

5.10.1 Iteration Updates

As previously discussed, the entropy correction is currently applied at every sixth iteration. This frequency was determined to provide an optimal balance between ensuring sufficient updates to the entropy correction field while allowing the solver enough room to converge.

To evaluate the solver's sensitivity to this parameter, tests were conducted with varying entropy correction update intervals for both airfoil cases. Including both cases highlights similarities and differences across distinct flow conditions. Specifically, the solvers were tested with entropy updates applied every third, sixth, tenth, and fifteenth iteration. The results for the NACA0012 and RAE2822 cases are presented in Table 5.11 and Table 5.12, respectively.

Table 5.11: The sensitivity of the NACA0012 airfoil case to the iteration update of the entropy field.

Update at nth iteration	Cl	Cd	Number of iterations	Wall-Clock time [s]
3	-	-	-	-
6	0.369	0.00692	17	2.9
10	0.373	0.00697	19	3.2
15	0.367	0.00686	23	3.9

The sensitivity of the NACA0012 airfoil was analysed at its previous conditions of $M_\infty = 0.752$ and $\alpha = 1.49$. The results indicate that at an update frequency of 3 iterations, the solver fails to converge. This does not imply divergence but rather an inability to achieve sufficiently small residual values. This behaviour is attributed to overly frequent entropy correction updates, which introduce numerical oscillations that hinder convergence. The update frequency of 6 iterations delivers the best results, offering the optimal balance between computational performance and solution accuracy. It is therefore recommended for future research.

For the RAE2822 case, the update frequency of 6 iterations also provides the best computational performance. Similar to the NACA0012 case, a frequency of 3 iterations results in non-convergence due to excessive entropy correction updates causing numerical instability. At an update frequency of 15 iterations, the results deviate significantly, producing slightly less accurate outcomes. This discrepancy arises from the infrequent updates, which fail to sufficiently account for the evolving entropy field of the solution.

Table 5.12: The sensitivity of the RAE2822 airfoil case to the iteration update of the entropy field.

Update at nth iteration	Cl	Cd	Number of iterations	Wall-Clock time [s]
3	-	-	-	-
6	0.969	0.0128	24	4.3
10	0.967	0.0126	45	8.1
15	1.009	0.0150	28	4.8

5.10.2 Symmetry Analysis for NACA0012 Airfoil

To assess whether *EC-Flow* preserves symmetry in the solution, the NACA0012 airfoil is analyzed with a free-stream Mach number of $M_\infty = 0.8$, and an angle of attack of $\alpha = 0$. The pressure coefficient distributions for both *Flow* and *EC-Flow* are shown in Figure 5.19, and the corresponding aerodynamic load coefficients are provided in Table 5.13.

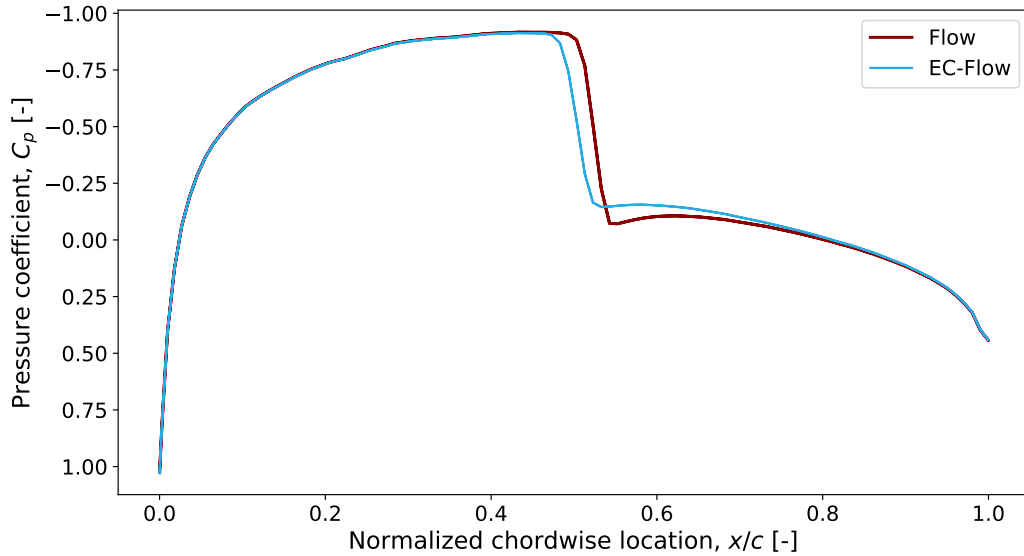


Figure 5.19: Pressure coefficient plot for the NACA0012 airfoil at $M_\infty = 0.8$ and $\alpha = 0$.

