A RANS-based CFD method for prediction of design and installation effects on automotive cooling fan noise

Master's thesis in Mobility Engineering

Sergio Moya Vinuesa

Department of Mechanics and Maritime Sciences

Chalmers University of Technology
Gothenburg, Sweden 2024
www.chalmers.se

Master's thesis 2024

A RANS-based CFD method for prediction of design and installation effects on automotive cooling fan noise

SERGIO MOYA VINUESA

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mechanics and Maritime Sciences
Division of Fluid Dynamics
Chalmers University of Technology
Gothenburg, Sweden 2024

A RANS-based CFD method for prediction of design and installation effects on automotive cooling fan noise

Supervisors: Asuka Gabriele Pietroniro and Zhongjie Huang, Volvo Car Corporation
Examiner: Niklas Andersson, Department of Mechanics and Maritime Sciences

Master's Thesis 2025
Department of Mechanics and Maritime Sciences
Division of Fluid Mechanics
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: CFD visualisation of various surface fields showing irradiated Curle acoustic power across the two studied fans for different rotational speeds.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2025

A RANS-based CFD method for prediction of design and installation effects on automotive cooling fan noise

Sergio Moya Vinuesa
Department of Mechanics and Maritime Sciences
Chalmers University of Technology

Abstract

Given the quietness of electrified powertrains, the concern about noise generated by axial cooling fans has gained importance in the automotive industry, especially in the premium sector. Besides, the need of substantial fan-driven airflow during standstill battery charging represents an added reason to develop silent cooling systems that meet the noise regulations for public and residential areas. This has caused an increasing interest in conducting further research on axial fan acoustics and in achieving more effective and efficient computational fluid dynamics (CFD) and computational aeroacoustics (CAA) analyses in the design process of this type of machines.

As part of the so-called *eFan2* project, this work aims at investigating the potential of steady-state Reynolds-averaged Navier-Stokes (RANS) simulations to yield profitable conclusions regarding the influence of the design and installation on the aeroacoustic noise generated by axial cooling fans. The moving reference frame (MRF) method was employed to recreate the effect of fan rotation on the computational fluid continuum. As the main core of the study, the impact of different mesh parameters and turbulence modelling on the total broadband acoustic power predictions was investigated on two fans with different geometry and architecture, for which validation data were produced during a dedicated experimental campaign.

A robust modelling methodology and valuable conclusions were achieved, highlighting the relationship between mesh resolution and predicted flow features and noise levels. Comparison with experimental data showed that the simulations reasonably well capture the trends of the total acoustic power over different operating conditions. Also, the simulation strategy for cases having greater tonal character is discussed.

Keywords: automotive axial cooling fans, aeroacoustics, aerodynamic performance, CFD, RANS, MRF, broadband acoustic source models, noise generation, turbulence models, mesh study, STAR-CCM+, ANSA.

Acknowledgements

This work would not have been possible without the invaluable supervision and guidance of Asuka Gabriele Pietroniro and Zhongjie Huang, whose expertise and support were instrumental throughout the project. I am also sincerely grateful to Magnus Knutsson for making this opportunity possible on behalf of Volvo Cars Corporation. The support and thoughtful input from Chalmers' examiner Niklas Andersson have been equally appreciated. From Volvo Trucks' industrial side, the realisation of the experimental test activities would not have been achievable without the collaboration of Sassan Etemad, Mikael Karlsson, and Peter Draksjö, as well as the assistance of Brian Fechner and Hamed Jamshidi from Volvo Cars. I express my sincere thanks to all of them.

To my family, and to my future self.

Sergio Moya Vinuesa, Gothenburg, 2025

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis, listed in alphabetical order:

3D	Three-dimensional
AMG	Algebraic multigrid
BEV	Battery electric vehicle
CAA	Computational aeroacoustics
CAD	Computer-aided design
CFD	Computational fluid dynamics
DES	Detached eddy simulation
DNS	Direct numerical simulation
EV	Electric vehicle
FDM	Finite difference method
FEM	Finite element method
FVM	Finite volume method
HPC	High-performance computing
HVAC	Heat, ventilation and air conditioning
IEA	International Energy Agency
IGES	Initial Graphics Exchange Specification
LES	Large eddy simulation
LBM	Lattice Boltzmann method
MRF	Moving reference frame
NVH	Noise, vibration and harshness
PL	Prism layer
RANS	Reynolds-averaged Navier-Stokes
RPM	Revolutions per minute
SDR	Specific dissipation rate
SST	Shear-stress transport
STEP	Standard for the Exchange of Product Data
STL	Stereolithography
SWL	Sound wattage level
TKE	Turbulent kinetic energy
VCC	Volvo Car Corporation (or Volvo Cars)
(V)GTT	(Volvo) Group Trucks Technology
WHO	World Health Organization

Contents

List of Figures

List of Tables

1

Introduction

With the aim of developing an immersive understanding of the project's context, this first chapter introduces key background on the topic at hand, offers a glimpse of relevant state-of-the-art research, and elaborates on the objective of the work.

