



UNIMORE

UNIVERSITÀ DEGLI STUDI DI
MODENA E REGGIO EMILIA

Department of Engineering "Enzo Ferrari"

Master's Degree in Vehicle Engineering

**Assessment of a novel Large Eddy Simulation
model in separating flows around a wing section**

Supervisor:

Andrea Cimorelli

Co-Supervisor:

Yves Tessier Urrecha

Candidate:

Marco Cesaroni

July 2026

*«Non si può descrivere la passione,
la si può solo vivere»*

Enzo Ferrari

Acknowledgements

This thesis represents the culmination of a long and demanding academic path. Since high school, every decision I have made has been guided by a profound passion for motorsport and Formula 1. This passion was kindled during my childhood, affectionately imparted by my grandfather. I dedicate my thoughts to him and my other grandparents who are watching over me from above: I know they would be incredibly proud and overjoyed to celebrate this achievement with me.

I wish to express my sincere gratitude to all those who have stood by me during the realisation of this work. Firstly, my deepest appreciation goes to my supervisor, Prof. Andrea Cimorelli, for his invaluable guidance, encouragement, and insightful feedback. I extend my heartfelt thanks to my co-supervisor, Yves Tessier Urrecha, for his constant availability and patience, as well as to Filippo Moroni for his constructive discussions and essential assistance. Their combined expertise and dedication have been instrumental in shaping this thesis.

Beyond my supervisors, I would also like to acknowledge my colleagues and friends at Lab. Zero. Sharing ideas, daily challenges, and moments of lightheartedness with you has profoundly enriched this experience.

On a more personal note, my endless gratitude belongs to my family, specifically my parents, my brother, my uncles, and my cousins. Their unconditional love, patience, and understanding have been my greatest source of strength throughout my entire university education.

Finally, a special thought goes to all my friends. To my lifelong friends in Perugia: thank you for always being my anchor and for standing by me over the years. To the friends I met here in Modena: thank you for the moments of leisure and the unforgettable experiences we shared together. You have all made this chapter of my life incredibly enjoyable, and I feel truly lucky to have you by my side.

Abstract

Simulating turbulent flows in engineering applications is challenging due to the high computational cost of resolving a wide range of time and length scales. While Direct Numerical Simulation (DNS) offers ideal accuracy, its computational demands make it impractical for industrial use. Consequently, Reynolds-Averaged Navier-Stokes (RANS) models remain the industry standard, offering rapid solutions by predicting only the mean flow. However, Large Eddy Simulation (LES) is becoming increasingly viable. By resolving large-scale fluctuations and modelling the rest, LES provides superior accuracy and is rapidly gaining industrial relevance thanks to advancements in high-performance computing and subgrid-scale modelling.

In the present work, a novel LES approach, the non-linear Tensorial Viscosity Model (TVM), and its Shear-Improved variant (SITVM) are tested and benchmarked against the established eddy-viscosity model, the Dynamic Smagorinsky Model (DSM). Relying on a new decomposition of the subgrid stress tensor, the TVM incorporates velocity gradients and the inertia properties of the grid element, while the SITVM further subtracts the mean velocity gradients to correctly yield zero subgrid stresses at the walls and in laminar regions. The models are evaluated on a NACA 4412 airfoil at a 15° Angle of Attack (AoA) and a chord Reynolds number of $Re_c = 66\,667$. Freestream turbulence is introduced to trigger the transition of the separated laminar boundary layer and promote flow reattachment. All simulations are performed using the Finite Volume Method (FVM) code T-Flows, and the results are compared against DNS and Implicit LES (ILES) data.

The findings demonstrate that, despite suffering from numerical stability issues under specific conditions, both tensorial models provide enhanced predictive capabilities compared to the traditional DSM and ILES, showing a closer agreement with high-fidelity DNS data in capturing laminar separation bubble and turbulent reattachment. However, zones characterised by permanent separation continue to pose significant challenges for all evaluated approaches, highlighting the inherent limitations of current subgrid-scale modelling and underscoring the need for further advancements in aerodynamic applications.

Contents

1	Introduction	1
2	Mathematical formulation of LES models	5
2.1	Governing equations	5
2.2	Smagorinsky model	6
2.3	Dynamic Smagorinsky model	7
2.4	Tensorial viscosity model	9
2.4.1	Theoretical framework	9
2.4.2	Modelling approach	11
2.5	Shear-Improved tensorial viscosity model	12
2.6	Implicit LES	13
3	Simulation setup	15
3.1	Computational grid	15
3.1.1	Domain definition	15
3.1.2	Mesh topology and resolution	15
3.1.3	Boundary conditions	19
3.2	Simulation parameters	21
3.3	Flow initialisation	22
3.4	Numerical schemes	22
3.5	Synthetic eddies	23
3.5.1	Turbulence intensity	26
3.5.2	Integral length scales	27
4	Simulation results	31
4.1	Instantaneous flow field	32
4.2	Mean flow topology	33
4.3	Wall stresses	37
4.4	Integral quantities and time scales	41
4.5	Turbulent kinetic energy and Reynolds stresses	44
4.6	Subgrid stresses	47

4.7	Subgrid dissipation rates	54
5	Conclusions	59
5.1	Accuracy	59
5.2	Numerical stability	62
5.3	Computational efficiency	63
5.4	Future improvements	64

List of Figures

3.1	Schematic representation of the computational domain and the NACA 4412 airfoil.	16
3.2	Overview of the computational mesh: entire domain (top), detailed view around the airfoil (middle), and close-up views of the LE and TE regions (bottom).	17
3.3	Synthetic eddy formation by convolution.	25
4.1	Q-criterion isosurfaces ($Q = 50$) superimposed with contours of the instantaneous streamwise velocity component u	33
4.2	Mean streamwise velocity component U	35
4.3	Mean spanwise vorticity component $\langle \omega_z \rangle$	36
4.4	Skin friction coefficient C_f distribution along the airfoil suction side. .	38
4.5	Pressure coefficient C_p distribution along the airfoil chord.	40
4.6	Lift coefficient spectra.	43
4.7	Mean resolved Turbulent Kinetic Energy (TKE) $\langle k \rangle$	45
4.8	Resolved Reynolds shear stress component $\langle u'v' \rangle$	46
4.9	Subgrid stress component $\langle \tau_{11} \rangle$	48
4.10	Subgrid stress component $\langle \tau_{22} \rangle$	49
4.11	Subgrid stress component $\langle \tau_{12} \rangle$	50
4.12	Total subgrid dissipation rate ϵ_{sgs}	56
4.13	Subgrid dissipation rate of mean kinetic energy E_{sgs}	57
4.14	Subgrid dissipation rate of turbulent kinetic energy ϵ'_{sgs}	57

List of Tables

3.1	Comparison of the a posteriori mesh resolution metrics: average and maximum non-dimensional grid spacings evaluated at the suction side of the airfoil.	18
3.2	Comparison of the average TI values obtained from the LESs with the reference value extracted from the DNS data.	27
3.3	Comparison of the average spanwise ILS values obtained from the LESs with the reference values extracted from the DNS data.	29
4.1	Aerodynamic coefficient statistics.	41

List of Acronyms

AoA	Angle of Attack
APG	Adverse Pressure Gradient
AR	Aspect Ratio
BARC	Benchmark on the Aerodynamics of a Rectangular Cylinder
BL	Boundary Layer
CFD	Computational Fluid Dynamics
CFL	Courant-Friedrichs-Lewy
DNS	Direct Numerical Simulation
DoF	Degrees of Freedom
DSM	Dynamic Smagorinsky Model
FFT	Fast Fourier Transform
FTT	Flow-Through Time
FVM	Finite Volume Method
ILES	Implicit Large Eddy Simulation
ILS	Integral Length Scale
LE	Leading Edge
LES	Large Eddy Simulation
NACA	National Advisory Committee for Aeronautics
NS	Navier-Stokes
PISO	Pressure-Implicit with Splitting of Operators
PRR	Primary Recirculation Region
PSD	Power Spectral Density
RANS	Reynolds-Averaged Navier-Stokes
SB	Separation Bubble

SEM	Synthetic Eddy Method
SIMPLE	Semi-Implicit Method for Pressure Linked Equations
SITVM	Shear-Improved Tensorial Viscosity Model
SM	Smagorinsky Model
SMART	Sharp and Monotonic Algorithm for Realistic Transport
SRR	Secondary Recirculation Region
TE	Trailing Edge
TI	Turbulence Intensity
TKE	Turbulent Kinetic Energy
TVM	Tensorial Viscosity Model
ZPG	Zero Pressure Gradient

Chapter 1

Introduction

In fluid dynamics, turbulence is a complex phenomenon arising from the amplification of flow perturbations, resulting in chaotic and irregular fluid motion. Although turbulent flows are described by deterministic equations, namely the Navier-Stokes (NS) equations, they remain unpredictable due to their extreme sensitivity to inevitable perturbations in fluid properties, as well as in initial and boundary conditions. Furthermore, turbulent flows are characterised by a wide range of spatial and temporal scales. This feature makes numerical simulations computationally expensive, as the number of Degrees of Freedom (DoF) to be resolved increases significantly with the Reynolds number (Re), specifically with $Re^{9/4}$. Consequently, while Direct Numerical Simulation (DNS) provides the most accurate and reliable results, it remains unfeasible for typical engineering applications due to its prohibitive computational cost.

To overcome this issue, turbulence modelling is required to significantly reduce the number of Degrees of Freedom (DoF) in the computed solution, thereby enabling simulations at relevant Reynolds numbers at a manageable computational cost. This reduction is typically achieved by applying a mathematical operator to the flow field variables. In practice, this operation can be of two types:

- filtering out all the turbulent fluctuations, which forms the basis of the Reynolds-Averaged Navier-Stokes (RANS) approach. In this methodology, the effects of the entire fluctuating field are modelled, while only the mean flow field is explicitly resolved;
- filtering out only the smallest scales of motion, which is the foundation of the Large Eddy Simulation (LES) approach. Here, only the small scales are modelled, while the large scales, being the most energetic and coherent structures, are explicitly resolved.

In both approaches, decomposing the flow variables into a resolved and an unresolved component introduces the so-called *turbulence closure problem*: the governing

equations are no longer closed, as the number of unknowns exceeds the number of equations. Consequently, the effect of the unresolved scales on the resolved ones must be accounted for by expressing the unknown terms as a function of the resolved flow variables. This practice is formally known as *turbulence modelling*.

In the LES framework, the filtering operation artificially truncates the energy cascade at a specific cut-off scale, known as the *filter length*. Without a proper closure, the kinetic energy from the resolved large scales cannot naturally transfer to the unresolved small scales, leading to an unphysical accumulation of energy near the cut-off scale. Therefore, a subgrid-scale model is required to extract this excess energy from the resolved field, thereby mimicking the forward energy cascade and preventing numerical instabilities. However, it is worth noting that the energy transfer is not strictly unidirectional: a reverse cascade, where energy flows from the subgrid scales back to the resolved ones, can also locally occur. This phenomenon is known as *backscatter*.

In the present work, two novel LES approaches, the Tensorial Viscosity Model (TVM) and its Shear-Improved variant (SITVM), are evaluated and benchmarked against the established Dynamic Smagorinsky Model (DSM). Unlike traditional closures based on a scalar subgrid viscosity, such as the Smagorinsky Model (SM) and the DSM, these novel models employ a tensorial viscosity. By coupling the filtered velocity gradients with the inertia tensor of the computational grid element, the TVM and SITVM effectively capture the anisotropic nature of the subgrid stresses. Moreover, this tensorial formulation naturally accounts for the previously described backscatter phenomenon, which purely dissipative eddy-viscosity models fundamentally fail to reproduce. A comprehensive discussion of these models and their mathematical formulation is presented in Chapter 2.

To evaluate their predictive capabilities, the TVM and SITVM are tested on a flow around a wing section, specifically a NACA 4412 airfoil with an infinite span (thereby excluding finite-wing induced phenomena such as tip vortices). Compared to the Benchmark on the Aerodynamics of a Rectangular Cylinder (BARC)¹, where the TVM model was tested for the first time on an external aerodynamic case² by Y. Tessier Urrecha in his Master Thesis [28], a wing section represents a completely different scenario since it has a non-fixed separation point, whereas the BARC has a fixed separation point. In the present setup, an Angle of Attack (AoA) of 15° and a chord Reynolds number of $Re_c = 66\,667$ are selected, producing a highly challenging flow topology on the suction side. As the attached laminar boundary layer encounters the strong Adverse Pressure Gradient (APG), it separates, forming a detached shear layer. This free shear layer subsequently undergoes transition and reattaches to the airfoil surface, enclosing a laminar separation bubble. Progressing downstream,

¹The BARC is a rectangular cylinder with a chord-to-thickness ratio of 5.

²The TVM model authors first tested it in a channel flow.

the newly developed turbulent boundary layer is subjected to a persistent APG, ultimately resulting in a permanent turbulent separation over the mid-chord region of the suction side [22].

All simulations are performed using T-Flows [14, 15], an open-source collocated Finite Volume Method (FVM) code designed for arbitrary computational grids. The obtained results are then compared against Implicit Large Eddy Simulation (ILES) and DNS data. The ILES approach allows for the evaluation of the artificial dissipation introduced by discretisation errors, thereby highlighting the additional fluctuation damping provided by the explicit LES models. Furthermore, while the DNS results cannot be considered a strict benchmark due to the different synthetic eddy generation approach adopted, they provide a highly accurate reference framework for this flow configuration. The DNS dataset was provided by R. Corsini³, and was generated using the high-order spectral element code *Nek5000*.

The remainder of this thesis is organised as follows: Chapter 2 presents the mathematical formulation of the investigated LES models; Chapter 3 describes the computational grid and the numerical methods, alongside the adopted synthetic turbulence generation technique; and Chapter 4 analyses the results obtained from the numerical simulations. Finally, Chapter 5 summarises the main findings, examines the numerical stability issues encountered with the tensorial models, evaluates their computational efficiency, and outlines potential directions for future research.

³Former researcher at Università degli Studi di Modena e Reggio Emilia.

Chapter 2

Mathematical formulation of LES models

2.1 Governing equations

The Navier-Stokes equations for an incompressible flow, neglecting thermal effects, read as:

$$\begin{cases} \frac{\partial u_i}{\partial x_i} = 0 \\ \frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} \end{cases} \quad (2.1)$$

where u_i denotes the instantaneous velocity components, p stands for the pressure field, x_i refers to the spatial position, while ρ and ν represent the fluid density and kinematic viscosity, respectively. The subscripts i and j indicate the Cartesian directions, with the standard Einstein summation convention applied over repeated indices. Naturally, the system (2.1) must be coupled with proper initial and boundary conditions.

In Large Eddy Simulation (LES), the instantaneous velocity field is decomposed as follows:

$$u_i(x_j, t) = \overline{u_i}(x_j, t) + u_i^{sgs}(x_j, t) \quad (2.2)$$

where $\overline{u_i}(x_j, t)$ is the resolved velocity field (representing the large scales) and $u_i^{sgs}(x_j, t)$ is the modelled subgrid-scale velocity field (representing the small scales). This decomposition relies on a spatial filtering operation, which is mathematically expressed as a convolution integral [19]:

$$\overline{u_i}(x_j, t) = \int G(r_j, x_j) u_i(x_j - r_j, t) dr_j \quad (2.3)$$

where $G(r_j, x_j)$ is the filtering function. Applying this filtering operation to the NS equations (2.1), and assuming a spatially homogeneous filter (i.e., $G = G(r_j)$), the following system of governing equations for the resolved flow field is obtained:

$$\begin{cases} \frac{\partial \bar{u}_i}{\partial x_i} = 0 \\ \frac{\partial \bar{u}_i}{\partial t} + \frac{\partial \bar{u}_i \bar{u}_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j \partial x_j} - \frac{\partial \tau_{ij}}{\partial x_j} \end{cases} \quad (2.4)$$

Compared to the original system (2.1), a new term appears. This allows for the definition of the subgrid stress tensor, formulated as:

$$\tau_{ij} = \bar{u}_i \bar{u}_j - \overline{u_i u_j} \quad (2.5)$$

The subgrid stresses are three-dimensional, unsteady, and highly chaotic variables. Physically, they can be thought of as an additional stress exerted on the resolved scales by the unresolved ones due to the filtering procedure. Since the non-linear term $\overline{u_i u_j}$ is symmetric, it introduces six independent unknown components, thereby giving rise to the well-known *turbulence closure problem*. Consequently, the subgrid stress tensor must be modelled. In the following sections, the LES models most relevant to the present work are introduced.

2.2 Smagorinsky model

Due to its simplicity and numerical robustness, the most widely used approach to model the subgrid stresses in LES is the Smagorinsky Model (SM) [27, 26]. The primary role of the subgrid stresses is to drain the kinetic energy from the resolved scales. Since a physically analogous energy dissipation is governed by fluid viscosity, it is common practice to model the subgrid stress tensor using the Boussinesq hypothesis:

$$\tau_{ij} - \frac{1}{3} \tau_{kk} \delta_{ij} = -2\nu_{sgs} \overline{S_{ij}} \quad (2.6)$$

where ν_{sgs} is the subgrid eddy viscosity and $\overline{S_{ij}}$ represents the resolved strain rate tensor.¹

By substituting equation (2.6) into the filtered momentum equations (2.4), the resolved NS equations can be rewritten as:

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial \bar{u}_i \bar{u}_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \frac{\partial [2(\nu + \nu_{sgs}) \overline{S_{ij}}]}{\partial x_j} \quad (2.7)$$

¹This is conceptually identical to the Boussinesq hypothesis employed in RANS for the Reynolds stress tensor, $\langle u'_i u'_j \rangle - \frac{1}{3} \langle u'_k u'_k \rangle \delta_{ij} = -2\nu_t \langle S_{ij} \rangle$, where ν_t is the turbulent eddy viscosity and $\langle S_{ij} \rangle$ the mean strain rate tensor.