From Figure 5.19, it is evident that both solvers produce perfectly overlapping pressure coefficient distributions on the upper and lower surfaces of the airfoil, preserving the expected symmetry. A slight variation in shock location and strength is observed, attributed to the influence of the entropy correction.

The aerodynamic load coefficients further confirm this symmetric behaviour, with the lift

Table 5.13: Load coefficients for the NACA0012 airfoil at $M_\infty = 0.8$ and $\alpha = 0$.

NACA0012	<i>Flow</i>	<i>EC-Flow</i>
C_l	0.0000	0.0000
C_d	0.00735	0.00723
$C_{m_{0.25c}}$	0.0000	0.0000

and moment coefficients remaining zero, as expected for a symmetric airfoil at zero angle of attack. The slight alteration in shock location and strength is also reflected in the drag coefficient, with *EC-Flow* predicting a marginally lower value compared to *Flow*.

Chapter 6

Discussion

The introduction of an entropy correction in the full potential flow solver aimed to address the model's limitations in capturing the non-isentropic nature of shock waves in transonic flow. This chapter reflects on the results, connecting them to the research questions outlined in chapter 3, and analysing their implications for solution quality, computational efficiency, and applicability to aeroelastic contexts.

6.1 Effects of the Entropy Correction on Solution Quality

The incorporation of an entropy correction fundamentally transforms the solver's behaviour, yielding key improvements in the resolution of transonic shock waves. By shifting shock locations closer to the leading edge and reducing their strength, the entropy correction enhances the solver's ability to align predictions with physical reality. This improvement is crucial for aeroelastic applications, where the accuracy of shock wave resolution directly impacts aerodynamic load distribution and, consequently, structural deformations and responses.

Additionally, the solver demonstrates decreased sensitivity to mesh resolution, offering a robust basis for aeroelastic models that require reliable aerodynamic input for multiple structural configurations. To examine the full capabilities of *EC-Flow* within aeroelastic analysis, the results of the newly formed solver should be used in aeroelastic analyses in further research.

While the entropy correction extends the solver's range of applicability by improving load predictions, extreme cases with strong shocks remain a limitation. Nevertheless, the enhanced physical realism of the results provides a strong foundation for further coupling with aeroelastic solvers.

6.2 Comparison with State-of-the-Art Models

In transonic flow analysis, the entropy-corrected solver serves as a middle ground between the standard full potential flow solver and higher-fidelity models like the Euler equations. Its ability to more accurately predict shock location and strength directly benefits aeroelastic analyses, where shock-induced flow changes significantly affect structural responses. By bridging the gap between computational efficiency and solution fidelity, *EC-Flow* provides a practical tool for early aeroelastic design phases.

For strong shocks, *EC-Flow* faces limitations. The absence of vorticity modelling restricts its accuracy in extreme cases, where vorticity effects play a significant role. This is however the case for both the standard solver and the entropy-corrected solver. Consequently, the results in such scenarios should be interpreted with caution, particularly when applied to scenarios involving strong shocks or complex geometries.

Despite these limitations, *EC-Flow* retains a distinct computational advantage. It achieves faster convergence and lower computational costs compared to both the Euler solver and the standard *Flow* solver. These advantages make it particularly suitable for early design phases, where rapid iteration and resource management are important.

The extension of *EC-Flow* to reliable three-dimensional solutions remains an area for future research. The challenges associated with 3D flow, most apparent in span wise shock variations, present opportunities to further enhance the model for entropy corrections. Addressing these complexities would unlock the full potential of *EC-Flow*, enabling it to tackle a broader range of aerodynamic problems.

6.3 Numerical Convergence and Stability

The numerical convergence behaviour of the entropy-corrected solver provides several valuable insights. Across both analysed cases, *EC-Flow* demonstrates faster convergence rates compared to the standard full potential solver. This improvement can be attributed to the smoother flow profiles resulting from the entropy correction, which mitigates abrupt changes in the flow field caused by strong shocks that are over-predicted. The reduced discontinuities improve stability, allowing the solver to achieve convergence more rapidly.