1.1 Background

In recent times, the solid and fast expansion of electromobility has opened the door to new demands in the automotive industry. The use of a newer powertrain concept does not only open horizons regarding energy management or infrastructure, but also unleashes consequences in terms of user experience. For that matter, apart from evaluating drivability, car manufacturers often aim at guaranteeing the comfort of drivers and passengers, especially when it comes to premium-brand makers. Part of it is achieved by eliminating annoying noise in the cabin, which may be generated by diverse sources – from the powertrain and related structural vibrations, to the external moving air that flows around the car.

Now that for electrified vehicles the noise of combustion is partially or totally demoted, less noticeable acoustic sources are brought under the spotlight. Some of them have their origin in cooling and climate systems, which deal with the thermal control of various elements of the vehicle. Due to the fact that they generate noise sources stemming from both mechanical excitations and complex fluid flow turbulence, their acoustic analysis often demands thorough work to the field of NVH (noise, vibration and harshness), becoming one of the main focus of study in the continuous search of quietness and comfort.

However, it must be noted that although on-board comfort is indeed addressed, a big acoustic concern for electrified plug-in vehicles does not only occur while continuous driving, but also during battery recharge. In such a circumstance, the car will most likely be parked either in the very vicinity of a residential area – in essence, in the yard – or by the access zone of a frequented facility, often in fast charging mode. In this operational state, the batteries heat up considerably, and given that the car is standstill, there will be a lack of external cooling stream that must be supplemented by major fan-powered airflow. The high level of exterior noise disturbance produced by cooling fans in these conditions propitiate the scenario on which the present investigation is based.

According to the International Energy Agency (IEA), the regions where car electrification has advanced the most are the Nordic countries, China and some

Central European states. In many of these areas, such as Norway, Belgium and the Netherlands, terraced and detached houses predominate over flats as dwelling options, which makes the problem of the charging noise in residential spots somewhat more troubling. The three aforementioned countries registered in their 2023 sales records an electric-vehicle (EV) share of 93%, 41% and 35% from the total cars sold, respectively, compared to the European average of 22% [1]. Coincidentally, almost 80% of their populations live in housing types that potentially accommodate an adjacent charging point [2].

On top of that, the entrenched culture of welfare and environmental awareness in northern Europe does not overlook the harm that acoustic disruption can cause in outdoor areas. In fact, as related studies claim, the prolonged exposure to excessive noise pollution can lead to cardiovascular and neurological disorders, ultimately affecting quality of life and even reducing longevity [3]. Accordingly, the World Health Organization (WHO) proposes a standard for automotive use whose most constraining specification states that no living space should be exposed to more than 45 dB during night time. Meanwhile, low-pressure axial fans employed in cars can reach overall sound levels up to twice that threshold when working around their upper rotational speed bound [4], which demonstrates the need for taking action on the matter at hand.

Within this framework, the increasing emphasis on quality standards and product liability in the fast-growing EV market drives the need to tackle emerging challenges through extensive analytical understanding, in order to find optimal solutions that allow effective and robust control of the problem in the long run [5, 6, 7]. In line with this, a collaboration between Chalmers and industrial partners Volvo Cars and Volvo Group Trucks Technology (VGTT), together with the funding from the Swedish Energy Agency, led to the creation of the *eFan2* programme. This initiative forms a cluster of research projects aimed at conducting in-depth studies to expand knowledge on automotive cooling fans and to improve their performance and functionality. Good examples are Ghosh' study on the influence of inlet geometry on blade tip noise [8], Vourakis' in-lab acoustic measurements of varied installation cases [9], as well as their work on computational aeroacoustics (CAA) of fans installed in parallel [10].

As part of the programme, the present thesis work complements such inquiries by delving into fan aeroacoustics, with a focus on the use of CFD (computational fluid dynamics) to explore the impact of impeller geometry and installation on noise levels. Despite the existing studies around this topic, the potential of steady-state flow simulations for assessing fan design [11] still remains relatively untapped, making it a convenient field to cover in the current dissertation.

1.2 Project goal and scope

As outlined above, the particular objective of this project is to explore the viability of inexpensive steady-state CFD modelling for predicting variations in the levels of irradiated aeroacoustic noise across different low-pressure axial fan designs and installation layouts. This is to be done by setting up simple, yet robust, numerical models for two fan specimens that exhibit significant differences in blade geometry and stator placement, and subsequently validating the obtained simulation results through comparison with experimental data.