This form of the momentum equation introduces an effective viscous coefficient, given by the sum of the kinematic fluid viscosity (ν) and the subgrid viscosity (ν_{sgs}). The role of the subgrid viscosity is essentially to locally increase the total diffusion of the flow field, thereby providing the necessary damping for the small-scale fluctuations.²

To calculate the subgrid viscosity, the SM relies on an analogy with Prandtl's mixing-length hypothesis, expressing ν_{sgs} as:

$$\nu_{sgs} = l^2 |\overline{S_{ij}}| \quad \text{with} \quad |\overline{S_{ij}}| = \sqrt{2\overline{S_{ij}} \overline{S_{ij}}} \quad (2.8)$$

where l is the characteristic length, defined as:

$$l = C_s \Delta \quad (2.9)$$

with C_s being the Smagorinsky constant (theoretically evaluated as $C_s \approx 0.17$ for homogeneous isotropic turbulence [21]) and Δ the filter length. Therefore, the modelled subgrid stresses take the form:

$$\tau_{ij} - \frac{1}{3}\tau_{kk}\delta_{ij} = -2(C_s\Delta)^2 |\overline{S_{ij}}| \overline{S_{ij}} \quad (2.10)$$

By definition, the subgrid dissipation rate is given by the contraction:

$$\epsilon_{sgs} = -\tau_{ij} \overline{S_{ij}} \quad (2.11)$$

Substituting equation (2.10) into this definition, it is possible to obtain:

$$\epsilon_{sgs} = \nu_{sgs} |\overline{S_{ij}}|^2 \geq 0 \quad (2.12)$$

Physically, the subgrid dissipation rate represents the energy transfer from the resolved large scales to the unresolved small ones. As demonstrated by equation (2.12), this quantity is strictly non-negative. Consequently, the SM is purely dissipative and can only reproduce a forward energy cascade. This inherent inability to account for backscatter, combined with the behaviour of the subgrid stresses in near-wall and laminar flow regions, represents the primary limitation of the model. To overcome these drawbacks, dynamic variants such as the Dynamic Smagorinsky Model (DSM) were developed, as introduced in the following section.

2.3 Dynamic Smagorinsky model

The primary objective of the Dynamic Smagorinsky Model (DSM) [11] is to dynamically adjust the Smagorinsky constant, and consequently the characteristic length

²Similarly, in RANS, the eddy viscosity artificially increases the total viscosity to damp the entire fluctuation spectrum, yielding the mean flow solution.

l , based on the local flow conditions. The foundation of the DSM relies on the introduction of a secondary spatial filter, known as the *test filter* (denoted by a tilde $\widetilde{\cdot}$), whose characteristic width $\widetilde{\Delta}$ is larger than the grid filter width Δ (typically $\widetilde{\Delta} \approx 2\Delta$). This double-filtering operation leads to a refined velocity decomposition:

$$u_i = \widetilde{u}_i + (\overline{u}_i - \widetilde{u}_i) + u_i^{sgs} \quad (2.13)$$

where $\widetilde{u}_i$ represents the velocity field of the largest resolved scales, $(\overline{u}_i - \widetilde{u}_i)$ represents the intermediate resolved scales (contained between the grid and test filters), and u_i^{sgs} is the unresolved subgrid velocity field. Applying the test filter to the original filtered NS equations (2.4) yields the subgrid stresses of the double filtered velocity field, defined as:

$$T_{ij} = \widetilde{\overline{u}_i \overline{u}_j} - \widetilde{u}_i \widetilde{u}_j \quad (2.14)$$

Concurrently, applying the test filter directly to the subgrid stress tensor defined in equation (2.5) gives:

$$\widetilde{\tau}_{ij} = \widetilde{\overline{u}_i \overline{u}_j} - \widetilde{u}_i \widetilde{u}_j \quad (2.15)$$

By subtracting equation (2.15) from equation (2.14), the term $\widetilde{\overline{u}_i \overline{u}_j}$ cancels out, yielding the well-known Germano identity:

$$\mathcal{L}_{ij} = T_{ij} - \widetilde{\tau}_{ij} = \widetilde{\overline{u}_i \overline{u}_j} - \widetilde{u}_i \widetilde{u}_j \quad (2.16)$$

This identity isolates the contribution to the Reynolds stresses from the scales whose length falls between the grid filter and the test filter, namely the small resolved scales [11]. Furthermore, the resolved turbulent stress tensor $\mathcal{L}_{ij}$ is a known computable quantity that provides a dynamic procedure to evaluate the Smagorinsky constant C_s . By applying the Boussinesq hypothesis at both filter levels, the subgrid stresses and the subgrid stresses of the double filtered velocity field are defined as:

$$\tau_{ij} = -2C_s \Delta^2 |\overline{S}_{ij}| \overline{S}_{ij} \quad (2.17)$$

$$T_{ij} = -2C_s \widetilde{\Delta}^2 |\widetilde{S}_{ij}| \widetilde{S}_{ij} \quad (2.18)$$

Noticeably, compared to the standard SM, the Smagorinsky constant C_s is not mathematically squared. This modification is theoretically designed to allow for negative values, thereby enabling the reproduction of backscatter phenomena (i.e., the reverse energy transfer from unresolved to resolved scales).

Substituting equations (2.17) and (2.18) into Germano identity yields an over-determined system of equations for the single unknown C_s . To evaluate this parameter, the least-squares minimization approach is universally adopted, computing the Smagorinsky constant dynamically as:

$$C_s = \frac{\langle \mathcal{L}_{ij} M_{ij} \rangle}{\langle M_{kl} M_{kl} \rangle} \quad \text{with} \quad M_{ij} = 2\Delta^2 |\widetilde{S}_{ij}| \widetilde{S}_{ij} - 2\widetilde{\Delta}^2 |\widetilde{S}_{ij}| \widetilde{S}_{ij} \quad (2.19)$$

where the brackets $\langle \cdot \rangle$ denote an averaging operation strictly required to avoid large localised fluctuations of C_s and ensure numerical stability.

By dynamically computing the constant, C_s correctly vanishes in laminar regions. Moreover, in near-wall regions, the model naturally recovers the proper asymptotic behaviour for the shear stress ($\tau_{12} \sim y^3$) without requiring empirical damping functions. However, despite its theoretical capability to handle negative viscosity, the DSM struggles to practically reproduce backscatter. From a numerical standpoint, a negative overall viscosity often leads to instabilities, forcing the implementation of a clipping procedure (i.e., strictly bounding $C_s \geq 0$) to preserve solver stability. Finally, being rooted in the Boussinesq hypothesis, the DSM remains a linear eddy-viscosity model. This imposes an alignment between the subgrid stress tensor and the resolved strain rate tensor, an assumption that fails in complex real flows.

2.4 Tensorial viscosity model

2.4.1 Theoretical framework

While the subgrid stress tensor is classically defined by equation (2.5), Germano [9] demonstrated that it can also be expressed as a function of the velocity increments between two points within the domain:

$$\tau(u_i, u_j) = \frac{1}{2} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [u_i(\boldsymbol{\xi}) - u_i(\boldsymbol{\eta})] [u_j(\boldsymbol{\xi}) - u_j(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.20)$$

By substituting the LES flow velocity field decomposition (equation 2.2) into this integral formulation, the subgrid stress tensor can be expanded into four distinct contributions:

$$\tau(u_i, u_j) = \tau(\overline{u_i}, \overline{u_j}) + \tau(\overline{u_i}, u_j^{sgs}) + \tau(u_i^{sgs}, \overline{u_j}) + \tau(u_i^{sgs}, u_j^{sgs}) \quad (2.21)$$

where, by using equation (2.20), the individual terms are computed as follows:

$$\tau(\overline{u_i}, \overline{u_j}) = \frac{1}{2} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [\overline{u_i}(\boldsymbol{\xi}) - \overline{u_i}(\boldsymbol{\eta})] [\overline{u_j}(\boldsymbol{\xi}) - \overline{u_j}(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.22)$$

$$\tau(\overline{u_i}, u_j^{sgs}) = \frac{1}{2} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [\overline{u_i}(\boldsymbol{\xi}) - \overline{u_i}(\boldsymbol{\eta})] [u_j^{sgs}(\boldsymbol{\xi}) - u_j^{sgs}(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.23)$$

$$\tau(u_i^{sgs}, u_j^{sgs}) = \frac{1}{2} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [u_i^{sgs}(\boldsymbol{\xi}) - u_i^{sgs}(\boldsymbol{\eta})] [u_j^{sgs}(\boldsymbol{\xi}) - u_j^{sgs}(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.24)$$

This specific decomposition provides the mathematical basis for various subgrid stress modelling approaches. For instance, the following approximation can be introduced:

$$\tau(u_i, u_j) \approx \tau(\overline{u_i}, \overline{u_j}) \quad (2.25)$$

which directly recovers the modelling approach given by similarity models, i.e., $\tau(u_i, u_j) \approx \overline{\overline{u_i}} \overline{\overline{u_j}} - \overline{\overline{u_i}} \overline{\overline{u_j}}$.

Alternatively, a different subgrid stress decomposition can be formulated, as proposed in [6]:

$$\tau(u_i, u_j) = \frac{1}{2} [\tau(\overline{u_i}, u_j) + \tau(u_i^{sgs}, u_j) + \tau(u_i, \overline{u_j}) + \tau(u_i, u_j^{sgs})] \quad (2.26)$$

where, applying equation (2.20), the terms are defined as:

$$\tau(\overline{u_i}, u_j) = \frac{1}{2} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [\overline{u_i}(\boldsymbol{\xi}) - \overline{u_i}(\boldsymbol{\eta})] [u_j(\boldsymbol{\xi}) - u_j(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.27)$$

$$\tau(u_i^{sgs}, u_j) = \frac{1}{2} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [u_i^{sgs}(\boldsymbol{\xi}) - u_i^{sgs}(\boldsymbol{\eta})] [u_j(\boldsymbol{\xi}) - u_j(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.28)$$

It is worth mentioning that both subgrid stress decompositions (equations 2.21 and 2.26) preserve Galilean invariance, owing to the Galilean invariance of their individual terms. Now, starting from the last decomposition (2.26), the following approximation can be performed:

$$\tau(u_i, u_j) \approx \frac{1}{2} [\tau(\overline{u_i}, u_j) + \tau(u_i, \overline{u_j})] \quad (2.29)$$

Furthermore, by considering that the large-scale motion $\overline{u_i}$ is sufficiently smooth at the filter scale, a first-order Taylor expansion of the resolved velocity increments can be performed:

$$\begin{cases} \overline{u_i}(\boldsymbol{\xi}) - \overline{u_i}(\boldsymbol{\eta}) \approx \frac{\partial \overline{u_i}}{\partial x_k} (\xi_k - \eta_k) \\ \overline{u_j}(\boldsymbol{\xi}) - \overline{u_j}(\boldsymbol{\eta}) \approx \frac{\partial \overline{u_j}}{\partial x_h} (\xi_h - \eta_h) \end{cases} \quad (2.30)$$

Substituting this expansion back into the integral definition leads to a novel formulation for the subgrid stresses:

$$\tau(u_i, u_j) \approx -\nu_{kj} \frac{\partial \overline{u_i}}{\partial x_k} - \nu_{ki} \frac{\partial \overline{u_j}}{\partial x_k} \quad (2.31)$$

where ν_{ij} represents a tensorial eddy viscosity, mathematically defined as:

$$\nu_{ki} = -\frac{1}{2} \tau(x_k, u_i) = -\frac{1}{4} \iint G(\mathbf{x}, \boldsymbol{\xi}) G(\mathbf{x}, \boldsymbol{\eta}) [\xi_k - \eta_k] [u_i(\boldsymbol{\xi}) - u_i(\boldsymbol{\eta})] d\boldsymbol{\xi} d\boldsymbol{\eta} \quad (2.32)$$

It is crucial to highlight that in the subgrid stress approximation (2.31), there is no longer a direct proportionality between the subgrid stress tensor and the filtered strain rate tensor. By breaking the linear alignment imposed by the Boussinesq hypothesis, this tensorial formulation fundamentally resolves one of the major limitations inherent to classical scalar eddy-viscosity models, such as the SM and DSM.

2.4.2 Modelling approach

Having established the theoretical framework, a practical modelling approach can be developed to derive the subgrid stress tensor. By applying a top-hat filter, defined by the kernel:

$$G(\mathbf{x}, \boldsymbol{\xi}) = \begin{cases} \frac{1}{\Omega} & \text{inside the grid element} \\ 0 & \text{otherwise} \end{cases}$$

and approximating the unfiltered velocity field via a first-order Taylor expansion around the filtered velocity field as:

$$u_i(\boldsymbol{\xi}) \approx \bar{u}_i(\mathbf{x}) + \frac{\partial \bar{u}_i(\mathbf{x})}{\partial x_h} (\xi_h - x_h) + O^2, \quad (2.33)$$

the subgrid viscosity tensor can be computed as:

$$\nu_{ki} = -\frac{1}{2\Omega} \frac{\partial \bar{u}_i(\mathbf{x})}{\partial x_h} \int_{\Omega} (\xi_k - x_k)(\xi_h - x_h) d\boldsymbol{\xi} \quad (2.34)$$

where Ω represents the volume of the computational grid element and $\mathbf{x}$ is its barycentric position (centroid).

Analysing equation (2.34), for $k \neq h$ the integral corresponds to the product of inertia of the grid element with respect to the barycentric axes in directions k and h :

$$\mathcal{I}_{kh} = - \int_{\Omega} (\xi_k - x_k)(\xi_h - x_h) d\boldsymbol{\xi}, \quad (2.35)$$

Conversely, for $k = h$, the integral yields:

$$\mathcal{I}_{kk} - \frac{1}{2} \text{Tr}(\mathcal{I}) = - \int_{\Omega} (\xi_k - x_k)^2 d\boldsymbol{\xi} \quad (2.36)$$

where $\mathcal{I}_{kk}$ is the moment of inertia of the grid element with respect to the k -th barycentric axis, and $\mathcal{I}$ is the barycentric inertia tensor.

Therefore, by defining a modified inertia tensor for the grid element as:

$$\mathcal{I}'_{kh} = \mathcal{I}_{kh} - \frac{1}{2} \delta_{kh} \text{Tr}(\mathcal{I}), \quad (2.37)$$

the subgrid viscosity tensor can be elegantly expressed as:

$$\nu_{ki} = \frac{1}{2\Omega} \frac{\partial \bar{u}_i}{\partial x_h} \mathcal{I}'_{kh} \quad (2.38)$$

This first-order approximation links the subgrid viscosity tensor directly to the filtered velocity gradients and the geometrical features of the grid element, specifically its

volume and inertia tensor. The fundamental physical premise behind this modelling approach is that the localised mesh structure directly influences, through the grid element inertia tensor, the anisotropic characteristics of the unresolved subgrid scales. Consequently, this geometric information can be effectively leveraged alongside the filtered velocity gradients to significantly improve the prediction of the subgrid viscosity tensor and, ultimately, the subgrid stress tensor.

In conclusion, by substituting equation (2.38) into the subgrid stress approximation (2.31), the final formulation for the subgrid stress tensor is obtained as:

$$\begin{aligned}\tau_{ij} &= - \left(\frac{1}{2\Omega} \frac{\partial \bar{u}_j}{\partial x_h} \mathcal{I}'_{kh} \right) \frac{\partial \bar{u}_i}{\partial x_k} - \left(\frac{1}{2\Omega} \frac{\partial \bar{u}_i}{\partial x_h} \mathcal{I}'_{kh} \right) \frac{\partial \bar{u}_j}{\partial x_k} = \\ &= - \frac{\mathcal{I}'_{kh}}{2\Omega} \left(\frac{\partial \bar{u}_j}{\partial x_h} \frac{\partial \bar{u}_i}{\partial x_k} + \frac{\partial \bar{u}_i}{\partial x_h} \frac{\partial \bar{u}_j}{\partial x_k} \right)\end{aligned}\quad (2.39)$$

As highlighted by the mathematical formulation, incorporating the geometric properties of the grid element enhances the representation of the unresolved subgrid scales, particularly when dealing with highly anisotropic and inhomogeneous mesh structures. This feature represents the primary advantage of the TVM over the classical SM and DSM, which are inherently constrained by a scalar subgrid viscosity definition. Furthermore, the TVM is mathematically non-linear, meaning that the subgrid stress tensor is no longer aligned with the resolved strain rate tensor, successfully overcoming a major physical limitation of traditional eddy-viscosity models. Finally, the TVM is theoretically capable of reproducing backscatter phenomena. However, due to its structurally under-dissipative nature, the model occasionally fails to drain a sufficient amount of kinetic energy from the resolved large scales, potentially triggering severe numerical instabilities during the simulation.