Despite this improved convergence, the implementation of the entropy correction introduces residual oscillations during the iterations where the correction field is updated. Nevertheless, between updates, the solver exhibits rapid reductions in residuals, maintaining an overall faster convergence trend compared to the standard solver.

While the solver shows robust convergence behaviour, the oscillations introduced during entropy correction updates suggest an area for further improvement. Strategies such as adaptive or conditional updates could help smooth residuals and further enhance performance.

Chapter 7

Conclusion

This thesis presented the development and evaluation of an entropy correction in a finite element full potential flow solver, aimed at improving shock wave resolution in transonic aerodynamic flows. By lifting the isentropic assumption inherent in the standard solver, the entropy correction enhances the solver's ability to capture non-isentropic effects and resolve shock waves with greater fidelity, while maintaining the computational efficiency of the original model.

The results demonstrate that the entropy correction significantly improves shock predictions, including shock location, strength, and downstream flow profiles. These capabilities were validated across the NACA0012 and RAE2822 test cases, where the corrected solver consistently produced more realistic aerodynamic load predictions. The enhanced accuracy and smoother flow fields provide a robust foundation for aeroelastic analyses, where accurate load distributions are critical for evaluating structural responses and stability.

From a computational perspective, the entropy-corrected solver serves as an intermediate-fidelity model, trying to bridge the gap between the standard full potential flow solver and higher-fidelity Euler-based methods. Although the entropy correction introduces additional complexity, the solver remains computationally efficient enough for early design phases. Its enhanced stability for high-fidelity meshes and stronger shocks underscores its robustness in transonic applications, and even provides faster computational performance than the original model.

Despite these advancements, there is still room for improvement in the performance of *EC-Flow*. While the solver is well-suited to steady transonic flows, extending its capabilities to handle unsteady and three-dimensional flows would significantly broaden its applicability, particularly for aeroelastic analyses involving dynamic shock behaviour. Furthermore, the smoothing of residuals and thread-splitting of the entropy field updates can provide further improvements in computational performance.

In conclusion, the entropy correction represents a significant enhancement to the full potential flow solver, offering a practical and computationally efficient solution for modelling transonic flows. Its ability to predict and resolve shock waves with better accuracy makes it particularly well-suited for early aerodynamic and aeroelastic design phases. By addressing current limitations and extending its scope, the entropy-corrected solver has the potential to play a role in advancing computational aeroelasticity in early-phase design.

Chapter 8

Recommendations

This chapter outlines potential improvements and future directions for the entropy-corrected full potential flow solver (*EC-Flow*) developed in this research. While the solver demonstrates promise in accurately resolving shock waves and enhancing solution quality for transonic aerodynamic flows, there remain opportunities to extend its capabilities and optimize its computational performance. The recommendations presented here focus on addressing the challenges encountered in the current work and exploring directions for further development, particularly in expanding the solver to three-dimensional flows, improving convergence rates, and implementing advanced numerical techniques such as multi-threading, and adaptive mesh refinement. These enhancements aim to increase the robustness, efficiency, and applicability of *EC-Flow*, increasing the potential for further applicability in aerodynamic, and aeroelastic analysis.

8.1 Current State of Three-Dimensional Flow Analysis

This work primarily focuses on two-dimensional cases, where the entropy correction method has demonstrated significant potential, both in terms of solution quality and computational efficiency, particularly for strong shocks. While the methodology was initially developed for 2D, efforts are underway to extend it to three-dimensional problems.

Several complications occur when moving from 2D to 3D, especially in applying the entropy correction field on the domain behind the shock wave. In 2D, it was stated that the entropy correction is applied on downstream elements, and all of its neighbouring elements. This method provides accurate results in two-dimensional cases, but it does not hold in three-dimensional bodies. This is due to the inclusion of span wise effects of shock waves, which has a detrimental effect on the solution if the same methods are applied. Therefore, another method is used, based on Equation 8.1.

$$d = \frac{|V \times r|}{|V|} \quad (8.1)$$

In Equation 8.1, d is the perpendicular distance from the neighbouring element to the velocity vector of the analysed element. Moreover, V is the velocity vector, and r is the position vector between the two elements. Based on these calculated distances, from the n neighbouring elements, the m closest elements are selected for an entropy correction. A visual representation of this method is provided in Figure 8.1. While illustrated in 2D for clarity, the concept applies directly to 3D cases.