The pursuit of runtime economy drives the methodology to rely on certain prerequisites. At its core, the employment of Reynolds-averaged Navier-Stokes (RANS) formulation is the basis of the flow modelling. In conjunction with this, the moving reference frame (MRF) approach is established as the method to virtually model the fan rotation, given its proven cost-effectiveness. Furthermore, in alignment with the RANS-based strategy, SST (shear stress tensor) k-ω will be the principal turbulence model used in the study due to its strong near-wall workability, which is particularly advantageous for a flow scenario where the boundary layer has a major impact. Nonetheless, the standard k-ε model will also be utilised as a secondary option, offering supplementary data for possible comparison analyses concerning computational efficiency.

Given the steady-state nature of the RANS approach, the model is unable to recreate the temporal fluctuations of the pressure field that stem from the volume displaced by the blade rotation. Therefore, the scope of the acoustic assessment is limited to accounting for modelled broadband noise originated by turbulence. Likewise, distribution of acoustic pressure over frequency spectra will not be computed in the simulations, and, as required by the objective of the investigation, only overall sound levels will be calculated via source models that derive from the RANS formulation.

Building on these premises, the robustness of the models will be addressed through a systematic evaluation of key aspects of the numerical setup, including boundary conditions, physics modelling, convergence tracking, and, most notably, a dedicated mesh independence study. To adapt the size of the task to the time constraint and intent of the project, the meshing strategy is confined to assessing the prism layer construction and the refinement level in the wake region. That will constitute a significant part of the modelling methodology, and will also play a critical role in identifying the mesh configuration that offers the most computationally efficient numerical setup for each fan design.

Once the preferred simulation setup for each fan is found, it will be adjusted to evaluate a range of operating conditions. Nevertheless, to align with the workload and timeframe allocated to the project, the methodology will focus on optimising the model for conditions of zero inlet-to-outlet pressure rise, avoiding any blockage elements immediately upstream or downstream of the fan other than the shroud.

Allowing free intake and discharge of flow through the fan will generate neither front nor back pressure build-up in the reservoir regions adjacent to the test section, ensuring a neutral pressure gradient from the installation point of view. Therefore, the operational range simulations will comprise solely a sweep of rotational speeds, spanning from 800 to 3200 RPM. Although the simplified "laboratory" conditions described above do not fully replicate the complexities of a real installation scenario within a vehicle during usual operation, the analysis serves to consolidate the methodology and establish a consistent foundation for prospective, more realistic work.

The specified conditions will also be replicated on the test rig during the experimental validation, where both the measured noise levels and the flow performance parameters will be compared against simulation results. This comparison will quantify the model's error across various rotational speeds, highlighting the ranges where it performs with greater accuracy and those where further assessment may be necessary. Ideally, if a strong correlation is achieved, the presented modelling process could be used as a CFD methodology flexibly applicable to further fan designs and working conditions. That could facilitate high-quality yet cost-effective analyses, enabling efficient evaluation during the early stages of fan design. As a consequence, the overall agility of fan development, both in terms of performance and acoustics, would be enhanced.

2

Theory

This chapter presents the theoretical background upon which the study is based, covering key concepts in cooling fan aerodynamics, flow acoustics, and computational fluid dynamics. Together, these topics establish the foundation for the analysis and modelling approaches employed throughout the study.

2.1 Fluid mechanics of cooling fans

The generation of power in vehicles – whether by internal combustion engines or electric drive units – inevitably involves energy losses, primarily in the form of heat. For optimal system performance and to avoid overheating, this excess heat must be continuously rejected from the power-generating components. According to the first law of thermodynamics, this process relies on a basic principle: in order to dissipate heat, a colder medium must absorb it. In automotive systems, ambient air acts as this medium. The heat rejection is then quantified by the following balance equation:

$$\dot{Q} = \dot{m} \cdot c_p \cdot \Delta T \tag{2.1}$$

Here, increasing the air's mass flow rate $\dot{m}$ and the temperature change ΔT enhances heat rejection $\dot{Q}$ for a given specific heat c_p. The effectiveness of the system thus depends on heating up and evacuating as much airflow as possible. To accomplish this, the vehicle's cooling system uses a module formed by an axial fan and a radiator working in series. A coolant fluid circulates through internal ducts in the vehicle, absorbing heat from the power-generating components when circulating in close contact with them to transport that thermal energy to the cooling module. The heated coolant then flows into the radiator, where the heat is transferred to the surrounding air. The axial fan promotes this exchange by forcing air through the radiator core, enhancing the rate of heat transfer. This cooling mechanism operates based on convection, where the movement of air in contact with the radiator's surface carries thermal energy away. By maintaining a steady airflow, the fan ensures that heat is continuously removed and replaced with cooler ambient air, keeping component temperatures within operational limits.

The efficiency of this cooling setup depends heavily on the fan's aerodynamic performance, which must be optimised through thorough geometric design and awareness of installation elements. A well-designed fan must balance airflow generation, energy consumption, and acoustic emissions. Therefore, analysing the fluid dynamic behaviour of the fan becomes essential for both cooling effectiveness and

overall vehicle performance.