2.5 Shear-Improved tensorial viscosity model

To address the limitations of the standard TVM, particularly its inaccurate near-wall scaling, a further refinement can be introduced by applying the Shear-Improved approach, originally developed and tested for the Smagorinsky model [20]. The fundamental concept behind the Shear-Improved TVM (SITVM) is to compute the subgrid viscosity tensor by explicitly subtracting the mean filtered velocity gradients from the instantaneous ones:

$$\nu_{ki} = \frac{1}{2\Omega} \left(\frac{\partial \bar{u}_i}{\partial x_h} - \frac{\partial \langle \bar{u}_i \rangle}{\partial x_h} \right) \mathcal{I}'_{kh} \quad (2.40)$$

where the brackets $\langle \cdot \rangle$ denote an averaging operation to extract the mean flow field. Consequently, substituting this modified viscosity into equation (2.31), the subgrid

stress tensor for the SITVM is formulated as [7]:

$$\tau_{ij} = - \left[\frac{1}{2\Omega} \left(\frac{\partial \bar{u}_j}{\partial x_h} - \frac{\partial \langle \bar{u}_j \rangle}{\partial x_h} \right) \mathcal{I}'_{kh} \right] \frac{\partial \bar{u}_i}{\partial x_k} - \left[\frac{1}{2\Omega} \left(\frac{\partial \bar{u}_i}{\partial x_h} - \frac{\partial \langle \bar{u}_i \rangle}{\partial x_h} \right) \mathcal{I}'_{kh} \right] \frac{\partial \bar{u}_j}{\partial x_k} \quad (2.41)$$

This structural modification intrinsically accounts for the non-isotropic and inhomogeneous nature of the flow near solid boundaries. In the standard TVM, the presence of strong mean velocity gradients (such as the wall-normal gradient of the streamwise velocity) artificially generates subgrid viscosity even in the laminar sublayer, causing some components of τ_{ij} to improperly remain non-zero at the wall. By adopting the Shear-Improved correction, the mean shear contribution is perfectly removed: as the wall is approached, the instantaneous gradients tend to the mean gradients, causing the dominant mean-shear-induced subgrid viscosity to physically vanish. Furthermore, this property ensures that in purely laminar flow regions, where fluctuations are entirely absent, the subgrid stresses naturally reduce to zero. Finally, in regions characterised by fully developed, locally homogeneous, and isotropic turbulence (where the mean gradients are negligible), the SITVM automatically recovers the formulation of the standard TVM.

2.6 Implicit LES

The final approach employed in this work, serving to further evaluate the performance of the previously discussed models, is the Implicit Large Eddy Simulation (ILES). Unlike explicit formulations analysed so far, this method does not introduce any explicit subgrid-scale stress tensor into the governing equations. Instead, it relies entirely on the inherent numerical dissipation and dispersion introduced by the numerical schemes used to discretise the Navier-Stokes equations on a finite-resolution grid.

The fundamental premise of ILES is that the truncation error inherent to the numerical scheme, represented as an implicit artificial stress tensor σ_{ij}^{num} , can effectively mimic the energy drain provided by a physical subgrid stress tensor τ_{ij} . Consequently, the explicit modelling term $\partial \tau_{ij} / \partial x_j$ is deliberately omitted, operating under the assumption that the numerical scheme intrinsically provides the necessary damping:

$$\frac{\partial \tau_{ij}}{\partial x_j} \approx \frac{\partial \sigma_{ij}^{num}}{\partial x_j} \quad (2.42)$$

By directly exploiting the dissipative properties of the discretisation scheme, the filtered momentum equations reduce to the raw, unfiltered form of the NS equations, solved directly on the computational grid:

$$\frac{\partial \tilde{u}_i}{\partial t} + \frac{\partial \tilde{u}_i \tilde{u}_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \tilde{p}}{\partial x_i} + \nu \frac{\partial^2 \tilde{u}_i}{\partial x_j \partial x_j} \quad (2.43)$$

In this paradigm, the computational mesh and the chosen numerical scheme act concurrently as the implicit filter and the subgrid model.

Chapter 3

Simulation setup

3.1 Computational grid

3.1.1 Domain definition

The numerical simulations are performed within a C-type fluid domain, extruded homogeneously in the spanwise direction. The geometry under investigation is a NACA 4412 airfoil with a unit chord length ($c = 1$). The airfoil coordinates are defined by 200 discrete data points sourced from [1], distributed using a cosine spacing function to ensure proper resolution near the extremities, and featuring a blunt (open) Trailing Edge (TE).

The origin of the Cartesian coordinate system is located exactly at the airfoil Leading Edge (LE). The computational domain has an overall streamwise length $L_x = 20c$, a vertical height $L_y = 10c$, and a spanwise width $L_z = 0.6c$. To establish the required flow conditions, the airfoil geometry is rotated around the LE to achieve an Angle of Attack (AoA) of $\alpha = 15^\circ$ relative to the free-stream velocity vector, which is aligned with the x -axis.

As illustrated in Figure 3.1, the airfoil is centred vertically within the domain. The semi-circular inlet boundary is located at a radial distance of $5c$ from the LE, while the outlet boundary extends $15c$ downstream to allow for full wake development. Finally, the extruded airfoil span covers the entire spanwise dimension of the domain ($b = 0.6c$).

3.1.2 Mesh topology and resolution

The computational domain is discretised using the open-source mesh generator Gmsh [12]. To ensure the generation of purely hexahedral elements, the two-dimensional topology is partitioned into several sub-regions. This planar base is subsequently extruded along the spanwise z -direction into 40 uniform layers, resulting in a fully




Figure 3.1: Schematic representation of the computational domain and the NACA 4412 airfoil.

three-dimensional structured grid. The final mesh consists of approximately 6.4×10^6 hexahedral cells, representing $\approx 0.65\%$ of the baseline DNS Degrees of Freedom.

The overall grid quality was rigorously verified to guarantee optimal mesh metrics. This verification is of paramount importance for the present study, particularly given the employment of the Tensorial Viscosity Model (TVM) and its Shear-Improved variant (SITVM). As detailed in Sections 2.4 and 2.5, these specific LES models compute the tensorial subgrid viscosity by relying directly on the geometric properties and the inertia tensor of the individual grid cells. Consequently, excessive mesh non-orthogonality or severe cell skewness would inevitably lead to inaccurate predictions of the anisotropic subgrid stresses, thereby acting as a primary trigger for numerical instabilities. Therefore, the grid generation process was meticulously tuned to minimise skewness and non-orthogonality throughout the entire domain. Furthermore, dedicated attention was focused on restricting the cell Aspect Ratios (ARs) both at the wall and in the near-wake region (see Figure 3.2).

To effectively manage this geometric constraint while ensuring a high-fidelity




Figure 3.2: Overview of the computational mesh: entire domain (top), detailed view around the airfoil (middle), and close-up views of the LE and TE regions (bottom).

representation of the Boundary Layer (BL), the baseline spatial discretisation was guided by the recommended targets for a wall-resolved LES [25]: $\Delta\xi^+ \approx 50 - 150$, $\Delta\eta^+ < 2$, and $\Delta z^+ \approx 15 - 40$. Here, $\Delta\xi^+$, $\Delta\eta^+$, and Δz^+ represent the non-dimensional cell spacings at the wall in the tangential, wall-normal, and spanwise directions, respectively. It is worth noting that $\Delta\eta^+$ refers to the orthogonal distance from the wall to the first cell centre (i.e., half of the near-wall cell height). These three quantities are made non-dimensional through the friction length $l_\nu = \nu/u_\tau$, where $u_\tau = \sqrt{\tau_w/\rho}$ is the friction velocity. To evaluate u_τ a priori during the grid generation process, the wall shear stress τ_w was estimated by leveraging the local skin friction coefficient (C_f) distribution extracted from the DNS data.

The actual a posteriori mesh resolution metrics, evaluated specifically on the suction side of the airfoil for each simulation, are summarised in Table 3.1.

	$\langle\Delta\xi^+\rangle$	$\langle\Delta\eta^+\rangle$	$\langle\Delta z^+\rangle$	$\Delta\xi_{max}^+$	$\Delta\eta_{max}^+$	Δz_{max}^+
DSM	7.43	0.91	39.7	21.4	2.71	284
TVM	9.11	1.10	51.7	23.9	3.02	317
SITVM	7.93	0.96	41.6	22.3	2.83	297
ILES	6.91	0.84	37.8	21.4	2.71	285
DNS	4.10	0.21	3.60	7.60	0.45	7.50

Table 3.1: Comparison of the a posteriori mesh resolution metrics: average and maximum non-dimensional grid spacings evaluated at the suction side of the airfoil.

As can be observed, on most of the airfoil chord, the average tangential spacing $\langle\Delta\xi^+\rangle$ is kept significantly below the standard recommended targets ($\Delta\xi^+ \approx 50 - 150$). Regarding the wall-normal resolution, while the maximum values $\Delta\eta_{max}^+$ exceed the recommended threshold ($\Delta\eta^+ < 2$), the average spacing $\langle\Delta\eta^+\rangle$ remains consistently around 1 across the suction side, guaranteeing an optimal overall resolution of the viscous sublayer. Conversely, the average spanwise spacing $\langle\Delta z^+\rangle$ lies at the upper boundary of the suggested target range ($\Delta z^+ \approx 15 - 40$), or slightly above it depending on the evaluated subgrid model. Furthermore, its maximum value Δz_{max}^+ exhibits a prominent peak. It is crucial to clarify that this local maximum for the spanwise spacing occurs exclusively at the LE. In this stagnation and subsequent strong flow acceleration region, the skin friction coefficient inherently spikes, causing the local friction velocity u_τ to increase dramatically. Consequently, the non-dimensional grid spacing Δz^+ inflates mathematically. This severe peak is unique to the spanwise direction because, unlike the tangential and wall-normal dimensions which are highly refined at the LE, the physical spanwise grid width Δz is uniform. Despite this localised mathematical inflation, the physical grid dimensions remain strictly controlled and perfectly adequate along most of the airfoil chord. Finally, the grid spacing metrics for the baseline DNS are included in Table 3.1 to

provide a clear reference framework, highlighting the substantially higher resolution across all spatial dimensions, particularly in the wall-normal and spanwise directions.

The pronounced grid refinement in the tangential direction is strictly required to restrict the cell ARs at the wall to values below 10, particularly on the suction side past mid-chord, as higher AR values lead to numerical instabilities when running the TVM and SITVM. From a theoretical and numerical perspective, the adoption of highly stretched cells at the wall severely degrades the accuracy of the discretisation schemes used to compute the resolved velocity gradients. Since the tensorial models heavily rely on these gradients, coupled with the grid element inertia tensor, to evaluate the subgrid viscosity, any numerical inaccuracies are directly propagated into the subgrid-scale modelling. Moreover, considering the inherently low subgrid dissipation typical of these non-linear models, they lack the capability to damp the unphysical high-frequency oscillations triggered by inaccurate gradients, ultimately resulting in a divergent simulation. By aggressively reducing $\Delta\xi^+$, the near-wall elements are maintained closer to an isotropic shape, thereby preserving the gradient computation accuracy and guaranteeing overall numerical stability.

Finally, to properly initialise the freestream turbulence, a dedicated plane was defined upstream of the LE at a streamwise distance of $x = -1.5c$ to accommodate the generation of synthetic eddies. In this specific region, which spans the entire domain width and extends vertically between $y = -0.5c$ and $y = 0.5c$, a strictly isotropic cell resolution of $\Delta x = \Delta y = \Delta z = 0.015c$ is imposed. This spatial discretisation was deliberately selected to be consistent with the grid resolution adopted by Zucchi [30], who recently characterised and validated a specific Synthetic Eddy Method (SEM) implementation within the T-Flows solver.

3.1.3 Boundary conditions

Inlet At the inlet boundary, a uniform, dimensionless freestream velocity U_∞ is prescribed. Mathematically, a Dirichlet boundary condition is applied to the resolved velocity field, implying that the only non-zero component is the streamwise one:

$$\bar{u} = U_\infty = 1, \quad \bar{v} = 0, \quad \bar{w} = 0 \quad (3.1)$$

Therefore, the flow enters the domain with a uniform velocity profile, devoid of any turbulent fluctuations. As mentioned in Subsection 3.1.2, rather than being injected at the domain inlet, the synthetic eddies are locally superimposed onto the freestream velocity field at a dedicated plane located $1.5c$ upstream of the airfoil LE, leveraging a specific SEM subroutine implemented in the T-Flows solver [30].

Outlet At the downstream outlet boundary, a convective boundary condition is applied. In the context of LES, the standard practice of imposing a homogeneous

Neumann boundary condition (i.e., vanishing streamwise gradients) is generally inadequate. This approach causes an artificial distortion of the exit velocity field and generates spurious reflections of outgoing waves that propagate upstream, altering the solution within the computational domain.

To effectively circumvent this issue, the convective boundary condition proposed by Bottaro [2] is employed. This formulation is specifically designed to smoothly advect the turbulent vortices out of the domain without non-physical reflections, thereby maintaining numerical stability and ensuring global mass conservation even in the presence of strong, unsteady wakes. For a generic resolved flow variable $\bar{\phi}$ (such as the filtered velocity components $\bar{u}, \bar{v}, \bar{w}$), this boundary condition is mathematically expressed as:

$$\frac{\partial \bar{\phi}}{\partial t} + U_c \frac{\partial \bar{\phi}}{\partial x} = 0 \quad (3.2)$$

where U_c is the constant bulk convective velocity, assumed to be approximately equal to the freestream velocity U_∞ .

Top & Bottom boundaries For both the top and bottom sides, a symmetry boundary condition is applied. Mathematically, this formulation enforces a zero normal velocity component ($\bar{v} = 0$), whilst the normal gradients of both the tangential velocity components ($\bar{u}$ and $\bar{w}$) and the pressure ($\bar{p}$) are set to zero:

$$\bar{v} = 0, \quad \frac{\partial \bar{u}}{\partial y} = 0, \quad \frac{\partial \bar{w}}{\partial y} = 0, \quad \frac{\partial \bar{p}}{\partial y} = 0 \quad (3.3)$$

From a physical standpoint, this formulation models a free-slip condition. It allows the fluid to flow freely parallel to the boundaries and prevents any mass flux across them. This assumption is physically accurate only if the top and bottom boundaries are located in the far-field, at a sufficient distance from the airfoil, so that the aerodynamic perturbations induced by the body have completely dissipated and the undisturbed freestream flow is fully restored.

Lateral boundaries On the lateral domain boundaries perpendicular to the spanwise z -axis, periodic boundary conditions are imposed. This formulation numerically couples the two opposite boundaries by imposing that the values of the flow variables exiting one face instantaneously re-enter the domain through the other. Mathematically, for a generic resolved flow variable $\bar{\phi}$ (representing the filtered velocity components $\bar{u}, \bar{v}, \bar{w}$, and the pressure $\bar{p}$), this condition is expressed as:

$$\bar{\phi}(x, y, z = 0) = \bar{\phi}(x, y, z = L_z) \quad (3.4)$$

From a physical and computational perspective, this approach allows three-dimensional turbulent structures to naturally develop and interact along the spanwise direction,

avoiding the need to extend the computational domain and thus optimising the total number of DoF of the simulation.

Airfoil At the airfoil surface, a no-slip boundary condition is applied. Mathematically, this implies that all the resolved velocity components must identically vanish at the wall, enforcing a homogeneous Dirichlet condition:

$$\bar{u} = 0, \quad \bar{v} = 0, \quad \bar{w} = 0 \quad (3.5)$$

For the pressure field, coupling this kinematic constraint with the wall-evaluated momentum equations implicitly yields a homogeneous Neumann boundary condition, dictating a zero wall-normal pressure gradient:

$$\frac{\partial \bar{p}}{\partial \eta} = 0 \quad (3.6)$$

This coupled formulation ensures strict consistency with the fluid dynamics at the wall and actively prevents any unphysical mass flux across the solid aerodynamic surface.

3.2 Simulation parameters

The simulations are performed assuming an incompressible flow regime. Throughout this work, all fluid properties and governing equations are treated in a fully non-dimensional form, implicitly adopting the airfoil chord c and the freestream velocity U_∞ as unit reference length and velocity scales, respectively. Specifically, a unit mass density is imposed ($\rho = 1$). Consequently, the dynamic viscosity (μ) becomes numerically identical to the kinematic viscosity (ν), which in turn is defined as the inverse of the Reynolds number ($\nu = 1/Re = 1/66\,667 \approx 1.5 \times 10^{-5}$). This rigorous non-dimensionalisation approach allows the CFD solver to directly compute and output normalised aerodynamic fields.

Regarding the temporal discretisation, a constant time step of $\Delta t = 0.001 T$ is adopted, where $T = c/U_\infty$ denotes the characteristic time. For comparison, the high-fidelity baseline DNS was advanced using a significantly finer time step of $\Delta t = 2.5 \times 10^{-5} T$. Concerning the LESs, this choice of temporal resolution ensures that the maximum Courant-Friedrichs-Lewy (CFL) number remains strictly below 2. The highest CFL values are observed exclusively at the LE, particularly on the suction side. This localised peak is a direct consequence of the underlying geometrical constraints: since a structured mesh topology is employed, the highly refined, isotropic cell resolution specifically imposed at the synthetic eddies generation plane (as detailed in Subsection 3.1.2) inevitably propagates downstream. This

topological propagation results in an extremely dense grid clustering around the leading-edge region. Nevertheless, in all other areas of the computational domain, the CFL number safely remains below 1.

3.3 Flow initialisation

Given the high AoA ($\alpha = 15^\circ$) and the complex, separated nature of the flow around the NACA 4412 airfoil, the simulations are initialised by computing a preliminary potential flow solution. The potential solver, configured with a convergence tolerance of 1.0×10^{-6} and a maximum of 900 iterations, computes a divergence-free velocity field that strictly satisfies the imposed kinematic boundary conditions. This initial state prevents severe numerical shocks and ensures a smooth transition to the fully turbulent regime.