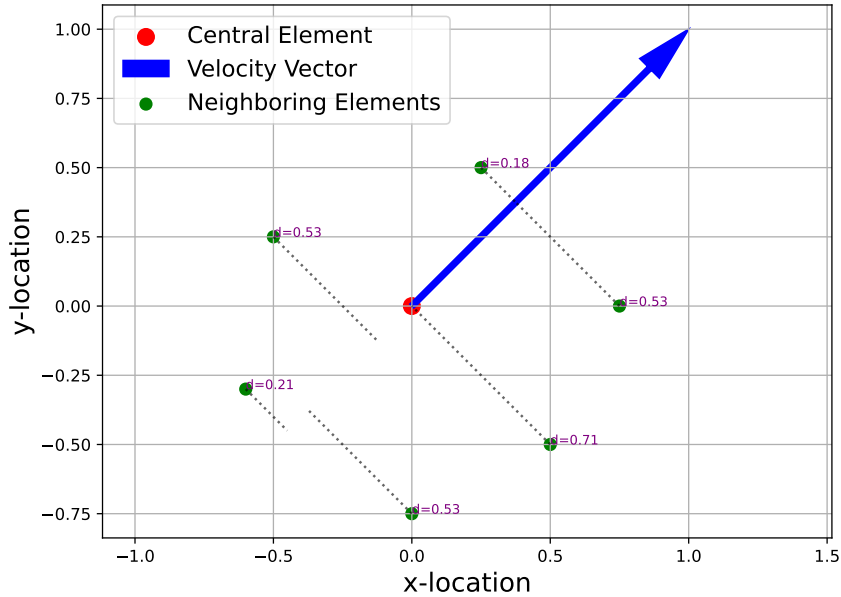


Figure 8.1: Distance based approach in 2D.

This method aims to mitigate the span-wise effects of shock waves, and already produces converged results for some cases. Moreover, the current implementation of the three-dimensional entropy correction excludes the tip of the aircraft wing, as this area provides unstable effects on the entropy correction. For example, for the AGARD445 wing, at a Mach number of $M_\infty = 0.9$, and an angle of attack of $\alpha = 2$, the produced load coefficients can be found in Table 8.1.

From Table 8.1, several things are apparent. First, the differences in three dimensions between the several solvers are way more subtle. The differences are small, and the entropy correction only affects the solution slightly. However, the lift and moment coefficients of *EC-Flow* do seem to better align with the Euler solution compared to *Flow*. In contrast, the drag coefficient

Table 8.1: Load coefficients for the AGARD445 wing at $M_\infty = 0.9$ and $\alpha = 2$.

AGARD445	Flow	EC-Flow	Euler
Cl	0.1380	0.1365	0.1365
Cd	0.00271	0.00261	0.00278
Cm	-0.1466	-0.1446	-0.1448

is actually larger in magnitude for the Euler compared to both full potential models. The entropy correction causes the drag coefficient to go down with respect to the standard model, which does seem not to be an improvement compared to the Euler solutions.

A serious limitation of *EC-Flow* in 3D is for cases with double shocks. For example, the ONERAM6 wing at a Mach number of $M_\infty = 0.839$ and angle of attack of $\alpha = 3.06$. In this scenario, flow particles traverse two shocks on the wing surface, necessitating the application of double corrections to the density field. While conceptually straightforward, the current implementation of *EC-Flow* struggles with differentiating between the two shocks due to the selection of spurious, false-positive shock locations in the region between them. This limitation prevents the accurate application of double corrections.

While the preliminary results in 3D are promising, they should be interpreted with caution. Further analysis and validation are necessary to ensure the robustness and accuracy of *EC-Flow* in three-dimensional applications.

8.2 Smoothing of Convergence Rates

As shown in section 5.4, the *EC-Flow* convergence rates show some oscillating behaviour. This can be attributed to the entropy correction field updates, and it causes negative effects on the residual convergence. Although the overall convergence of *EC-Flow* is faster than *Flow*, improvements are possible.