Fans operate by imparting momentum to the surrounding air, producing a pressure difference that drives airflow across the heat exchanger. The process entails complex three-dimensional phenomena that must be analysed using the principles of fluid mechanics. This involves a tight relationship between the magnitudes describing the dynamic state of the continuum and the intrinsic properties of the fluid, such as viscosity, density, and thermal conductivity, which stem from its molecular and intermolecular characteristics.

At the core of this analysis are the Navier-Stokes equations, which describe how fluid motion and properties evolve in space and time. In their differential form for compressible constant-viscosity flows, these governing equations are written as:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \tag{2.2}$$

$$\frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) = -\nabla p + \nabla \cdot \tau + \rho \mathbf{f} \tag{2.3}$$

$$\frac{\partial (\rho E)}{\partial t} + \nabla \cdot [\mathbf{u}(\rho E + p)] = \nabla \cdot (k \nabla T) + \Phi \tag{2.4}$$

In these equations, ρ denotes the fluid density, $\mathbf{u}$ is the velocity vector, p represents the static pressure, the tensor τ stands for the shear stresses, and $\mathbf{f}$ encompasses body forces such as gravity. E is the total energy per unit mass, thermal conductivity is given by k, T is the temperature, and the symbol Φ accounts for viscous dissipation. The differential character of the equations is appreciated through the first terms of the left-hand sides, which express partial derivatives with respect to time t, and through the nabla operator ∇, which introduces gradients and divergences with respect to the available space directions. Importantly, Cauchy's shear stress τ depends in turn on the fluid's viscosity μ and on the velocity gradients, as represented in simplified form in Eq. 2.6.

These equations respectively describe the conservation of mass, momentum, and energy. The continuity equation (Eq. 2.2), with a vectorial form, ensures that the change of mass within a control volume is balanced with the overall inflow or outflow of fluid. The momentum equation (Eq. 2.3), also vectorial, applies Newton's second law to a fluid element, linking the rate of change of momentum to the forces acting on it, including pressure gradients, viscous stresses, and body forces. The second term on its left-hand side, known as the convective term, provides the Navier-Stokes equation set with a non-linear relation, making their analytical solution extremely challenging in most practical scenarios. The energy equation (Eq. 2.4), of scalar nature, tracks the total energy in the system, including kinetic and internal energy, and incorporates the effects of heat conduction and viscous dissipation.

In applications involving gases, where flows are typically compressible, the system of equations is closed with an additional thermodynamic equation of state

that relates pressure, density, and temperature. In most cases, this is achieved through the ideal gas law:

$$p = \rho R T \tag{2.5}$$

In this equation, which applies for common monoatomic and diatomic gases under moderate pressure and temperature conditions, the ideal gas constant R has a value of 287 J/(kg · K) and provides a linear relation between the state variables. However, in many engineering applications – particularly those involving nearly isothermal or incompressible flows – thermal effects may be negligible, allowing the energy equation, and subsequently the equation of state, to be excluded from the analysis for simplification.

In any case, the presented principles form a set of equations with a high degree of mathematical complexity. Notably, due to the aforementioned non-linearities of the governing physics, the fluid motion becomes unsteady and chaotic – i.e. turbulent – in most flows, especially in those subject to highly agitated conditions as often encountered in rotating machinery. The interaction between moving components and the surrounding medium induces a large spectrum of turbulent phenomena that strongly influence the aerodynamic performance and acoustic characteristics of the system.

One of the most relevant turbulent features, especially in the generation of flow-induced noise and in the development of other unsteady aerodynamic phenomena, is the boundary layer. This is a thin flow region that forms near solid surfaces, where the wall-tangent velocity components of the fluid increase rapidly from zero at the wall – due to the no-slip condition – to the freestream velocity of the outer flow. The steepness of this velocity gradient depends on the fluid's viscosity and on the inertial characteristics of the flow, and it gives rise to a shear stress between the wall and the freestream. Thus, the following general relation for two-dimensional flow is involved, establishing the definition of viscosity as the relation between the arisen velocity gradient and the generated shear stress:

$$\tau = \mu \frac{du}{dy} \tag{2.6}$$

If this equation is applied to near-wall flow, τ is the shear stress, u is the wall-tangent velocity, and y is the wall-normal coordinate. This shear stress plays a critical role in the momentum exchange between adjacent fluid layers and strongly influences both the development and stability of the boundary layer. In this sense, the layer can form as either an orderly laminar regime, characterised by parallel streamlines with negligible transverse mixing, or develop as a swirling turbulent regime, where chaotic velocity fluctuations in the form of eddies dominate. This will depend on the influence of the inertial state of the flow over its viscous character. Such balance is quantified by the Reynolds number Re, a dimensionless parameter defined as the ratio of inertial to viscous forces:

$$Re = \frac{\rho U L}{\mu} \tag{2.7}$$

U denotes the characteristic velocity of the flow and L represents a characteristic length, associated with the size scale of the flow scenario. A higher Reynolds number indicates that inertial effects outweigh viscous damping, which promotes instability and triggers turbulence. The boundary layer typically grows along a surface as the fluid moves downstream, initially forming as laminar with a low local Reynolds number, and potentially transitioning to turbulent as the value increases. Thus, the Reynolds number serves as a critical indicator for predicting the nature and evolution of the boundary layer. Although there are no universal threshold values to determine the flow regime, these are general guidelines for turbomachinery applications:

- Laminar flow: $Re \lesssim 2 \cdot 10^5$
- Transitional flow: $2 \cdot 10^5 \lesssim Re \lesssim 5 \cdot 10^5$
- Turbulent flow: $Re \gtrsim 5 \cdot 10^5$

Turbulent boundary layers can be divided into distinct sublayers, each governed by different dominant forces. Closest to the wall lies the viscous sublayer, where viscous stresses are predominant and the velocity profile is approximately linear. Above this region is the buffer layer, a transitional zone where both viscous and turbulent effects are significant. Further away from the wall is the *log-law* layer, where turbulent mixing becomes dominant and the velocity profile follows a logarithmic distribution. Beyond this group of sublayers known as the inner region, an outer region can be additionally identified, in which the effects of the wall diminish and the flow matches the external freestream very gradually.

In essence, the velocity profile in turbulent boundary layers tends to follow a similar shape across a wide range of flows, mostly only varying the overall scale of the magnitudes. Subsequently, the structure and behaviour of the boundary layer can be generalised regardless of fluid properties, geometry, or flow scale, allowing us to express this profile in a universal form. This is done through the use of normalised variables, which make it possible to compare flows with different scales or fluid properties by collapsing their near-wall profiles onto a common framework. The most important of these variables are the dimensionless wall-tangent velocity u^+ and the dimensionless wall distance y^+. They are defined as a function of another variable frequently used in boundary layer theory, the friction velocity u_τ:

$$u^+ = \frac{u}{u_\tau} \tag{2.8}$$

$$y^+ = \frac{\rho u_\tau y}{\mu} \tag{2.9}$$

The friction velocity is defined in Eq. 2.10, which implies simultaneously the use of Eq. 2.11, a semi-empirical formula for the estimation of the wall shear stress τ_w caused by the boundary layer. To enable physical correlation, the latter includes the wall friction coefficient c_f, which retains a direct proportionality to the Reynolds number that varies across different flow scenarios.

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}} \tag{2.10}$$

$$\tau_w = \frac{\rho U^2 c_f}{2} \tag{2.11}$$

With this, the velocity profile across the boundary layer can be described by means of u^+ and y^+, and universal wall-distance limits defining the three sublayers can be identified:

- Viscous sublayer: $y^+ < 5$. Since the viscous effects dominate and the velocity increases linearly with wall distance, the profile in this region exhibits the following law:

$$u^+ = y^+ \tag{2.12}$$

- Buffer layer: $5 < y^+ < 30$. With both viscous and turbulent effects becoming significant, no simple analytical expression fully captures the behaviour, and particular empirical correlations are often used depending on the flow case.
- Log-law layer: $y^+ > 30$. The turbulence-dominant regime produces a logarithmic velocity profile described in Eq. 2.13 [12], where the von Kármán constant κ has a value of 0.41 and C^+ is an empirical case-dependent constant:

$$u^+ = \frac{1}{\kappa} \ln y^+ + C^+ \tag{2.13}$$

As seen in Section 2.2.3, this framework is particularly useful in CFD wall treatment and turbulence modelling, where accurate prediction of wall-bounded flow behaviour is critical. As it has been said, turbomachinery, external aerodynamics, and internal cooling systems are indeed applications where near-wall phenomena significantly influence global performance. In these cases, the boundary layer development can be additionally involved in the generation of further turbulent structures, which highlights the importance of a good near-wall flow prediction. Two good examples of these can be flow separation and blade tip vortices. The former occurs when the boundary layer fails to remain attached to the surface, typically due to adverse pressure gradients or abrupt changes in geometry, often happening near the trailing edges or blade tips. The separated flow forms recirculating regions that contribute to the outbreak of unstable and unpredictable patterns. On the other hand, at the blade tips, the pressure difference between the suction and pressure sides of the blade causes fluid to roll around the tip, forming tip vortices that likewise result detrimental for the overall machine performance.