3.4 Numerical schemes

Spatial discretisation schemes In LES, minimising numerical dissipation is mandatory. A highly dissipative scheme would artificially damp the resolved velocity fluctuations, overwhelming the physical contribution of the subgrid-scale models. For this reason, the convective terms in the momentum equations are discretised employing the Sharp and Monotonic Algorithm for Realistic Transport (SMART) scheme [8]. This high-order (third-order accurate on Cartesian grids), bounded scheme effectively blends different differencing strategies to maintain high accuracy and stability while preventing unphysical oscillations.

Additionally, the computation of the velocity and pressure gradients is performed by applying the Gauss theorem, since the Least Squares method led to instabilities. The gradient evaluation is handled iteratively with a maximum of 12 iterations and a strict tolerance set to 1.0×10^{-3} , ensuring mathematical consistency even in the presence of highly skewed cells.

Time integration scheme To properly capture the dynamics of turbulent structures, the temporal integration of the governing equations (2.4) is performed using the Parabolic scheme. It is a second-order, fully implicit backward scheme that utilises values at three consecutive time levels (ϕ^{n-1} , ϕ^n , and ϕ^{n+1}) [13]:

$$\frac{\partial \phi}{\partial t} \approx \frac{3\phi^{n+1} - 4\phi^n + \phi^{n-1}}{2\Delta t} \quad (3.7)$$

where the superscripts $n + 1$, n , and $n - 1$ denote the new (unknown), current, and previous time steps, respectively. This scheme ensures high temporal accuracy

and robustness, providing an optimal balance between numerical stability and the preservation of the resolved turbulent fluctuations, without introducing the severe artificial dissipation typical of first-order methods.

Pressure-momentum coupling The pressure-momentum coupling is handled through a hybrid PISO-SIMPLE algorithm (commonly referred to as the PIMPLE approach [13]), combining the advantages of both procedures to ensure robust pressure-velocity coupling and numerical stability at each temporal increment.

Specifically, the solution at each time step is computed iteratively via an outer semi-implicit SIMPLE loop, constrained between a minimum of 3 and a maximum of 25 iterations. Within each outer iteration, the algorithm sequentially resolves the momentum predictor step to obtain an intermediate velocity field. This is immediately followed by an inner PISO loop, executed over 3 consecutive pressure-correction steps, to correct the velocities and strictly enforce mass conservation. To stabilise the non-linear iterative process against the complex dynamics of the separated flow, under-relaxation factors of 0.6 and 0.4 are applied to the momentum and pressure fields, respectively. The outer loop repeats until the global residuals reliably fall below the target tolerance of 1.0×10^{-4} (with the minimum of 3 iterations being strictly enforced), ensuring a fully converged, divergence-free solution before advancing to the next temporal increment.

3.5 Synthetic eddies

The BL viscous separation strongly depends on the upstream level of freestream turbulence, particularly in low-Reynolds-number simulations such as those performed in the present work. In this specific regime, the transition to turbulence occurs within the separated shear layer over the laminar separation bubble, ultimately dictating the potential reattachment of the turbulent BL. Indeed, without an adequate level of Turbulence Intensity (TI), the flow around the airfoil would remain entirely separated. Furthermore, the length scale of the fluctuations plays a critical role in triggering this shear layer transition. Specifically, the shear layer is highly receptive to turbulent structures of a size comparable to its own thickness, whilst its receptivity significantly decreases for larger scales [3].

To ensure a physically realistic representation of the incoming turbulent flow, synthetic eddies are generated at a designated plane located upstream of the airfoil LE, specifically at a streamwise distance of $1.5c$ and with a vertical extent of $-0.5c < y < 0.5c$, as detailed in Subsection 3.1.2. The choice of this specific injection location, rather than the domain inlet, is driven by two main physical and numerical requirements. Firstly, introducing the velocity fluctuations closer to the airfoil exploits the locally refined computational grid. Conversely, the inlet region

is characterised by a coarser spatial resolution, which would introduce excessive numerical dissipation, leading to an unphysical damping of the smallest turbulent structures [3]. Secondly, synthetic turbulence generation methods intrinsically require a downstream development distance (often referred to as *fetch length*) to allow the injected structures to non-linearly interact with the mean flow and reach a statistically stationary condition. As demonstrated by Zucchi [30], this transition region is directly proportional to the eddy diameter. Therefore, given the small turbulent scales adopted for these simulations, a fetch length of $1.5c$ is proven sufficient to ensure that the turbulent field has correctly developed and stabilised by the time it impinges on the airfoil LE.

The synthetic eddies are generated using a specific formulation of the Synthetic Eddy Method (SEM) [16], namely the Vortex-Based SEM, developed and implemented in T-Flows by Zucchi [30]. This method relies on the superposition of coherent circular vortices onto the uniform freestream velocity field $(U_\infty, 0, 0)$ in a plane perpendicular to the streamwise direction. Considering a single vortex centred at the point (y_0, z_0) , the velocity components induced by this eddy are evaluated as (see Figure 3.3):

$$\begin{cases} v'(\mathbf{x}) = s \frac{z - z_0}{r} f(\mathbf{x}; \mathbf{x}_0, \sigma) I \\ w'(\mathbf{x}) = s \frac{y - y_0}{r} f(\mathbf{x}; \mathbf{x}_0, \sigma) I \end{cases} \quad (3.8)$$

where $s = \pm 1$ denotes the random vortex rotation direction, r is the vortex radius, f is the shape function governing the spatial distribution of the vortices, and I is an intensity parameter. Specifically, the shape function is defined by an anisotropic Gaussian distribution:

$$f(\mathbf{x}; \mathbf{x}_0, \sigma) = \exp \left[-\frac{(x - x_0)^2}{2\sigma_x^2} - \frac{(y - y_0)^2}{2\sigma_{yz}^2} - \frac{(z - z_0)^2}{2\sigma_{yz}^2} \right] \quad (3.9)$$

where $\sigma_x = 3r$ controls the streamwise extent of the vortex, while $\sigma_{yz} = r/2$ dictates its vertical and spanwise extents.

To guarantee the periodic boundary conditions along the spanwise direction, the vortices are generated in spatially correlated pairs: one vortex is located at (y_0, z_0) and its periodic counterpart at $(y_0, z_0 \pm L_z)$. This ensures that the velocity field induced by the eddies remains perfectly consistent with the topological periodicity of the computational domain.

Finally, the streamwise velocity fluctuation (u') induced by an eddy is automatically computed by the CFD solver to satisfy the kinematic incompressibility constraint ($\nabla \cdot \mathbf{u} = 0$), thus guaranteeing a divergence-free filtered velocity field. This numerical choice is motivated by the physics of real turbulence, where streamwise velocity fluctuations are intrinsically driven by vortex stretching and non-linear




Figure 3.3: Synthetic eddy formation by convolution.

interactions among the transverse velocity components. Therefore, allowing the solver to dynamically compute u' based on the induced v' and w' yields a far more physically consistent representation of the incoming turbulence.

Control parameters

The Vortex-Based SEM formulation allows governing the synthetic eddies generation by tuning the following three parameters:

- **Eddy diameter (D):** this parameter defines the physical size of the generated vortices, which is directly related to the Integral Length Scale (ILS) of the resulting turbulent flow.
- **Eddy intensity (I):** this parameter multiplies the velocity components induced by the vortices, thus controlling the Turbulent Kinetic Energy (TKE) and the overall TI of the generated flow.
- **Eddy generation frequency (f):** this parameter dictates the discrete temporal interval (expressed in time steps) between the generation of two consecutive sets of vortices, and is computed as:

$$f = \frac{k L}{\Delta t} = \frac{3 k D}{\Delta t} \quad (3.10)$$

where $k = 1$ for the "Low" frequency case and $k = 0.5$ for the "High" frequency case, whilst $L = 3 D$. By tuning this parameter, it is possible to control the temporal and spatial vortex density, directly altering the spectral content of the generated turbulence.

In the present work, the above parameters are calibrated to approximate the turbulence characteristics of the baseline DNS data, as detailed in Subsections 3.5.1 and 3.5.2. Initially, an in-depth analysis of the results obtained by Zucchi [30] on the Vortex-Based SEM was leveraged to identify an optimal starting combination. However, his characterisation was performed in an empty rectangular domain, periodic in both the vertical and spanwise directions. The introduction of the solid airfoil in the present computational domain severely alters the flow field, representing a fundamentally different aerodynamic topology. Therefore, a further, dedicated tuning of these parameters was required. The preliminary simulations, together with the theoretical constraints highlighted by Zucchi [30], demonstrated the physical impossibility of simultaneously matching both the TI and the ILSs of the baseline DNS data. Consequently, a targeted engineering compromise was made: priority was given to matching the turbulence intensity, as it is the most critical parameter governing the BL separation, whilst accepting lower integral length scale values compared to the DNS target. The final adopted parameters are: $D = 0.045$, $I = 2$, and $f = 135$ time steps, obtained by imposing a "Low" frequency setting ($k = 1$).

3.5.1 Turbulence intensity

The turbulence intensity of the generated synthetic eddies is a crucial parameter, as it directly affects the BL separation and the overall flow behaviour around the airfoil [3]. The TI I_∞ of the LESs is evaluated at a specific monitoring plane located upstream of the airfoil LE, specifically at a streamwise distance of $0.1c$. Furthermore, this monitoring plane spans the entire domain width and extends vertically between $y = -0.3c$ and $y = 0.3c$, thus covering the entire region where the incoming turbulence interacts with the airfoil BL.

To accurately compute the TI, the three-dimensional discrete flow fields extracted from the LESs are post-processed by performing a statistical analysis based on the Reynolds decomposition. To ensure methodological consistency, both this TI statistical evaluation and the ILSs analysis detailed in the subsequent Subsection 3.5.2 are conducted over the identical temporal sampling interval configured for the downstream aerodynamic results presented in Chapter 4. Exploiting the spanwise homogeneity of the flow, the statistical moments are evaluated by averaging over both time (t) and the spanwise direction (z). The space-time variances of the velocity fluctuations ($\langle u'^2 \rangle$, $\langle v'^2 \rangle$, $\langle w'^2 \rangle$) are therefore computed directly over the specified sampling window and the entire spanwise extent for each vertical coordinate y . These variances allow for the evaluation of the Turbulent Kinetic Energy (TKE), defined as:

$$k = \frac{1}{2} (\langle u'^2 \rangle + \langle v'^2 \rangle + \langle w'^2 \rangle) \quad (3.11)$$

Then, the TI along the vertical profile is calculated by normalising the square root

of the TKE with respect to the freestream velocity U_∞ :

$$I_\infty = \frac{\sqrt{\frac{2}{3}k}}{U_\infty} \quad (3.12)$$

Finally, the computed TI profile is spatially averaged across the vertical direction to obtain a single representative metric. The average TI $\langle I_\infty \rangle$ values obtained are reported in Table 3.2 for all the performed simulations, together with the reference value extracted from the DNS data in a monitoring plane located at the same streamwise distance from the airfoil LE but with a vertical extent of $-1.0c < y < 1.0c$.

	DSM	TVM	SITVM	ILES	DNS
$\langle I_\infty \rangle$ [%]	4.8	4.9	5.0	5.0	4.9

Table 3.2: Comparison of the average TI values obtained from the LESs with the reference value extracted from the DNS data.

As shown in Table 3.2, the simulation setup successfully replicates the reference TI with excellent accuracy. The marginal deviations observed among the LES approaches are physically consistent and stem from their respective subgrid-scale formulations, alongside standard statistical scatter. From a modelling perspective, numerical artificial dissipation acts as the primary source of turbulence decay during the downstream convection of the synthetic eddies. The DSM, being a purely dissipative linear eddy-viscosity closure, introduces an additional turbulent viscosity that further damps the fluctuations, resulting in the most pronounced TI reduction. Conversely, the non-linear tensorial models (TVM and SITVM) natively account for energy backscatter, effectively mitigating this excessive subgrid damping.

3.5.2 Integral length scales

The ILSs of the generated turbulence are another fundamental parameter, as they dictate the characteristic size of the energy-containing eddies within the flow. As thoroughly discussed in Section 3.5, simultaneously matching both the ILSs and the TI of the baseline DNS data proved physically unfeasible with the Vortex-Based SEM. Consequently, a targeted compromise was adopted: priority was given to the TI due to its dominant role in the BL separation dynamics, intrinsically accepting lower ILSs values compared to the DNS targets.

To evaluate the effective size of the synthetic eddies introduced into the computational domain, the spatial autocorrelation of the velocity fluctuations is analysed. This evaluation is performed on a monitoring plane located upstream of the LE at a

streamwise distance of $x = -0.1c$, extending across the entire spanwise width and vertically between $y = -0.3c$ and $y = 0.3c$.

Exploiting the spanwise homogeneity of this specific flow configuration, the normalised two-point spatial autocorrelation functions of the velocity fluctuations (ρ_{uu} , ρ_{vv} , and ρ_{ww}) are computed along the spanwise direction for each vertical coordinate y and time instant t :

$$\rho_{uu}(x, y, r) = \frac{\langle u'(x, y, z, t)u'(x, y, z + r, t) \rangle}{\langle u'^2(x, y, t) \rangle} \quad (3.13)$$

$$\rho_{vv}(x, y, r) = \frac{\langle v'(x, y, z, t)v'(x, y, z + r, t) \rangle}{\langle v'^2(x, y, t) \rangle} \quad (3.14)$$

$$\rho_{ww}(x, y, r) = \frac{\langle w'(x, y, z, t)w'(x, y, z + r, t) \rangle}{\langle w'^2(x, y, t) \rangle} \quad (3.15)$$

where r represents the spatial separation distance between two points along the z -axis, and the angle brackets $\langle \cdot \rangle$ denote spatial averaging over the entire spanwise extent of the domain at a specific time instant t .

From a theoretical standpoint, the ILS of turbulence is rigorously defined as the mathematical integral of the spatial autocorrelation function with respect to the separation distance r . However, in the methodology implemented for this study, to ensure coherence with the provided DNS data, the ILSs L_{uu} , L_{vv} , and L_{ww} are evaluated by identifying them directly with the separation distance r at which the autocorrelation curve first intersects the abscissa axis. This condition, known as the first zero-crossing, is a robust and widely accepted convention.

Once the instantaneous L_{uu} , L_{vv} , and L_{ww} values are computed, they are time-averaged across all the evaluated snapshots. Following the same procedure adopted for the TI, these resulting ILS vertical profiles are then spatially averaged along the y -coordinate, yielding a single global metric for the monitoring plane. To provide a rigorous baseline, the DNS target values were derived by averaging the data extracted from three specific probes within the spatial sampling window evaluated in the LES models ($x = -0.1c$, $y \in [-0.3c, 0.3c]$). A direct comparison between these reference targets and the spatial averages computed for the different LES approaches is presented in Table 3.3.

As can be observed from the tabulated data, a general under-prediction of all generated ILSs compared to the DNS baseline is evident. However, while the LES models provide a closer approximation for the streamwise and wall-normal fluctuations ($\langle L_{uu} \rangle$ and $\langle L_{vv} \rangle$), a significantly more pronounced discrepancy is observed for the spanwise ILS, $\langle L_{ww} \rangle$.

This specific deviation could primarily be driven by the spanwise spatial discretisation. As highlighted by the mesh metrics in Table 3.1, the baseline DNS data

	DSM	TVM	SITVM	ILES	DNS
$\langle L_{uu} \rangle$	0.064	0.063	0.064	0.061	0.078
$\langle L_{vv} \rangle$	0.059	0.057	0.058	0.059	0.080
$\langle L_{ww} \rangle$	0.094	0.092	0.092	0.093	0.190

Table 3.3: Comparison of the average spanwise ILS values obtained from the LESs with the reference values extracted from the DNS data.

was generated on a remarkably finer spanwise grid, whereas the LES methodologies inherently rely on a coarser resolution. This fundamental grid constraint physically limits the ability of the models to resolve and sustain the extended spatial coherence of the velocity fluctuations along the spanwise direction, causing the severe under-prediction of $\langle L_{ww} \rangle$.

Chapter 4

Simulation results

In this chapter, the results obtained with the different LES models are presented and discussed. The predictive accuracy of these subgrid-scale formulations is evaluated through a detailed analysis of the mean flow topology and the turbulent flow statistics.

To obtain reliable results, the instantaneous flow fields are averaged both in time and along the spanwise direction. Specifically, the LESs are advanced in time up to a characteristic time of $60 T$, which corresponds to exactly 3 Flow-Through Times (FTTs), given the total domain length of $L_x = 20c$. The first $30 T$ (1.5 FTTs) are discarded to bypass the initial numerical transient and ensure that the flow has reached a fully developed, statistically stationary state. Subsequently, the remaining $30 T$ are employed to save the three-dimensional instantaneous flow fields at regular intervals of $0.5 T$, yielding an ensemble of 59 independent snapshots. The onset of this statistical sampling phase was deliberately chosen after an evaluation of the time histories of the aerodynamic coefficients (C_L and C_D), which confirmed that the flow field had fully stabilised around a steady mean value for all the tested LES models. By contrast, for the DNS baseline, the statistics were collected over a total simulated time of $18 T$, with instantaneous fields saved every $0.2 T$.