Currently, the entropy correction field is updated at every sixth iteration. As mentioned, each entropy field update causes an increase in relative residual values. Some general directions in which this can be analysed/improved will be given in the remainder of this section.

8.2.1 Threshold-Based Updates

The general idea of threshold-based updates can best be explained by Equation 8.2. Here, R_n is the relative (or absolute if wanted) residual value at iteration n , and ϵ is the threshold that is applied. It basically entails that the entropy correction update is only applied if the inequality in Equation 8.2 holds.

$$\frac{R_{n+1} - R_n}{R_n} < \epsilon \tag{8.2}$$

When the inequality does not hold, the solver stops updating the entropy correction field. This prevents unnecessary field updates when the solver is converging smoothly.

8.2.2 Conditional Field Updates

This is a rather simple idea, with possible beneficial effects on the residual convergence. It is nothing more than applying an update frequency dependent on the solving phase. Early phases might require more updates, as the solution changes more significantly than in later stages. E.g., in numerical terms, in early phases update the field every 2 iterations, while in later stages, only apply the updates at every 8 iterations.

The update frequency can also be made dependent on the actual residual values themselves.

8.2.3 General Remarks

Although the topic of this smoothing is rather subjective, there is potential for a lot of gain here. For example, if the convergence rates were smoothed for the RAE2822 case in this report, it would have the potential to converge as soon as iterations in the 20s. Especially when expanding this method to 3D the gain could provide significant.

8.3 Add Thread Splitting for Entropy Field Update

As discussed, the solver currently calculates the entropy correction field values for the entire fluid domain using a single thread. While *Flow* is designed to leverage multi-threading, the current implementation of the entropy correction does not utilize this capability. Implementing thread splitting for this process could reduce computational time, particularly for large and complex flow domains.

However, careful attention must be paid to avoid potential overlaps or conflicts in calculations. Ensuring that the domain is partitioned efficiently and that each thread operates independently within its assigned region is critical to maintaining numerical accuracy and stability. Strategies such as domain decomposition or task-based parallelism could be explored to achieve optimal performance. Proper synchronization mechanisms should also be implemented to manage dependencies between threads and ensure consistency in shared data.

By integrating multi-threading into the entropy field update process, the solver can enhance

its computational efficiency even further, making it even more suitable for iterative analyses in early design phases.

8.4 Adaptive Mesh Refinement

The current implementation of *Flow* primarily relies on the far-field mesh size to control the overall mesh density in the computational domain, including the shock domain of influence. However, this approach results in uniformly increased mesh density across the entire domain, even in regions where such refinement is unnecessary. This inefficiency increases computational costs, as the denser mesh is only critically required in the vicinity of the shock wave.

To address this, implementing an adaptive mesh refinement (AMR) strategy could offer significant advantages. An adaptive mesh would focus refinement efforts on the shock domain of influence, where higher resolution is essential to accurately capture the entropy correction. Simultaneously, coarser mesh elements could be used in regions where the flow is not affected by the entropy correction.