Together, these features give rise to wake formation downstream of the impeller, recognisable as a large region of non-uniform flow field and characterised by low mean velocity and elevated turbulence levels. As these structures are convected through the domain, they impinge upon solid components as well as interact with each other, generating further disturbances and contributing to the inherently unsteady and complex nature of the flow. This challenge highlights the importance of thorough aerodynamic design to manage fluid flows and unleashes the need for CFD and turbulence modelling techniques to predict them in view of advanced analysis.

2.2 Computational fluid dynamics

Computational fluid dynamics refers to the numerical study of fluid flow phenomena by solving the governing equations, i.e. the Navier-Stokes equations, over a spatial domain through computational methods. Several numerical methods exist for solving fluid dynamics problems in CFD, including the finite difference method (FDM), the finite element method (FEM), and the lattice Boltzmann method (LBM). However, the most commonly used approach in industrial and engineering applications is the finite volume method (FVM). The FVM discretises the spatial domain into small control volumes and applies the conservation equations in integral form. This ensures the conservation of physical quantities across each volume, making it particularly suitable and robust for simulating fluid flows with complex geometries and boundary conditions. Therefore, with FVM as the method employed in this project, the current section provides insight into the technical foundations of this modelling approach.

2.2.1 Fluid flow modelling

To properly describe the physics of the fluid flow, models must be selected that accurately represent the key physical processes involved in the system. Thus, in the CFD software environment, the selectable layers of continuum physics models include:

- Spatial dimensionality: the problem can be defined in one-dimensional, two-dimensional, or three-dimensional form depending on the geometry, symmetry, and required resolution of the physical domain.
- Time resolution: flow simulations may be steady, where the output is forced to target time-invariant conditions, or unsteady, where it evolves temporally. Within unsteady simulations, different numerical schemes such as explicit, implicit, or adaptive time stepping can be specified to resolve the time-dependent behaviour. Since this investigation focuses on evaluating a steady-state approach, time dependency is excluded, and the temporal terms in the governing equations are consequently omitted.
- Continuum type and thermodynamic assumptions: the matter state of the physical medium, such as solid, liquid, gas, or diverse multiphase systems, is to be selected. The corresponding thermodynamic behaviour is further specified through the choice of an equation of state – common formulations include constant density for incompressible flows, the ideal gas law for moderate compressible conditions, real gas equations incorporating empirical fits or tabulated data, complex polynomial or segmented expressions, and custom or user-defined state equations tailored to specific needs. This also includes defining the values of the fluid properties such as thermal conductivity, specific heat, and density, whether constant or dependent on fluid behaviour.
- Viscous regime: the flow regime can be set as inviscid, laminar, or turbulent. As with thermodynamic properties, viscosity must be specified – typically as a constant for Newtonian fluids or through a functional form in the case of

non-Newtonian behaviour.
- Turbulence modelling approach: options include Reynolds-averaged Navier-Stokes (RANS), large eddy simulation (LES), detached eddy simulation (DES), or hybrid methods. Since this work is based on the use of the former, the governing equations rely on Reynolds decomposition to model the turbulent fluctuations within the mean flow field. Moreover, the concrete turbulence models and specific formulation is also set up. This will be addressed in greater depth in Section 2.2.3.

2.2.2 Fan rotation modelling

In simulations of turbomachinery, accurately capturing the effects of rotors requires appropriate modelling of their motion. The interaction between rotating and stationary parts influences the flow field and must be taken into account. The main methods employed for simulating rotating machinery in CFD are the moving reference frame (MRF) approach and the sliding mesh method.

The sliding mesh technique is used in transient simulations, where the mesh physically rotates, re-adapting to the rest of the domain after each timestep. Therefore, it delivers physical, time-marching flow solutions and captures transient phenomena such as unsteady wakes or blade-passing pressure pulsations. In contrast, the MRF method offers a computationally efficient alternative for steady-state simulations. This is particularly useful during the early design stages or when the interest lies in overall performance characteristics such as pressure rise, head, torque, or overall broadband noise level, rather than transient flow behaviour.

In the MRF approach, the mesh itself remains stationary, although it is also divided into two regions: the rotating zone – immediately close to the impeller – and the stationary zone. In this case, each zone has its own reference frame, and interfaces are defined to connect them. These interfaces assume a steady relative position between the rotating and stationary components. The impeller's region is referenced to a rotating coordinate system, whereas the rest of the domain obeys the global static reference frame. While the momentum equations from the stationary region remain in a regular directional form, those defining the fluid physics of the rotating region are modified to account for the rotational effects, allowing the flow field to be solved as steady-state from the perspective of the rotating frame. This modified version of the Navier-Stokes equations includes additional source terms, namely the Coriolis force and the centrifugal force, which must be considered to represent the rotating system dynamics. Thus, Cauchy's vectorial momentum equation in the rotating reference frame is expressed as:

$$\frac{\partial(\rho\vec{v}_r)}{\partial t} + \nabla \cdot (\rho\vec{v}_r\vec{v}_r) = -\nabla p + \nabla \cdot \vec{\tau} + \rho\vec{g} - \rho\vec{\Omega} \times (\vec{\Omega} \times \vec{r}) - 2\rho\vec{\Omega} \times \vec{v}_r \qquad (2.14)$$

Here, $\vec{v}_r$ represents the velocity relative to the rotating frame, $\vec{\Omega}$ is the angular velocity vector, and $\vec{r}$ is the position vector from the axis of rotation. The last

two terms on the right-hand side of the equation denote the centrifugal and Coriolis forces, respectively. These forces virtually act on the fluid particles due to the non-inertial nature of the rotating frame and are essential for capturing the correct pressure and velocity fields within the impeller region.