As anticipated in the previous chapter, the LES results are expected to exhibit some deviations with respect to the DNS target data. The different mathematical methodologies adopted to generate the upstream synthetic turbulence, alongside the resulting discrepancies in the average TI and ILSs, limit the feasibility of a strict quantitative correlation. Indeed, both of these synthetic eddy parameters exert a dominant physical influence on the boundary layer transition process and the subsequent reattachment of the separated shear layer [3]. Even if minimal, the differences induced by the inlet conditions cannot be neglected. Therefore, the comparison between the LESs and the DNS data will be predominantly qualitative, focusing on assessing whether the primary macroscopic flow structures characterising this specific aerodynamic scenario are correctly captured by the subgrid-scale approaches.

For the sake of notational simplicity, the filter symbol ($\bar{\cdot}$) identifying the resolved

flow field quantities is hereafter omitted throughout this chapter.

4.1 Instantaneous flow field

The instantaneous flow field provides a fundamental understanding of the main macroscopic features of the developed flow around the airfoil. To properly identify and visualise the coherent vortical structures, the Q-criterion is adopted. This criterion defines an eddy as a region characterised by a positive second invariant Q of the velocity gradient tensor $\nabla \mathbf{u}$ [17]. The second invariant Q is mathematically defined as:

$$Q = \frac{1}{2} [\text{Tr}(\mathbf{\Omega}\mathbf{\Omega}^T) - \text{Tr}(\mathbf{S}\mathbf{S}^T)] = \frac{1}{2} (\|\mathbf{\Omega}\|^2 - \|\mathbf{S}\|^2) \quad (4.1)$$

where $\mathbf{\Omega}$ is the antisymmetric rotation rate tensor and $\mathbf{S}$ is the symmetric strain rate tensor.

From a physical perspective, Q quantifies the local balance between the rotation rate and the deformation rate of the fluid. A strictly positive value ($Q > 0$) indicates that the rotation rate dominates the strain rate. As a result, the flow is locally governed by a swirling motion rather than by shearing or stretching, marking the presence of a coherent vortical structure.

A primary advantage of employing the Q-criterion rather than the simple vorticity magnitude ($|\boldsymbol{\omega}|$) is its robustness near the wall. In BLs, the vorticity magnitude reaches its peak at the solid wall due to the intense viscous shear, which does not necessarily imply a swirling motion. The Q-criterion overcomes this limitation since Q evaluates to zero exactly at a stationary wall [17]. This mathematical property makes it particularly suited for isolating true separated vortical structures from the attached BL shear on the airfoil surface.

In the present analysis, the coherent structures are visualised by extracting isosurfaces at a threshold value of $Q = 50$ at the characteristic time of $40T$ (2 FTTs). Furthermore, these isosurfaces are coloured according to the local instantaneous streamwise velocity component (Figure 4.1). The macroscopic flow structures are qualitatively consistent across all the LES models, which is to be expected since all computations simulate the identical flow configuration under the same operating conditions. In the very first portion of the airfoil suction side, a laminar free-shear layer is clearly identifiable. Moving downstream, this layer subsequently transitions into a fully turbulent state due to the Kelvin-Helmholtz instability driven by the sharp velocity gradient across the separated shear layer. This laminar-to-turbulent transition appears to occur earlier for the SITVM and ILES simulations. Moreover, the resolved vortical structures are noticeably larger for the DSM and ILES compared to the tensorial models (TVM and SITVM). Indeed, in accordance with the previous investigations on the BARC [28], the tensorial models demonstrate the capability to




Figure 4.1: Q-criterion isosurfaces ($Q = 50$) superimposed with contours of the instantaneous streamwise velocity component u .

reproduce a richer variety of small-scale flow structures, typical of non-linear models, primarily due to their low-dissipative nature. Further downstream, driven by the strong mean shear, the initial spanwise axes of these vortical structures gradually align with the streamwise direction. This kinematic process forms the well-known hairpin vortices, which subsequently continue to stretch until they eventually break down and cascade into smaller-scale turbulent structures with no preferential orientation.

4.2 Mean flow topology

Expected flow structures

At the adopted Reynolds number, the aerodynamic performance of the airfoil is heavily governed by the complex behaviour of the boundary layer, specifically involving separation, transition, and reattachment phenomena. When the attached laminar BL encounters the strong APG located just downstream of the suction

peak (minimum pressure point), it separates from the airfoil surface. This separated laminar shear layer is highly unstable and extremely receptive to the incoming freestream turbulence. Consequently, the transition to turbulence typically initiates a short distance downstream of the detachment point. Following this transition, the newly formed turbulent shear layer actively entrains high-momentum fluid from the external freestream. This intense turbulent mixing re-energises the near-wall flow, causing the shear layer to grow rapidly and endure the local pressure recovery, ultimately forcing the flow to reattach to the solid surface. The closed recirculating region bounded by the initial laminar separation and the subsequent turbulent reattachment is classically defined as a Separation Bubble (SB) [24]. Subsequently, as the fluid travels further downstream along the suction side, the fully developed turbulent BL encounters a sustained APG, eventually leading to a second separation. This downstream detachment forms the Primary Recirculation Region (PRR).

Simulated flow structures


The analysis of the mean streamwise velocity field, $U = \langle u \rangle$, provides valuable insight into the macroscopic flow topology predicted by the different LES approaches, facilitating a direct comparison with the DNS results (Figure 4.2).

The mean flow pattern reported for the LESs and the DNS generally confirms the expected macroscopic flow structures, but notable discrepancies emerge both between the LES models themselves and in their alignment with the DNS data. Specifically, all evaluated models successfully capture the formation of the SB at the suction side near the LE. However, a critical distinction arises when comparing the tensorial models against both the eddy-viscosity baseline (DSM) and the implicit approach (ILES).

Both the DSM and the ILES predict an additional flow structure enveloped by the SB, identified as a Secondary Recirculation Region (SRR) and driven by a reverse BL development. This secondary recirculation is characterised by near-zero, positive streamwise velocity values (visible as white contours within the separation region in Figures 4.2a and 4.2d). This structural similarity suggests that the inherent numerical dissipation of the ILES spatial discretisation acts on the mean flow similarly to the subgrid viscosity introduced by the Smagorinsky formulation.

Conversely, in accordance with the DNS target data and contrary to the previous numerical results obtained by Moroni [23], both tensorial approaches successfully avoid reproducing the SRR. This highlights the enhanced capability of the non-linear tensorial formulation in accurately resolving the complex near-wall physics inside the separated shear layer without introducing further recirculation zones.

Moving further downstream, the SB (represented by the blue contours highlighting negative velocities) effectively closes as the turbulent flow reattaches. Subsequently, the fully turbulent boundary layer evolves along the suction side until it encounters


Figure 4.2: Mean streamwise velocity component U .

the main separation point, giving rise to the PRR consistently captured by all the models.

The flow structures discussed so far can be further analysed by examining the average spanwise vorticity $\langle \omega_z \rangle$, as shown in Figure 4.3. This quantity is mathematically defined as:

$$\langle \omega_z \rangle = \left\langle \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right\rangle \quad (4.2)$$

The average spanwise vorticity provides further details on the SB topology, complementing the mean velocity field observations. Starting from the LE, the laminar BL separation is clearly captured by all the LES models, highlighted by the

Figure 4.3: Mean spanwise vorticity component $\langle \omega_z \rangle$.

intense negative vorticity strip (dark blue contours). However, moving downstream, the separated shear layer develops with notable differences between the models. Specifically, both the DSM and the ILES predict a notably thicker separated shear layer that deflects significantly away from the airfoil surface, resulting in a delayed turbulent reattachment further downstream. Conversely, the tensorial models (TVM and SITVM) exhibit a more coherent and thinner shear layer that remains remarkably closer to the wall, leading to a visibly earlier reattachment.

Furthermore, the near-wall region inside the SB is characterised by a reverse BL developing upstream toward the LE, identified by the thin layer of positive vorticity (red contours) shared by all the models. However, confirming the anticipation provided by the mean streamwise velocity contours, a crucial discrepancy emerges regarding the inner structure of the bubble. In Figures 4.3a and 4.3d, both the DSM

and the ILES reveal an embedded region of slightly negative vorticity near the wall, located beneath the main shear layer; this constitutes the clear footprint of the SRR. In contrast to the DSM and ILES approaches, and in strict accordance with the DNS data, both tensorial formulations maintain a continuous reverse BL without triggering this secondary recirculation.

These different near-wall features, including the turbulent reattachment position and the presence or absence of the SRR, will be further quantified and definitively confirmed by the analysis of the skin friction coefficient (C_f) distribution along the airfoil suction side surface, detailed in Section 4.3.

4.3 Wall stresses

Skin friction coefficient

The skin friction coefficient distribution allows for the characterisation of the boundary layer behaviour along the airfoil surface, distinctly highlighting the attached and separated flow regions. The skin friction coefficient is the non-dimensional version of the wall shear stress τ_w , and it is defined as:

$$C_f = \frac{\tau_w}{q_\infty} = \frac{\mu \frac{\partial u_{\parallel}}{\partial \eta}}{\frac{1}{2} \rho U_\infty^2} \quad (4.3)$$

where q_∞ is the freestream dynamic pressure, μ is the dynamic viscosity of the fluid, and $\partial u_{\parallel}/\partial \eta$ is the wall-normal gradient of the local wall-parallel velocity component.

In the present work, the wall shear stress τ_w is computed through the average velocity field evaluated at the centroid of the first cell adjacent to the wall. Specifically, the mean wall-parallel velocity is extracted as:

$$U_{\parallel} = \langle \bar{\mathbf{u}} \rangle \cdot \hat{t} \quad (4.4)$$

where $\langle \bar{\mathbf{u}} \rangle$ is the average-resolved velocity field at the first cell centroid and $\hat{t}$ is the unit vector tangent to the airfoil surface. Applying the no-slip condition at the wall, the wall shear stress is numerically approximated as:

$$\tau_w = \mu \frac{\Delta U_{\parallel}}{\Delta \eta/2} = \mu \frac{U_{\parallel}}{\Delta \eta/2} \quad (4.5)$$

where $\Delta \eta/2$ represents the normal distance from the wall to the first cell centroid (with $\Delta \eta$ being the local wall-normal dimension of the cell). This formulation intrinsically assumes a linear velocity profile between the solid boundary and the cell

centroid. From a physical perspective, this is valid since the near-wall mesh resolution ensures that the first cell lies entirely within the viscous sublayer. Furthermore, such an assumption is perfectly consistent with the second-order spatial accuracy of the cell-centred FVM employed, which relies on a linear interpolation scheme to reconstruct gradients at the cell faces.

In the study of separated flows, such as the one developing on the suction side of the NACA 4412 airfoil at a high AoA, the spatial evolution of the skin friction coefficient serves as a fundamental aerodynamic indicator. Specifically, the exact location where the C_f curve crosses the zero-axis, transitioning from positive to negative values, unequivocally identifies the BL mean separation point. Conversely, the point where the coefficient returns to positive values marks the mean reattachment point. Identifying these coordinates enables the evaluation of the spatial extent of the SB (and the SRR, where present), as well as the PRR.




Figure 4.4: Skin friction coefficient C_f distribution along the airfoil suction side.

In Figure 4.4, the results from the simulated LES models are directly compared against the high-fidelity DNS data. The macroscopic discrepancies between the modelling approaches become immediately evident, particularly concerning the spatial extent of the SB. While all evaluated methodologies successfully predict an early BL separation near the LE (characterised by the first zero-crossing of the C_f curve), their subsequent downstream behaviour differs drastically. Both the implicit approach (ILES) and the standard eddy-viscosity model (DSM) severely over-predict the length of the separated region. Furthermore, in both cases, a small region exhibiting a positive skin friction coefficient is observed for x/c approximately between 0.04 and 0.09, clearly highlighting the presence of the SRR. Consequently,

their mean turbulent reattachment points are pushed significantly downstream ($x/c > 0.2$) compared to the DNS target.

In sharp contrast, the tensorial models (TVM and SITVM) demonstrate a remarkable capability to capture the correct flow topology in the first portion of the airfoil, strictly aligning with the DNS data. Indeed, both models successfully suppress the positive C_f region, thus correctly avoiding the formation of the SRR, in full accordance with the DNS results.¹

Furthermore, the mean reattachment coordinate predicted by the tensorial approaches is in excellent agreement with the DNS data. Both the TVM and SITVM curves cross the zero-axis and return to positive values almost exactly at the DNS reattachment point. Following the SB, the tensorial models successfully recover the fully turbulent, attached BL regime. Although the absolute magnitude of the positive C_f peak is under-predicted compared to the DNS data, the short spatial extent of the separation bubble is well captured. Subsequently, the gradual decay of the friction coefficient towards the TE is correctly recovered by all simulations, including the DSM and the ILES, once their respective flow fields finally reattach.

Despite the excellent agreement in the leading-edge region, a significant discrepancy emerges between all the LES computations and the DNS regarding the mean turbulent separation point towards the TE. The DNS data indicates a mean separation at $x/c \approx 0.65$. Conversely, in the performed simulations, identifying the exact onset of this separation, and thus the start of the PRR, proves challenging due to residual statistical oscillations in the C_f profiles. Nevertheless, a clear macroscopic trend is visible: both tensorial models predict a premature mean separation, located between $x/c \approx 0.3$ and 0.4 . This is significantly upstream compared to both the DNS and the DSM and ILES methodologies, where this secondary separation is delayed between $x/c \approx 0.4$ and 0.6 .

Pressure coefficient

The pressure coefficient C_p is defined as:

$$C_p = \frac{P - P_\infty}{q_\infty} \quad (4.6)$$

where $P = \langle p \rangle$ is the average pressure at the first cell centroid adjacent to the wall, P_∞ is the mean reference freestream pressure evaluated at the inlet boundary, and q_∞ is the freestream dynamic pressure.

The spatial distribution of the pressure coefficient along the airfoil chord governs the overall aerodynamic force and provides deep insights into the BL evolution.

¹It is worth noting that the SITVM exhibits marginally positive C_f values extremely close to the LE ($x/c \approx 0.005$). This negligible deviation can be attributed to residual numerical noise inherently present in a region of strong geometric curvature, rather than to a physical flow feature.

Specifically, the presence of an APG drives flow separation, while a localised plateau in the C_p curve on the suction side represents the clear footprint of a SB. Inside this recirculation region, the pressure remains nearly constant until the separated flow transitions to turbulence and reattaches, triggering a sharp pressure recovery.




Figure 4.5: Pressure coefficient C_p distribution along the airfoil chord.

Figure 4.5 illustrates the pressure coefficient distribution for the simulated LES models, compared against the DNS data. In accordance with standard aerodynamic conventions, the vertical axis is inverted, meaning that higher suction (lower pressure) is plotted upwards.

Observing the pressure side of the airfoil (lower set of curves, $C_p > 0$), all numerical models exhibit excellent agreement with the DNS data. The flow in this region remains attached and predominantly laminar, allowing all subgrid closures, as well as the implicit approach (ILES), to accurately predict the smooth spatial pressure distribution.

Conversely, the suction side (upper set of curves, $C_p < 0$) reveals pronounced discrepancies, particularly in the leading-edge region, where the flow is subjected to intense acceleration and subsequent separation. The DNS curve is characterised by a distinct suction peak, immediately followed by a short pressure plateau. The TVM exhibits a severe over-prediction of this suction peak. The SITVM effectively mitigates this anomaly; although it still slightly overestimates the peak magnitude compared to the DNS, it recovers a much more realistic pressure distribution than

the baseline tensorial model. By contrast, the DSM and ILES fail to capture the sharpness of the suction peak, returning an under-predicted pressure drop.

Moving downstream of the suction peak, the C_p distribution perfectly corroborates the topological findings derived from the skin friction analysis. The DSM and ILES produce an extended pressure plateau that dominates the initial portion of the suction side. This flat pressure profile physically corresponds to the elongated SB previously observed in Figure 4.4. In stark contrast, both the TVM and SITVM correctly capture the short spatial duration of the pressure plateau, demonstrating a rapid transition to turbulence and a prompt pressure recovery.

Finally, within the PRR extending towards the TE, all numerical models tend to converge. They successfully capture the gradual pressure rise in the aft portion of the airfoil, despite the different mean turbulent separation locations discussed previously.

4.4 Integral quantities and time scales

The overall behaviour of external flows interacting with solid bodies can be further investigated by evaluating the dimensionless aerodynamic coefficients of drag and lift, C_D and C_L , respectively. These coefficients are defined as:

$$C_D = \frac{F_x}{q_\infty A}, \quad C_L = \frac{F_y}{q_\infty A} \quad (4.7)$$

where F_x and F_y are the streamwise and vertical aerodynamic force components respectively, q_∞ is the freestream dynamic pressure, and $A = bc$ represents the reference planform area of the airfoil. To ensure the initial transient phase is completely discarded, the time histories of the aerodynamic coefficients are sampled after $30T$, adopting a constant sampling interval of $T_s = 10 \Delta t = 0.01T$.

Table 4.1 summarises the time-averaged aerodynamic coefficients, their relative deviations from the DNS data, and their standard deviations (σ_{C_D} and σ_{C_L}), which quantify the unsteadiness of the global aerodynamic loads.

	$\langle C_D \rangle$	$\langle C_L \rangle$	ΔC_D [%]	ΔC_L [%]	σ_{C_D}	σ_{C_L}
DSM	0.121	1.481	+25.1	+5.4	0.0078	0.0340
TVM	0.079	1.555	-18.6	+10.7	0.0059	0.0214
SITVM	0.097	1.484	-0.4	+5.6	0.0051	0.0241
ILES	0.120	1.489	+23.4	+6.0	0.0074	0.0348
DNS	0.097	1.405	-	-	0.0018	0.0098

Table 4.1: Aerodynamic coefficient statistics.