Bibliography

- J Anderson. *Fundamentals of Aerodynamics*. McGraw-Hill Education, New York, NY, 6 edition, 2017.
- O O Bendiksen. Review of unsteady transonic aerodynamics: Theory and applications. *Progress in Aerospace Sciences*, 47(2):135–167, 2011. ISSN 0376-0421. doi: <https://doi.org/10.1016/j.paerosci.2010.07.001>. URL <https://www.sciencedirect.com/science/article/pii/S0376042110000473>.
- E Crawley, H Curtiss, D Peters, R Scanlan, F Sisto, and E Dowell. *A Modern Course in Aeroelasticity*. Springer, Dordrecht, 3rd edition, 1995.
- A Crovato. *Steady Transonic Aerodynamic and Aeroelastic Modeling for Preliminary Aircraft Design*. PhD thesis, University of Liege, 2020.
- T D Economon, F Palacios, S R Copeland, T W Lukaczyk, and J J Alonso. SU2: An open-source suite for multiphysics simulation and design. *AIAA Journal*, 54(3):828–846, 2016. doi: 10.2514/1.J053813. URL <https://doi.org/10.2514/1.J053813>.
- M Hafez and D Lovell. Entropy and vorticity corrections for transonic flows. *International Journal for Numerical Methods in Fluids*, 8(1):31–53, 1 1988. ISSN 0271-2091. doi: <https://doi.org/10.1002/flid.1650080104>. URL <https://doi.org/10.1002/flid.1650080104>.
- G H Klopfer and D Nixon. Nonisentropic Potential Formulation for Transonic Flows. *AIAA Journal*, 22(6):770–776, 6 1984. ISSN 0001-1452. doi: 10.2514/3.48516. URL <https://doi.org/10.2514/3.48516>.
- K Ma, J Van Rosendale, and W Vermeer. 3D shock wave visualization on unstructured grids. In *Proceedings of 1996 Symposium on Volume Visualization*, pages 87–94, 1996. doi: 10.1109/SVV.1996.558049.
- A Parrinello and P Mantegazza. Improvements and Extensions to a Full-Potential Formulation Based on Independent Fields. *AIAA Journal*, 50(3):571–580, 3 2012. ISSN 0001-1452. doi: 10.2514/1.J051270. URL <https://doi.org/10.2514/1.J051270>.
- W Whitlow. Application of a nonisentropic full potential method to AGARD standard airfoils. In *26th Aerospace Sciences Meeting*, Aerospace Sciences Meetings. American Institute of Aeronautics and Astronautics, 1 1988. doi: doi:10.2514/6.1988-710. URL <https://doi.org/10.2514/6.1988-710>.

- W Whitlow, M M Hafez, and S J Osher. An entropy correction method for unsteady full potential flows with strong shocks. *Journal of Fluids and Structures*, 1(4):401–414, 1987. ISSN 0889-9746. doi: [https://doi.org/10.1016/S0889-9746\(87\)90286-6](https://doi.org/10.1016/S0889-9746(87)90286-6). URL <https://www.sciencedirect.com/science/article/pii/S0889974687902866>.
- Z Wu, Y Xu, W Wang, and R Hu. Review of shock wave detection method in CFD post-processing. *Chinese Journal of Aeronautics*, 26(3):501–513, 2013. ISSN 1000-9361. doi: <https://doi.org/10.1016/j.cja.2013.05.001>. URL <https://www.sciencedirect.com/science/article/pii/S1000936113001143>.

Appendix A

NACA0012 Case with Experimental Data

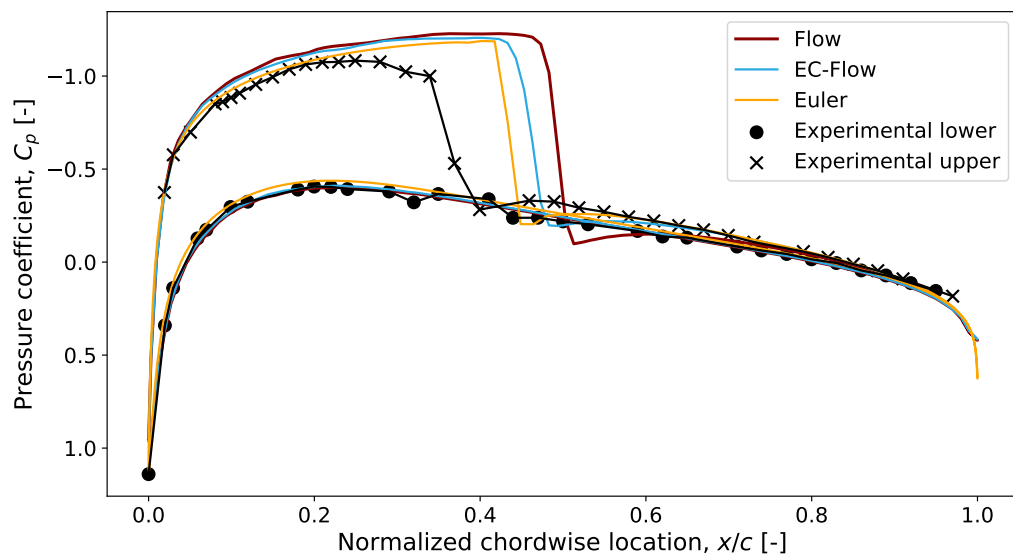


Figure A.1: NACA0012 airfoil at $M_\infty = 0.752$ and $\alpha = 1.49$ with experimental data.