2.2.3 Turbulence modelling

Turbulence is a complex, chaotic state of fluid motion characterized by vortical structures, fluctuating velocity fields, and enhanced mixing. It spans a wide range of time and length scales, making its direct numerical simulation (DNS) computationally demanding, especially for high Reynolds number flows encountered in practical applications. As a result, turbulence modelling becomes essential to approximate its effects on the mean flow in a tractable manner.

The most used approach when modelling turbulence in a steady-state environment is the Reynolds-Averaged Navier–Stokes (RANS) formulation. This framework involves decomposing the instantaneous flow variables into mean and fluctuating components, a process known as Reynolds decomposition. For velocity and pressure, this is expressed as:

$$u_i = \overline{u_i} + u_i' \tag{2.15}$$

$$p = \overline{p} + p' \tag{2.16}$$

Substituting these decomposed variables into the Navier-Stokes equations and averaging them over expressed leads to the RANS equations. As a sample, the RANS moment equation in tensor notation is presented as follows:

$$\frac{\partial \overline{u_i}}{\partial t} + \overline{u_j}\frac{\partial \overline{u_i}}{\partial x_j} = -\frac{1}{\rho}\frac{\partial \overline{p}}{\partial x_i} + \nu\frac{\partial^2 \overline{u_i}}{\partial x_j^2} - \frac{\partial \overline{u_i' u_j'}}{\partial x_j} + f_i \tag{2.17}$$

As it can be seen, this formulation includes an additional term known as the Reynolds stress $\overline{u_i' u_j'}$, which represents the momentum transfer due to turbulent fluctuations and requires modelling to close the system. The Boussinesq hypothesis provides a closure by relating the Reynolds stresses to the mean velocity gradients through the eddy viscosity ν_t, introduced as a new complementary variable analogous to the molecular viscosity. The Boussinesq hypothesis is then defined as:

$$-\overline{u_i' u_j'} = \nu_t \left(\frac{\partial \overline{u_i}}{\partial x_j} + \frac{\partial \overline{u_j}}{\partial x_i} \right) - \frac{2}{3} k \delta_{ij} \tag{2.18}$$

Here, $k = \frac{1}{2}\overline{u_i' u_i'}$ is the turbulent kinetic energy. The core challenge now shifts to modelling ν_t, for which various eddy-viscosity models have been developed. Among this family, the k-ε model and the SST k-ω model are widely utilised in industrial fluid dynamics simulations, as they balance computational efficiency with the ability to capture key turbulent effects. In alignment with this, they are also the used turbulence models in the present project.

2.2.3.1 The standard k-ε model

The standard k-ε model, introduced by Launder and Spalding [13], is a two-equation eddy-viscosity model that introduces transport equations for the turbulent kinetic energy k and its rate of dissipation ε:

$$\frac{\partial k}{\partial t} + \overline{u_j}\frac{\partial k}{\partial x_j} = P_k - \varepsilon + \frac{\partial}{\partial x_j}\left[\left(\nu + \frac{\nu_t}{\sigma_k}\right)\frac{\partial k}{\partial x_j}\right] \tag{2.19}$$

$$\frac{\partial \varepsilon}{\partial t} + \overline{u_j}\frac{\partial \varepsilon}{\partial x_j} = C_{\varepsilon_1}\frac{\varepsilon}{k}P_k - C_{\varepsilon_2}\frac{\varepsilon^2}{k} + \frac{\partial}{\partial x_j}\left[\left(\nu + \frac{\nu_t}{\sigma_\varepsilon}\right)\frac{\partial \varepsilon}{\partial x_j}\right] \tag{2.20}$$

The production term P_k is typically given by:

$$P_k = \nu_t\left(\frac{\partial \overline{u_i}}{\partial x_j} + \frac{\partial \overline{u_j}}{\partial x_i}\right)\frac{\partial \overline{u_i}}{\partial x_j} \tag{2.21}$$

The eddy viscosity is then computed as:

$$\nu_t = C_\mu \frac{k^2}{\varepsilon} \tag{2.22}$$

Moreover, in this formulation we find newly introduced empirical constants: σ_k, σ_ε, $C_{\varepsilon 1}$, $C_{\varepsilon 2}$, and C_μ. Their values might slightly change depending on the flow case.