As a general trend, all evaluated LES models over-predict the mean lift coefficient $\langle C_L \rangle$ compared to the DNS. The TVM exhibits the largest deviation; this is likely driven by the severe over-prediction of the leading-edge suction peak previously discussed. Conversely, the DSM, SITVM, and ILES provide closer estimates with tightly clustered relative errors.

The most significant discrepancies among the models arise in the drag prediction. At 15° AoA, the total drag is predominantly governed by the pressure force (form drag) rather than skin friction. Thus, $\langle C_D \rangle$ is highly sensitive to variations in the pressure distribution along the entire chord. Specifically, in the leading-edge region, the intense suction peak acts on the highly curved, forward-facing surface of the profile. While it primarily contributes to the overall lift generation, it simultaneously generates a forward-pointing force component (the so-called leading-edge thrust) that counteracts a portion of the drag [4]. The SITVM correctly captures the magnitude of this peak and the rapid closure of the laminar SB, achieving a highly accurate mean drag estimation. Conversely, the TVM severely over-predicts the suction peak. While likely not the sole mechanism at play, this localised excessive suction produces a forward pull that can be identified as a contributing factor to its under-predicted total drag. In sharp contrast, the DSM and the ILES fail to capture the flow reattachment point compared to the DNS, resulting in an elongated SB. By smoothing out the sharp suction peak and distributing a moderate suction (the pressure plateau) over a flatter section of the airfoil, which is nearly parallel to the freestream direction, it contributes mainly to the lift but fails to generate a forward-pointing leading-edge thrust. Interestingly, this leading-edge thrust deficit seems to dominate the overall drag production. As clearly visible in the C_p distribution (Figure 4.5), beyond $x/c \approx 0.25$ all curves practically overlap, indicating that the pressure recovery towards the TE is remarkably similar across all models. Consequently, while the delayed turbulent separation predicted by the DSM and ILES should theoretically reduce form drag, this benefit is completely cancelled out by the massive drag penalty generated at the LE. Therefore, while other mechanisms may be involved, this lack of a compensatory forward pull represents a driver for the over-prediction of the total drag observed for the eddy-viscosity and implicit formulations.

Regarding the signal unsteadiness, SITVM and TVM yield the lowest standard deviations, indicating a relatively stable separated region. Conversely, DSM and ILES produce higher fluctuation levels.

To further investigate the dynamic nature of this flow configuration, the Power Spectral Density (PSD) of the lift coefficient is computed by applying the Fast Fourier Transform (FFT) to the signal:

$$S_{C_L} = \widehat{C_L} \widehat{C_L}^* \quad (4.8)$$

where $\widehat{C}_L$ is the Fourier transform of the lift coefficient time history, and $\widehat{C}_L^*$ is its complex conjugate. In Figure 4.6, the PSD is plotted against the non-dimensional frequency, the modified Strouhal number St , which can be computed as [18]:

$$St = \frac{f c \sin(\alpha)}{U_\infty} \quad (4.9)$$

where f is the dimensional frequency [Hz].




Figure 4.6: Lift coefficient spectra.

The resulting lift coefficient spectra for the evaluated LES approaches are presented in Figure 4.6. As a point of reference, the baseline DNS data exhibits a visible energetic peak corresponding to a vortex shedding modified Strouhal number of $St = 0.963$. In contrast, the spectra obtained from the present simulations do not exhibit a distinct, highly energetic peak, but rather a broadband distribution. Due to the lack of an evident dominant frequency, a specific vortex shedding modified Strouhal number is not extracted for the evaluated LES formulations.

This broadband behaviour indicates the absence of a purely periodic von Kármán-like vortex shedding mechanism in the simulated flow field. Unlike perfectly symmetric blunt bodies, such as the BARC geometry analysed by Tessier [28], where, except

for the TVM, a dominant shedding frequency clearly drives the alternating lift fluctuations around a near-zero mean, the flow around the NACA 4412 at a high AoA is characterised by strong geometrical and aerodynamic asymmetries.

Furthermore, the unsteady separation onset continuously alters the effective aerodynamic shape of the body and the thickness of the separated wake. This highly dynamic behaviour effectively smears the shedding frequencies over a wider band, masking the underlying periodicity observed in the DNS. This broadband energy distribution is in excellent agreement with the findings reported by Moroni [23] for an analogous flow configuration, where the highly chaotic separated region similarly hindered the precise identification of a single shedding frequency.

4.5 Turbulent kinetic energy and Reynolds stresses


The mean Turbulent Kinetic Energy (TKE) distribution around the airfoil, specifically on the suction side, provides fundamental physical insights into the flow topology and the energetic dynamics predicted by the analysed LES models. The TKE represents the kinetic energy content of the turbulent fluctuations of the resolved flow field, and it is defined as:

$$\langle k \rangle = \frac{1}{2} (\langle u'^2 \rangle + \langle v'^2 \rangle + \langle w'^2 \rangle) \quad (4.10)$$

where $\langle u'^2 \rangle$, $\langle v'^2 \rangle$, and $\langle w'^2 \rangle$ are the variances of the resolved velocity components, corresponding to the trace of the resolved Reynolds stress tensor.

As depicted in Figure 4.7, all LES models successfully capture the transition process within the separated shear layer, highlighted by a sudden increase in resolved TKE downstream of the LE. Furthermore, the near-wall region inside the SB is characterised by near-zero TKE values (white contours) across all models, confirming the laminar nature of the reversed BL in this area.

Despite these qualitative similarities, notable quantitative discrepancies emerge regarding the intensity and spatial distribution of the turbulent fluctuations. Both the DSM and the ILES predict the highest TKE levels (with peaks reaching $\langle k \rangle \approx 0.40$), developing a thick shear layer that noticeably deflects away from the solid wall. Conversely, the TVM shows a significant reduction in the resolved TKE peak ($\langle k \rangle \approx 0.25$) and preserves a visibly thinner shear layer that remains remarkably closer to the airfoil surface. Interestingly, the SITVM exhibits an intermediate behaviour: while it successfully maintains the thin and attached shear layer topology of the TVM, it predicts higher TKE peak values, energetically positioning itself between the DSM and ILES approaches and the baseline tensorial formulation. This intermediate energetic state is directly tied to the mathematical treatment of the velocity gradients within the respective subgrid closures. The TVM evaluates the subgrid viscosity using the total filtered velocity gradients. Consequently, in regions

Figure 4.7: Mean resolved Turbulent Kinetic Energy (TKE) $\langle k \rangle$.

dominated by an intense mean shear, such as the separated shear layer bounding the SB, it tends to predict large subgrid stresses, leading to a high dissipation of the resolved turbulent fluctuations. In contrast, the SITVM is specifically designed to subtract the mean filtered velocity gradients from the total filtered ones. By neutralising the influence of the intense mean shear on the subgrid stress production, the SITVM mitigates the excessive dissipative behaviour that characterises the TVM, allowing the model to retain a higher amount of resolved TKE.

This overall reduction in the resolved TKE for the TVM and SITVM models indicates that the tensorial approaches distribute turbulent energy differently. Since the resolved fluctuations are visibly damped compared to the DSM and the ILES, it is highly likely that the tensorial formulations compute higher subgrid stresses. These stresses act as a damping mechanism for the resolved fluctuations, directly

limiting the intensity of the large scales.

The analysis of the resolved turbulent fluctuations is completed by evaluating the primary resolved Reynolds shear stress component, $\langle u'v' \rangle$, reported in Figure 4.8.




Figure 4.8: Resolved Reynolds shear stress component $\langle u'v' \rangle$.

Consistent with the TKE observations, all LES models capture the intense turbulent activity within the separated shear layer, identified by a distinct region of negative $\langle u'v' \rangle$ values (blue contours). However, notable quantitative differences among the models are evident. Both the DSM and the ILES predict the highest magnitude of resolved shear stress, distributing it over a noticeably thicker layer that extends further from the wall. Conversely, the TVM significantly damps the resolved shear stress intensity, confining it to a thinner region that strictly follows the airfoil profile. As previously observed for the mean TKE, the SITVM confirms

its intermediate nature: it retains a thin shear layer topology, similar to the TVM, but recovers slightly higher stress magnitudes. This confirms that subtracting the intense mean shear in the subgrid evaluation not only preserves the resolved TKE, but also prevents the excessive damping of the resolved Reynolds shear stress.


This coherent reduction in both normal (quantified by $\langle k \rangle$) and shear ($\langle u'v' \rangle$) resolved Reynolds stresses for the tensorial approaches strongly corroborates the physical hypothesis introduced earlier. The non-linear tensorial formulations appear to generate locally higher subgrid stresses ($\langle \tau_{ij} \rangle$) within the intense shear region. This aspect can explain the lower mean TKE and Reynolds shear stress observed for the tensorial models, since these intense subgrid stresses act as a damping mechanism for the resolved fluctuations compared to both the standard eddy-viscosity model and the implicit approach.

4.6 Subgrid stresses

The spatial distribution and intensity of the subgrid stresses differ significantly among the studied LES models. The most energetic components of the subgrid stress tensor, that are the normal components $\langle \tau_{11} \rangle$ and $\langle \tau_{22} \rangle$, and the shear component $\langle \tau_{12} \rangle$, are analysed.

As anticipated in Section 4.5, the tensorial models predict more intense subgrid stresses compared to the DSM. This behaviour quantitatively confirms the physical hypothesis derived from the analysis of the resolved TKE and the Reynolds shear stress $\langle u'v' \rangle$. Indeed, the higher magnitude of the subgrid stresses generated by both tensorial approaches exerts a stronger damping effect on the resolved large scales. This enhanced damping mechanism directly results in the lower intensity of the resolved fluctuations observed previously.

The discrepancies between the eddy-viscosity and the tensorial approaches are evident, particularly concerning the normal subgrid stresses (Figures 4.9 and 4.10). The first fundamental difference emerges when observing the DSM (Figure 4.9a and Figure 4.10a). This model relies on the Boussinesq hypothesis, which assumes that the subgrid stress tensor is aligned to the resolved strain rate tensor. Because the diagonal strain rate components are much smaller than the off-diagonal ones within a separated shear layer, this linear limitation of the eddy-viscosity models severely hinders their ability to capture the subgrid anisotropy. Consequently, the DSM significantly underpredicts the intensity of the normal subgrid stresses. As visually confirmed by the predominantly white contour maps, the model returns near-zero values for both normal subgrid stresses across most of the domain. Specifically, the near-wall region presents near-zero negative values for $\langle \tau_{11} \rangle$, shifting to strictly positive values within the separated shear layer and the SB region. A similar trend,


Figure 4.9: Subgrid stress component $\langle \tau_{11} \rangle$.

albeit with an opposite sign, is observed for the component $\langle \tau_{22} \rangle$.

In stark contrast, the tensorial models (TVM and SITVM) abandon the Boussinesq hypothesis. Being non-linear models, they natively account for the grid-induced anisotropy through the grid inertia tensor. Because of this architectural difference, they are able to predict highly intense diagonal subgrid stresses, as depicted in Figures 4.9b, 4.9c, 4.10b, and 4.10c.

However, when comparing the TVM and the SITVM, the impact of the mathematical treatment of the filtered velocity gradients (discussed in Section 4.5) manifests in more subtle, yet physically crucial, differences. While the global behaviour of both normal subgrid stress components ($\langle \tau_{11} \rangle$ and $\langle \tau_{22} \rangle$) appears structurally similar for the two tensorial models, two distinct features emerge. Firstly, the TVM exhibits a narrower band of positive normal stresses along the suction side shear layer. Secondly, and most importantly, the SITVM correctly predicts near-zero normal subgrid stresses along the laminar pressure side. This perfectly aligns with the theoretical formulation of the SITVM: by computing the subgrid stresses based solely on the fluctuating filtered velocity gradients, it successfully prevents activation in laminar regions subjected to mean shear.

Moreover, the diagonal components of the subgrid stress tensor define the generalised subgrid turbulent kinetic energy, k_{sgs} , originally formulated within the filtering


Figure 4.10: Subgrid stress component $\langle \tau_{22} \rangle$.

approach proposed by Germano [10]:

$$k_{sgs} = \frac{1}{2}(\tau_{11} + \tau_{22} + \tau_{33}) \quad (4.11)$$

According to [29], it is highly desirable for this quantity to remain strictly non-negative, since it physically represents the kinetic energy content of the unresolved scales. For the eddy-viscosity baseline, the sum of the normal subgrid stresses mathematically evaluates to zero. As previously noted, the individual diagonal components exhibit both positive and negative values (e.g., negative $\langle \tau_{11} \rangle$ near the wall and negative $\langle \tau_{22} \rangle$ in the shear layer) that cancel each other out. This yields a zero subgrid turbulent kinetic energy ($k_{sgs} = 0$), which unphysically contradicts its definition as a measure of the energy content of the unresolved scales. In stark contrast, the tensorial formulations overcome this limitation in a more physically consistent manner. By natively abandoning the Boussinesq assumption, they estimate strictly positive values for the normal stresses, thus ensuring a positive k_{sgs} throughout the fluid domain and guaranteeing the physical realizability of the unresolved kinetic energy.

Moving the analysis to the off-diagonal component, Figure 4.11 illustrates the averaged subgrid shear stress $\langle \tau_{12} \rangle$. In sharp contrast to the normal stresses, the DSM (Figure 4.11a) demonstrates a significant and physically coherent activation. A

Figure 4.11: Subgrid stress component $\langle \tau_{12} \rangle$.

distinct band of negative values (blue contours) is clearly visible along the separated shear layer. This behaviour aligns with the theoretical expectations of the Boussinesq hypothesis since the resolved off-diagonal strain rate is the dominant component within a shear layer. Furthermore, the DSM shows positive values of the shear subgrid stress in the near-wall region of the SB.

Likewise, the tensorial models, TVM and SITVM (Figures 4.11b and 4.11c), exhibit a spatial distribution for $\langle \tau_{12} \rangle$ that is more similar to the DSM compared to the normal subgrid stresses. A distinctive feature of the tensorial approaches is the presence of an intense, highly localised region of positive $\langle \tau_{12} \rangle$ (deep red contours) exactly at the leading-edge region, followed by negative values (blue contours) in the separated shear layer which persist as the flow moves downstream.

Finally, when comparing the TVM and the SITVM for this specific off-diagonal component, the global discrepancies in the separated shear layer are minimal, with both models returning comparable magnitudes and spatial distributions. However, perfectly consistent with the normal components, the SITVM successfully predicts near-zero values along the laminar pressure side, suppressing the non-physical activation observed in the TVM. This indicates that while the subtraction of the mean filtered velocity gradients in the SITVM formulation has a less pronounced impact on the subgrid shear stress above the airfoil suction side, it remains crucial for maintaining physical consistency in attached laminar BL.

Near-wall behaviour

The near-wall behaviour of the subgrid stresses provides crucial insights into the numerical stability of the tensorial models. Indeed, both the TVM and SITVM exhibit a near-wall asymptotic behaviour that differs fundamentally from the DSM. At the airfoil surface, the no-slip condition is applied to the resolved flow field (see Subsection 3.1.3); therefore, the filtered strain rate tensor simplifies to:

$$\bar{S} = \frac{1}{2} \begin{bmatrix} 0 & \frac{\partial \bar{u}}{\partial \eta} & 0 \\ \frac{\partial \bar{u}}{\partial \eta} & 0 & \frac{\partial \bar{w}}{\partial \eta} \\ 0 & \frac{\partial \bar{w}}{\partial \eta} & 0 \end{bmatrix} \quad (4.12)$$

Consequently, recalling equation (2.17), the subgrid stresses at the wall for the DSM evaluate to:

$$\begin{cases} \tau_{11,w} = \tau_{22,w} = \tau_{33,w} = \tau_{13,w} = 0 \\ \tau_{12,w} = -\nu_{sgs} \frac{\partial \bar{u}}{\partial \eta} \\ \tau_{23,w} = -\nu_{sgs} \frac{\partial \bar{w}}{\partial \eta} \end{cases} \quad (4.13)$$

where the subscripts 1, 2 and 3 denote the tangential, normal and spanwise directions at the airfoil surface, respectively. Thus, for the diagonal components of the subgrid stress tensor and $\tau_{13,w}$, the correct physical behaviour at the wall is guaranteed by the null components of the filtered strain rate tensor. For the remaining components, in order to achieve null results to represent the absence of fluctuations in the viscous sublayer, the subgrid viscosity must be zero. As already explained in Chapter 2, this aspect is ensured through the dynamic procedure applied to the SM.

On the other hand, the TVM unphysically predicts non-zero components of the subgrid stress tensor at the wall. By expanding the tensorial formulation (2.39) and substituting the non-zero filtered velocity gradients from the filtered strain rate

tensor (4.12), the subgrid stresses at the wall become:

$$\begin{cases} \tau_{22,w} = \tau_{12,w} = \tau_{23,w} = 0 \\ \tau_{11,w} = -\frac{\mathcal{I}'_{22}}{\Omega} \left(\frac{\partial \bar{u}}{\partial \eta} \right)^2 \\ \tau_{33,w} = -\frac{\mathcal{I}'_{22}}{\Omega} \left(\frac{\partial \bar{w}}{\partial \eta} \right)^2 \\ \tau_{13,w} = -\frac{\mathcal{I}'_{22}}{\Omega} \left(\frac{\partial \bar{u}}{\partial \eta} \frac{\partial \bar{w}}{\partial \eta} \right) \end{cases} \quad (4.14)$$

This demonstrates that several components do not naturally vanish at the wall, introducing spurious stresses that act as a primary source of numerical instability. Assuming rectangular hexahedral grid elements at the wall, the modified inertia tensor component corresponding to the wall-normal direction is:

$$\mathcal{I}'_{22} = -\frac{\Omega}{12} \Delta \eta^2 \quad (4.15)$$

where $\Omega = \Delta \xi \Delta \eta \Delta z$ is the cell volume, with $\Delta \xi$, $\Delta \eta$, and Δz being the local tangential, wall-normal, and spanwise grid dimensions, respectively. Following the same discretisation approach used in equation (4.5) for the wall-normal filtered velocity gradient, the subgrid stress component $\tau_{11,w}$ can be approximated as:

$$\tau_{11,w} = \frac{u_{\parallel,w}^2}{3} \quad (4.16)$$

where $u_{\parallel,w}$ is the resolved wall-parallel velocity evaluated at the first cell centroid. This formulation highlights a critical numerical trade-off. This subgrid stress is directly proportional to the square of the resolved parallel velocity. Consequently, to minimise $\tau_{11,w}$, $u_{\parallel,w}$ must be as close to zero as possible, which requires placing the first cell centroid extremely close to the wall (i.e., achieving a very low y^+ target). However, reducing y^+ while maintaining a constant wall-parallel resolution ($\Delta \xi$) inevitably increases the cell ARs. This conflicts with the stability requirements of the gradient solving schemes, which severely degrade at high ARs. Although reducing $\Delta \xi$ effectively mitigates the AR growth, it comes at the expense of a higher, yet feasible, computational effort. Note that this analysis implicitly assumes a constant spanwise dimension Δz .

Similarly, the remaining non-zero components for the tensorial model can be discretised at the wall as:

$$\tau_{33,w} = \frac{w_w^2}{3} \quad (4.17)$$

and

$$\tau_{13,w} = \frac{u_{\parallel,w} w_w}{3} \quad (4.18)$$

where w_w is the resolved spanwise velocity at the first cell. In the context of the present flow configuration, the $\tau_{11,w}$ component is the most likely trigger for numerical instabilities, as the resolved wall-parallel velocity gradient heavily dominates the near-wall flow field. This hypothesis is confirmed by the 2D field of the subgrid stress $\langle \tau_{11} \rangle$ (see Figure 4.9).

To overcome these near-wall limitations and the associated numerical constraints, the SITVM introduces a physical correction based on the mean filtered flow field (see Section 2.5). Applying the Shear-Improved formulation to the near-wall region, the subgrid stresses become:

$$\begin{cases} \tau_{22,w} = \tau_{12,w} = \tau_{23,w} = 0 \\ \tau_{11,w} = -\frac{\mathcal{I}'_{22}}{\Omega} \left(\frac{\partial \bar{u}}{\partial \eta} - \frac{\partial \langle \bar{u} \rangle}{\partial \eta} \right) \frac{\partial \bar{u}}{\partial \eta} \\ \tau_{33,w} = -\frac{\mathcal{I}'_{22}}{\Omega} \left(\frac{\partial \bar{w}}{\partial \eta} \right)^2 \\ \tau_{13,w} = -\frac{\mathcal{I}'_{22}}{2\Omega} \left[\frac{\partial \bar{u}}{\partial \eta} \frac{\partial \bar{w}}{\partial \eta} + \left(\frac{\partial \bar{u}}{\partial \eta} - \frac{\partial \langle \bar{u} \rangle}{\partial \eta} \right) \frac{\partial \bar{w}}{\partial \eta} \right] \end{cases} \quad (4.19)$$

Following the same discretisation approach previously adopted, the subgrid stress component $\tau_{11,w}$ for the SITVM evaluates to:

$$\tau_{11,w} = \frac{u_{\parallel,w}(u_{\parallel,w} - U_{\parallel,w})}{3} \quad (4.20)$$

where $U_{\parallel,w}$ represents the resolved mean wall-parallel velocity evaluated at the first cell centroid. From a mathematical and physical standpoint, this formulation yields a significantly attenuated subgrid stress compared to the standard TVM. Within the viscous sublayer adjacent to the solid boundary, the resolved velocity closely tracks the mean resolved velocity profile. Consequently, the resolved fluctuation term $(u_{\parallel,w} - U_{\parallel,w})$ asymptotically approaches zero. This mechanism suppresses the subgrid stress $\tau_{11,w}$, allowing the model to recover a physically coherent near-wall behaviour.

In contrast to the aforementioned subgrid stress component, $\tau_{33,w}$ has the same mathematical formulation both for the TVM and SITVM since, for the present flow configuration, the mean resolved spanwise velocity is identically zero. Therefore, this diagonal subgrid stress component is:

$$\tau_{33,w} = \frac{w_w^2}{3} \quad (4.21)$$

Thus, the SITVM naturally reduces to the baseline model, isolating its corrective action exclusively to the wall-parallel and wall-normal directions, governing the BL development.

Finally, for the off-diagonal subgrid stress component, following the same discretisation approach adopted for $\tau_{11,w}$, the equation evaluates to:

$$\tau_{13,w} = \frac{w_w[u_{\parallel,w} + (u_{\parallel,w} - U_{\parallel,w})]}{6} = \frac{w_w(2u_{\parallel,w} - U_{\parallel,w})}{6} \quad (4.22)$$

This result reveals a distinctive asymptotic behaviour of the SITVM for the off-diagonal subgrid stress component. Unlike the diagonal component $\tau_{11,w}$, where the streamwise fluctuation term $(u_{\parallel,w} - U_{\parallel,w})$ acts as a global multiplier that forces the entire subgrid stress to vanish as the velocity approaches the mean profile, the $\tau_{13,w}$ component behaves differently. Indeed, since the term $(u_{\parallel,w} - U_{\parallel,w})$ vanishes near the wall, the SITVM formulation $\tau_{13,w}$ tends to $(u_{\parallel,w}w_w)/6$. Compared to the standard TVM, which intrinsically predicts a stress approaching $(u_{\parallel,w}w_w)/3$, the SITVM effectively halves the magnitude of the unphysical off-diagonal subgrid stress near the wall. While it does not suppress the term entirely, this significant reduction contributes to mitigating the near-wall numerical instabilities.

4.7 Subgrid dissipation rates

A relevant feature that a subgrid modelling approach should reproduce is the correct energy transfer between the resolved large scales and the unresolved small ones. In a LES context, the resolved kinetic energy is defined as:

$$k^f = \frac{1}{2} \overline{u_i} \overline{u_i} \quad (4.23)$$

Multiplying each term of the filtered momentum equation in (2.4) by $\overline{u_i}$, it is possible to derive the following resolved kinetic energy equation:

$$\frac{\partial k^f}{\partial t} + \frac{\partial k^f \overline{u_j}}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \overline{p}}{\partial x_i} \overline{u_i} + \nu \frac{\partial^2 k^f}{\partial x_j \partial x_j} - \epsilon^f - \frac{\partial \tau_{ij} \overline{u_i}}{\partial x_j} - \epsilon_{sgs} \quad (4.24)$$

Focusing on the turbulent dissipation terms, two main components can be distinguished: the filtered dissipation rate

$$\epsilon^f = \nu \frac{\partial \overline{u_i}}{\partial x_j} \frac{\partial \overline{u_i}}{\partial x_j} \quad (4.25)$$

which represents how much resolved kinetic energy is dissipated at large resolved scales, and the subgrid dissipation rate

$$\epsilon_{sgs} = -\tau_{ij} \frac{\partial \overline{u_i}}{\partial x_j} = -\tau_{ij} \overline{S_{ij}} \quad (4.26)$$


which represents the resolved kinetic energy transfer between the resolved large scales and the modelled ones. Contrary to the filtered dissipation rate, which is always a sink term, the subgrid dissipation rate can be both a source (negative) or sink (positive) term. Indeed, in the near-wall region, the turbulence production occurs at intermediate scales, so depending on the filter length Δ (and therefore on the mesh discretisation), there could be an energy transfer towards both the small and large scales. The energy transfer towards small scales is the well-known forward energy cascade, while the energy transfer towards large scales is the reverse energy cascade, also known as the *backscatter* phenomenon.

The subgrid dissipation rate ϵ_{sgs} can be further decomposed into the subgrid dissipation rate of mean kinetic energy and the subgrid dissipation rate of turbulent kinetic energy of the resolved scales. By applying the Reynolds decomposition and the average operator $\langle \cdot \rangle$ to the subgrid stress tensor and to the filtered strain rate tensor, it is possible to show that:

$$\epsilon_{sgs} = -\langle \tau_{ij} \rangle \langle \overline{S_{ij}} \rangle + \langle -\tau'_{ij} \overline{S'_{ij}} \rangle = E_{sgs} + \epsilon'_{sgs} \quad (4.27)$$



thus making explicit the decomposition of the subgrid dissipation rate. Both decomposition terms are positive throughout most of the relevant fluid regions, meaning that the mean kinetic energy is drained to feed the subgrid-scale turbulent motion, and turbulent kinetic energy cascades from the resolved large scales towards the modelled small ones. To thoroughly assess the predictive capabilities of the selected models, particularly their ability to capture local backscatter phenomena in separating flows, the spatial distributions of the total subgrid dissipation rate and its components (Figures 4.12, 4.13, and 4.14) are analysed in detail below.

As observed in Figure 4.12a, the DSM predicts exclusively positive values of ϵ_{sgs} , which are densely concentrated within the separated shear layer near the LE where transition occurs. This strictly dissipative behaviour is a direct mathematical consequence of the Boussinesq hypothesis. By constraining the subgrid stress tensor to align perfectly with the resolved strain rate tensor, the model acts purely as an energy sink, forcing a forward energy cascade regardless of the local flow physics. Conversely, the TVM (Figure 4.12b) reveals a negative subgrid dissipation rate in the initial portion of the shear layer prior to transition. This successfully reproduces the energy backscatter that is entirely missed by the DSM. The SITVM (Figure 4.12c), however, exhibits a distinct behaviour compared to its baseline counterpart. The Shear-Improved version predicts positive values even in the initial stages of the shear layer, precisely in the region immediately above where the TVM exhibited negative values. Furthermore, only a few localised, low-intensity negative spots are present in regions where the TVM previously captured strong backscatter. Ultimately, both tensorial models, due to their non-linear formulations, yield a generally lower overall dissipative effect. This physically bridges the observations from Section 4.6: although

Figure 4.12: Total subgrid dissipation rate ϵ_{sgs} .

the tensorial models generate highly intense subgrid stresses, these non-linear stresses are not forced to align with the resolved strain rate tensor. Consequently, the tensorial contraction $-\tau_{ij}\overline{S}_{ij}$ can yield low or negative values (backscatter) even in the presence of strong subgrid activity. This confirms that the intense subgrid stresses in the TVM and SITVM act primarily as an anisotropic damping mechanism for the momentum, rather than as a purely dissipative energy sink. Furthermore, all three models predict positive total subgrid dissipation rates in the first portion of the very near-wall region of the reversed BL.

Decomposing the total dissipation rate yields deeper physical insight into the predictive behaviour of the models. Figure 4.13 illustrates the subgrid dissipation rate of the mean kinetic energy, E_{sgs} . In separating and reattaching flows, it is well established that the transitional region of the shear layer immediately downstream of the LE exhibits negative turbulence production [5]. Within these regions, the fluctuating field transfers energy back to the mean field; thus, the subgrid scales act physically as a source of mean kinetic energy ($E_{sgs} < 0$). The non-linear formulation of the TVM successfully captures these localised source regions with significantly greater intensity than the SITVM, mirroring the trends previously observed for ϵ_{sgs} . In addition to the separated shear layer, a closer inspection reveals that both tensorial models predict slightly negative values of E_{sgs} in the initial portion of the near-wall reversed BL. This indicates a weak, localised transfer of energy back to the mean

Figure 4.13: Subgrid dissipation rate of mean kinetic energy E_{sgs} .Figure 4.14: Subgrid dissipation rate of turbulent kinetic energy ϵ'_{sgs} .

field within the recirculation zone. Conversely, the DSM introduces a robust and unphysical sink of mean kinetic energy, driven by the intense shear in the transitional region. This excessive energy dissipation can artificially distort the development of the shear layer, thereby compromising the prediction of downstream flow dynamics, including the shape and extent of the recirculation bubble, as evidenced by the skin friction coefficient distribution (see Figure 4.4).

Furthermore, Figure 4.14 illustrates the subgrid dissipation rate of the turbulent kinetic energy, ϵ'_{sgs} , which characterises the energy cascade at the filter scale. Unlike the DSM, both the TVM and SITVM are capable of capturing localised regions of backscatter ($\epsilon'_{sgs} < 0$) within the separated shear layer. It should be noted that, by comparing the two decomposed terms, the main contribution to the total backscatter in the shear layer is represented by the E_{sgs} component. Ultimately, the complete inability of the DSM to reproduce the inverse energy transfer in either of these components underscores its inherent physical limitations when applied to transitional flow regimes.

Chapter 5

Conclusions

5.1 Accuracy

The investigation of the complex aerodynamic flow around the NACA 4412 airfoil at a high AoA ($\alpha = 15^\circ$) and a chord Reynolds number of $Re_c = 66\,667$ has highlighted major physical advantages as well as inherent limitations of the novel non-linear Tensorial Viscosity Model (TVM) and its Shear-Improved variant (SITVM). Their predictive accuracy has been evaluated against the traditional linear Dynamic Smagorinsky Model (DSM) and an Implicit Large Eddy Simulation (ILES) approach, using high-fidelity Direct Numerical Simulation (DNS) data as baseline. However, it is worth noting that a strict quantitative correlation with the DNS dataset is hindered by the different synthetic eddy generation methodologies adopted, which allowed for an accurate matching of the inflow TI but not of the ILSs. Consequently, the DNS data is primarily employed to assess whether the macroscopic flow structures are correctly captured.

Flow topology and boundary layer dynamics

In the leading-edge region of the airfoil suction side, the flow is characterised by a laminar BL separation under the influence of an intense APG. This separated laminar shear layer undergoes transition to turbulence and subsequently reattaches to the solid surface, enclosing a laminar SB. Both the traditional eddy-viscosity model (DSM) and the implicit approach (ILES) severely over-predict the spatial extent of this recirculating region, pushing the mean turbulent reattachment point significantly downstream compared to the DNS framework. Furthermore, within the SB, the DSM and ILES predict a flow structure identified as a SRR, driven by an reverse BL development.

In sharp contrast, the tensorial formulations (TVM and SITVM) demonstrate a remarkable capability to capture the correct flow topology in the initial portion of the

airfoil chord. Both models, in strict accordance with the high-fidelity DNS results, successfully suppress the formation of the SRR. Owing to their low-dissipative nature, the tensorial approaches preserve a thinner and more coherent separated shear layer that remains significantly closer to the airfoil surface, ensuring a prompt transition to turbulence and an accurate prediction of the mean reattachment coordinate.

Progressing further downstream along the suction side, the fully developed turbulent BL encounters a sustained APG, which eventually induces a permanent detachment known as the PRR. While the leading-edge separation bubble dynamics are accurately resolved by the tensorial models, the prediction of the downstream permanent separation represents a shared challenge. Compared to the DNS data, both the TVM and SITVM predict a notably premature detachment. This early separation occurs significantly upstream of the DNS baseline and also stands in contrast to the DSM and ILES methodologies, where the onset of the PRR is appropriately delayed and much closer to the DNS behaviour. This discrepancy underscores that, despite their advanced mathematical structure, current tensorial subgrid-scale models still face difficulties in properly balancing the effects of moderate but prolonged APGs in fully turbulent regimes.

Wall stresses and aerodynamic loads

The macroscopic features of the BL development are directly reflected in the local wall stresses and the resulting global aerodynamic performance. The spatial distribution of the pressure coefficient (C_p) confirms that the standard TVM suffers from a severe over-prediction of the leading-edge suction peak. This localised numerical anomaly exerts a noticeable impact on the total drag calculation; while other mechanisms are likely involved, the excessive suction acting on the forward-facing, highly curved surface of the LE generates an artificial thrust component that contributes to counteracting the form drag. This ultimately results in a substantial under-prediction of the mean drag coefficient.

The Shear-Improved variant elegantly mitigates this limitation. By explicitly subtracting the mean filtered velocity gradients from the instantaneous field, the SITVM eliminates the subgrid viscosity generated by the strong mean shear in the attached laminar BL. Without this viscous blockage distorting the local flow acceleration around the LE, the SITVM effectively dampens the excessive suction peak, yielding a highly realistic pressure distribution. This accurate capturing of the pressure plateau and the rapid closure of the laminar SB allows the SITVM to achieve an exceptional mean drag estimation compared to the DNS baseline.

Conversely, the DSM and ILES formulations underestimate the suction drop and damp the sharpness of the peak. This flattened pressure profile, combined with the drag penalty stemming from the elongated separation bubble, leads to a substantial over-prediction of total drag. Regarding global lift performance, all evaluated subgrid

formulations exhibit a general trend towards over-predicting the mean lift coefficient. The TVM displays the largest discrepancy, while the SITVM aligns closely with the DSM and ILES errors.

Subgrid stresses and energy transfer

The fundamental structural differences between linear eddy-viscosity models and non-linear tensorial formulations are mathematically isolated through the analysis of the resolved TKE ($\langle k \rangle$), resolved Reynolds shear stresses ($\langle u'v' \rangle$), and the subgrid stress tensor (τ_{ij}). The traditional DSM relies entirely on the linear Boussinesq hypothesis, which constrains the subgrid stress tensor to be perfectly aligned with the resolved strain rate tensor. Since the diagonal components of the strain rate tensor are negligible within a separated shear layer compared to the off-diagonal terms, the DSM fundamentally struggles to capture subgrid anisotropy. Apart from a localised activation within the transitional region of the shear layer near the LE, the model returns near-zero normal subgrid stresses across the rest of the fluid domain, severely underpredicting these components compared to the tensorial formulations.

By abandoning the Boussinesq alignment and directly coupling the filtered velocity gradients with the geometric inertia tensor of the grid cell, the TVM and SITVM natively resolve the anisotropic nature of the small-scale turbulent structures. As a consequence, both tensorial approaches predict highly intense diagonal subgrid stress components. These intense modelled stresses exert a robust damping effect on the resolved flow field within the intense shear regions. This mechanism explains why the normal and shear resolved Reynolds stresses are significantly reduced in the TVM and SITVM simulations compared to the highly energetic fluctuations observed in the DSM and ILES fields.

A key physical distinction between the two tensorial variants emerges on the laminar pressure side of the airfoil. While the baseline TVM unphysically activates in this zone due to the dominant mean shear, the SITVM correctly predicts near-zero normal and shear subgrid stresses along the attached laminar BL, validating the theoretical framework of the Shear-Improved correction. Furthermore, the analysis of the subgrid dissipation rates demonstrates that the non-linear formulations successfully reproduce local energy backscatter phenomena inside the transitional region of the separated shear layer. This reverse energy transfer, whereby energy flows from the subgrid scales back into the resolved mean and fluctuating fields, is a physical feature of transitional flows that strictly dissipative, purely scalar formulations like the DSM are mathematically incapable of reproducing.

5.2 Numerical stability

Alongside the evaluation of predictive accuracy, the present investigation highlighted significant challenges regarding the numerical stability of non-linear tensorial closures. The implementation of the TVM natively introduces a considerable degree of numerical fragility, primarily stemming from its inaccurate near-wall behaviour. As previously discussed, the TVM generates unphysical, non-zero subgrid stresses in the viscous sublayer. When coupled with the inherently low-dissipative nature of the model, these artificial stresses struggle to be numerically damped. This contrasts with traditional linear eddy-viscosity formulations, such as the DSM, whose higher inherent dissipation naturally suppresses such numerical fluctuations. While the TVM exhibits the highest degree of numerical fragility, it is crucial to specify that the SITVM does not entirely eliminate these stability challenges. By explicitly filtering out the mean shear contribution from the subgrid-scale computation, the Shear-Improved formulation mathematically mitigates the unphysical subgrid stresses near the wall. This correction renders the solver relatively more robust compared to the baseline TVM under identical numerical setups. However, the SITVM still retains the inherent sensitivity typical of non-linear models, confirming that numerical fragility remains a shared limitation of both tensorial approaches.

For both formulations, the primary trigger for these instabilities was identified in the geometric proportions of the near-wall cells, specifically the AR. Tensorial models rely heavily on the precise computation of the resolved velocity gradients. High AR cells at the solid boundary severely degrade the accuracy of the gradient-solving schemes. This numerical error directly propagates into the subgrid-scale modelling and is subsequently amplified by the non-linear formulation. Consequently, achieving a stable computation required an extensive iterative process of mesh generation and refinement. Numerous grids were tested to strictly optimise the near-wall AR while simultaneously maintaining an adequate y^+ resolution to resolve the viscous sublayer.

During the initial simulation attempts, the solver frequently exhibited divergent behaviour. Notably, this divergence did not occur during the early time steps of the transient flow initialisation, but rather at a mature stage of the simulation, specifically when the convected synthetic turbulent structures reached the airfoil and began interacting with the BL. This interaction triggered an unphysical, exponential growth of the near-wall velocity field, typically originating downstream of the mid-chord. To isolate the root cause, an extensive parametric study on the numerical schemes was conducted. Replacing the SMART advection scheme with a first-order Upwind formulation successfully prevented divergence; however, its inherently high numerical dissipation over-damped the resolved turbulent scales, rendering it strictly unsuitable for LES applications. Similarly, employing a first-order backward Euler time integration (the so-called Linear in T-Flows) proved more stable than the second-order Parabolic scheme. This increased stability stemmed directly from

numerical diffusion, which artificially damped both spurious oscillations and physical velocity fluctuations.

Further attempts were made to stabilise the computation by acting on the gradient discretisation, which directly feeds the non-linear subgrid models. The Least Squares method was tested but proved numerically unstable, even when increasing the number of iterations beyond the default value of four; consequently, the Gauss Theorem was selected, although attempts to prevent divergence by lowering the gradient calculation tolerances proved ineffective. Within the pressure-velocity coupling algorithm, increasing the number of PISO corrections from three to five merely increased the computational cost without mitigating the instability. This confirmed that the divergence was not driven by a lack of pressure-velocity convergence, but rather by the anomalous near-wall subgrid stresses generated by the tensorial formulation. Ultimately, while heavily lowering the under-relaxation factors was attempted to artificially suppress the fluctuations, it became evident that the root cause lay in the spatial discretisation. Only through the meticulous optimisation of the near-wall AR was it possible to truly stabilise the solver. The final structured mesh design successfully constrained the gradient calculations, allowing the simulations to run stably with the more accurate Parabolic and SMART schemes, while employing the default T-Flows under-relaxation factors.

It is also important to contextualise these findings within recent research on tensorial models. Previous studies on the NACA 4412 airfoil using the TVM [23] reported additional sources of numerical instability, such as large velocity gradients at the LE and abrupt variations in cell inertia distribution, particularly at the interfaces between structured and unstructured grid domains in the far wake. In the present work, however, these specific numerical issues were not encountered. The LE gradient instability was naturally mitigated by the exceptionally fine spatial resolution imposed in that region, which was inherently required to accommodate the resolution of the upstream synthetic eddy generation plane. Furthermore, the inertia-related stability issues observed in previous far-wake unstructured regions were entirely bypassed by employing a fully structured grid topology throughout the entire computational domain, which ensured smooth cell transitions and a consistent spatial distribution of the geometric inertia tensor.

5.3 Computational efficiency

In addition to predictive accuracy and numerical stability, the overall computational efficiency of the implemented subgrid-scale models was evaluated. A crucial practical aspect for the industrial viability of novel turbulence closures is their impact on simulation runtimes. It is noteworthy that the introduction of the non-linear tensorial formulations (TVM and SITVM) did not incur a perceptible increase in computational

cost compared to the traditional eddy-viscosity approach (DSM) or the implicit methodology (ILES).

Since all simulations were executed on identical hardware architectures and maintained consistent parallelisation strategies, a direct comparison of the execution times was possible. The analysis revealed that the marginal differences in the computational effort required by the distinct subgrid closures were entirely overshadowed by the physiological variability observed between adjacent run segments of the same simulation. Consequently, the present investigation provides no evidence to suggest that any specific modelling approach was notably faster or slower than the others. This confirms that, provided the near-wall grid constraints discussed in Section 5.2 are properly addressed, the tensorial models can be integrated without a detrimental impact on the overall computational efficiency of the solver.

5.4 Future improvements

Since the application of tensorial models to external aerodynamic flows is still in its early stages of development, the present investigation highlights several areas where further research is required. To address the inherent numerical instability of the tensorial approaches while preserving its superior predictive capabilities for complex flow topologies, two main developmental pathways are proposed for future works.

Mixed subgrid-scale model

The tensorial formulations could be employed as the non-linear component within a more comprehensive mixed subgrid-scale model. Generally, a mixed model combines a traditional eddy-viscosity model, which provides a scalar baseline for the subgrid stress tensor, with a non-linear model. In the context of the present research, a promising approach is the direct coupling of the DSM with the SITVM.

The proposed mixed model formulation would be expressed as:

$$\tau_{ij} = -2\nu_{sgs}\overline{S_{ij}} - \frac{\mathcal{I}'_{kh}}{2\Omega} \left[\left(\frac{\partial \overline{u_j}}{\partial x_h} - \frac{\partial \langle \overline{u_j} \rangle}{\partial x_h} \right) \frac{\partial \overline{u_i}}{\partial x_k} + \left(\frac{\partial \overline{u_i}}{\partial x_h} - \frac{\partial \langle \overline{u_i} \rangle}{\partial x_h} \right) \frac{\partial \overline{u_j}}{\partial x_k} \right] \quad (5.1)$$

where the first term represents the DSM contribution and the second term encapsulates the non-linear SITVM modelling.

This specific pairing exploits their complementary near-wall behaviours. While the SITVM, compared to the TVM, substantially mitigates the subgrid stress activity at the solid boundary by filtering out the mean shear, it does not completely suppress all the unphysical non-zero components. On the other hand, the dynamic procedure of the DSM reduces the eddy viscosity to zero in the viscous sublayer. By coupling these two approaches, the mixed model would maintain strict physical

consistency at the wall. In the separated and fully turbulent regions, the SITVM would accurately capture the subgrid anisotropy and backscatter phenomena, whilst the DSM component would introduce the strictly necessary background numerical dissipation to stabilise the non-linear tensorial fluctuations. This synergy could potentially bypass the need for extreme grid optimisation at the wall, rendering the solver significantly more robust.

Near-wall damping function

An alternative, computationally efficient solution to mitigate the unstable behaviour and the unphysical prediction of subgrid stresses at the wall is the implementation of mathematical damping functions. Specifically, the classic Van Driest damping function, traditionally defined as:

$$D(y^+) = 1 - e^{-\frac{y^+}{A^+}} \quad (5.2)$$

where y^+ is the dimensionless wall distance and A^+ is an empirical constant (conventionally $A^+ \approx 25$), could be applied to the subgrid stress components of both the baseline TVM and the SITVM. This formulation would artificially force the modelled stresses to zero within the viscous sublayer. By suppressing the spurious near-wall subgrid activity mathematically rather than geometrically, the solver would be relieved from the strict necessity of maintaining an extremely low near-wall AR. This would significantly ease the mesh generation process and relax the stringent structural constraints discussed in Section 5.2.

While this approach could significantly enhance near-wall numerical stability, it requires careful calibration. The standard Van Driest formulation is fundamentally based on the assumption of a Zero Pressure Gradient (ZPG). However, as extensively documented in Section 5.1, external aerodynamic applications are heavily governed by sustained APGs, which challenge the models and lead to a premature prediction of the PRR. Therefore, a highly relevant future improvement would involve the implementation and tuning of refined, APG-aware damping functions that explicitly incorporate local pressure gradient evaluations. Such an advancement could simultaneously address two of the primary limitations highlighted in this work: it would provide the necessary near-wall damping to guarantee numerical stability for both tensorial closures, while correctly adapting to the complex boundary layer dynamics to appropriately delay the premature turbulent detachment downstream.

Bibliography

- [1] Airfoil Tools. *NACA 4 digit airfoil generator*. <http://airfoiltools.com/airfoil/naca4digit>. 2026 (cit. on p. 15).
- [2] Alessandro Bottaro. ‘Note on open boundary conditions for elliptic flows’. In: *Numerical Heat Transfer, Part B Fundamentals* 18.2 (1990), pp. 243–256 (cit. on p. 20).
- [3] Michael Breuer. ‘Effect of inflow turbulence on an airfoil flow with laminar separation bubble: An LES study’. In: *Flow, Turbulence and Combustion* 101.2 (2018), pp. 433–456 (cit. on pp. 23, 24, 26, 31).
- [4] Harry W Carlson and Robert J Mack. ‘Studies of leading-edge thrust phenomena’. In: *Journal of Aircraft* 17.12 (1980), pp. 890–897 (cit. on p. 42).
- [5] A Cimarelli et al. ‘Resolved dynamics and subgrid stresses in separating and reattaching flows’. In: *Physics of Fluids* 31.9 (2019) (cit. on p. 56).
- [6] Andrea Cimarelli, Antonella Abbà and Massimo Germano. ‘General formalism for a reduced description and modelling of momentum and energy transfer in turbulence’. In: *Journal of Fluid Mechanics* 866 (2019), pp. 865–896 (cit. on p. 10).
- [7] Vlad Fluieraru. ‘Large Eddy Simulation of a turbulent channel flow using a tensorial eddy viscosity model’. M.Sc. Thesis. Motorvehicle University of Emilia-Romagna (MUNER) (cit. on p. 13).
- [8] Philip H. Gaskell and A. K. C. Lau. ‘Curvature-compensated convective transport: SMART, a new boundedness-preserving transport algorithm’. In: *International Journal for numerical methods in fluids* 8.6 (1988), pp. 617–641 (cit. on p. 22).
- [9] Massimo Germano. ‘A direct relation between the filtered subgrid stress and the second order structure function’. In: *Physics of Fluids* 19.3 (2007) (cit. on p. 9).
- [10] Massimo Germano. ‘Turbulence: the filtering approach’. In: *Journal of Fluid Mechanics* 238 (1992), pp. 325–336 (cit. on p. 49).

- [11] Massimo Germano et al. ‘A dynamic subgrid-scale eddy viscosity model’. In: *Physics of Fluids A: Fluid Dynamics* 3.7 (July 1991), pp. 1760–1765. ISSN: 0899-8213 (cit. on pp. 7, 8).
- [12] Christophe Geuzaine, Jean-Francois Remacle and P Dular. ‘Gmsh: a three-dimensional finite element mesh generator’. In: *International Journal for Numerical Methods in Engineering* 79.11 (2009), pp. 1309–1331 (cit. on p. 15).
- [13] Christopher Greenshields and Henry Weller. *Notes on Computational Fluid Dynamics: General Principles*. 1st. CFD Direct Limited, 2022. URL: <https://cfd.direct> (cit. on pp. 22, 23).
- [14] M. Hadziabdic and B. Niceno. *T-Flows*. <https://github.com/DelNov/T-Flows>. 2018 (cit. on p. 3).
- [15] M. Hadziabdic and B. Niceno. *T-Flows: The Open-Source Computational Fluid Dynamics Code on Arbitrary Computational Grids*. International University of Sarajevo, 2022 (cit. on p. 3).
- [16] Nicolas Jarrin et al. ‘A synthetic-eddy-method for generating inflow conditions for large-eddy simulations’. In: *International Journal of Heat and Fluid Flow* 27.4 (2006), pp. 585–593 (cit. on p. 24).
- [17] Jinhee Jeong and Fazle Hussain. ‘On the identification of a vortex’. In: *Journal of fluid mechanics* 285 (1995), pp. 69–94 (cit. on p. 32).
- [18] Kemal Koca et al. ‘Identification of flow phenomena over NACA 4412 wind turbine airfoil at low Reynolds numbers and role of laminar separation bubble on flow evolution’. In: *Energy* 144 (2018), pp. 750–764 (cit. on p. 43).
- [19] Athony Leonard. ‘Energy cascade in large-eddy simulations of turbulent fluid flows’. In: *Advances in geophysics*. Vol. 18. Elsevier, 1975, pp. 237–248 (cit. on p. 5).
- [20] Emmanuel Lévêque et al. ‘Shear-improved Smagorinsky model for large-eddy simulation of wall-bounded turbulent flows’. In: *Journal of Fluid Mechanics* 570 (2007), pp. 491–502 (cit. on p. 12).
- [21] D Lilly. ‘The representation of small-scale turbulence in numerical simulation experiments’. In: *NCAR Report* (1966) (cit. on p. 7).
- [22] Ivan Mary and Pierre Sagaut. ‘Large eddy simulation of flow around an airfoil near stall’. In: *AIAA journal* 40.6 (2002), pp. 1139–1145 (cit. on p. 3).
- [23] Filippo Moroni. ‘Large eddy simulation of a blunt body with sharp corners using a tensorial eddy viscosity model’. M.Sc. Thesis. Università degli Studi di Modena e Reggio Emilia (UNIMORE), 2023 (cit. on pp. 34, 44, 63).

- [24] Thomas J Mueller and Stephen M Batill. ‘Experimental studies of separation on a two-dimensional airfoil at low Reynolds numbers’. In: *AIAA journal* 20.4 (1982), pp. 457–463 (cit. on p. 34).
- [25] Ugo Piomelli and Jeffrey Robert Chasnov. ‘Large-eddy simulations: theory and applications’. In: *Turbulence and Transition Modelling*. Springer, 1996, pp. 269–336 (cit. on p. 18).
- [26] Stephen B Pope. ‘Turbulent flows’. In: *Measurement Science and Technology* 12.11 (2000), pp. 2020–2021 (cit. on p. 6).
- [27] Joseph Smagorinsky. ‘General circulation experiments with the primitive equations: I. The basic experiment’. In: *Monthly weather review* 91.3 (1963), pp. 99–164 (cit. on p. 6).
- [28] Yves Tessier Urrecha. ‘Large eddy simulation of a blunt body with sharp corners using a tensorial eddy viscosity model’. M.Sc. Thesis. Motorvehicle University of Emilia-Romagna (MUNER), 2022 (cit. on pp. 2, 32, 43).
- [29] Bert Vreman, Bernard Geurts and Hans Kuerten. ‘Realizability conditions for the turbulent stress tensor in large-eddy simulation’. In: *Journal of Fluid Mechanics* 278 (1994), pp. 351–362 (cit. on p. 49).
- [30] Manuel Zucchi. ‘Metodi numerici per la generazione di turbolenza in ingresso a simulazioni Large Eddy Simulation’. M.Sc. Thesis. Università degli Studi di Modena e Reggio Emilia (UNIMORE), 2026 (cit. on pp. 19, 24, 26).