2.2.3.2 The SST k-ω model

Menter's shear stress transport (SST) k-ω model combines the k-ω formulation near the wall with the k-ε behaviour in the free stream through a blending function [14]. This hybrid approach improves performance in adverse pressure gradient flows and boundary layer separation. The governing transport equations are:

$$\frac{\partial k}{\partial t} + \overline{u_j}\frac{\partial k}{\partial x_j} = P_k - \beta^* k\omega + \frac{\partial}{\partial x_j}\left[(\nu + \sigma_k \nu_t)\frac{\partial k}{\partial x_j}\right] \tag{2.23}$$

$$\frac{\partial}{\partial x_j}(\bar{v}_j\omega) = \frac{\partial}{\partial x_j}\left[\left(\nu + \frac{\nu_t}{\sigma_\varepsilon}\right)\frac{\partial \omega}{\partial x_j}\right] + \alpha\frac{\omega}{k}P^k - \beta\omega^2 + \frac{2}{k}\left(\nu + \frac{\nu_t}{\sigma_\varepsilon}\right)\frac{\partial k}{\partial x_i}\frac{\partial \omega}{\partial x_i} \tag{2.24}$$

The turbulent viscosity for SST is defined as:

$$\nu_t = \frac{k}{\omega} \tag{2.25}$$

The newly introduced coefficients are β^*, β and α. Their values are determined according to the blending functions, which in turn are defined by the conditional treatment of further coefficients and enable adaptation either to the k-ε behaviour or to the k-ω operation depending on the wall distance.

2.2.4 Noise modelling

The field of acoustics studies how mechanical vibrations propagate as waves through physical media. In the context of fluids, it concerns primarily the propagation of pressure disturbances through a medium such as air. This work focuses specifically on flow-induced noise, whose discipline is known as flow acoustics. In particular, it addresses the noise generated by the interaction of a fluid stream with solid elements – in this case, a rotating fan and its spokes.

When the periodic vibration pattern of a source, or of a response within the system, is harmonic, it can be mathematically decomposed into a sum of sinusoidal components ordered by frequency. This process, governed by the Fourier transform, allows complex acoustic signals to be analysed by identifying the contributing frequencies, their amplitudes, and intensities. Through this, the sources of various origins or nature can be identified and quantified within the overall sound field.

In fan-related aeroacoustic phenomena, the frequency content of the emitted sound is a key aspect for characterising its nature and origin. Acoustic signals can generally be classified as tonal, narrowband, or broadband, depending on how their energy is distributed across the frequency spectrum. Tonal components are marked by sharp, well-defined frequencies and typically arise from periodic sources, such as the blade-passing noise generated by the regular displacement of fluid by rotating blades. This source is inherently linked to the rotational frequency and its harmonics, and can only be resolved through time-dependent simulation approaches, as it is not present in steady-state fields. Narrowband noise, in contrast, arises from the interaction of coherent large-scale turbulent structures or vortex shedding mechanisms, leading to concentrated acoustic energy in relatively limited spectral bands; this phenomenon also requires unsteady modelling to be captured. Lastly, broadband noise covers a continuous range of frequencies and is associated with small-scale turbulent fluctuations and near-wall pressure variations. While more difficult to resolve in detail, such broadband content can be estimated through empirical or semi-empirical source models applied to steady-state flow solutions, such as those obtained from RANS simulations.

From a parallel perspective, Lighthill's work [15] proposes sorting the sound sources into three idealised categories based on their mode of generation and the nature of their acoustic radiation:

- Monopole sources: associated with time-varying mass injection or extraction, such as that of a pulsating sphere.
- Dipole sources: related to unsteady forces on solid surfaces, such as eddies impinging on a blade.
- Quadrupole sources: related to turbulence-induced pressure fluctuations in the bulk of the fluid, such as those caused by the isotropic eddies in a wake.

The prediction of fan noise in this RANS-based work will then focus on the contribution of broadband content, generated by dipole and quadrupole sources. As said, unlike tonal and narrowband components, which are linked to coherent,

periodic flow structures and require unsteady simulations, broadband noise can be estimated from steady-state CFD data. These estimations are possible thanks to the so-called broadband noise source models, which stem from Lighthill's acoustic analogy. The acoustic analogies interpret the turbulent flow field as an equivalent source term in the theoretical acoustic wave equation. Among the classical models derived from Lighthill's analogy, two of the most widely used are the Curle and Proudman models, suitable for dipole- and quadrupole-dominated broadband noise, respectively.

2.2.4.1 The Curle source model

The Curle model [16] applies Lighthill's theory including rigid solid walls, capturing dipole-type noise generated by unsteady surface forces, such as pressure fluctuations on fan blades. The total acoustic power P_S radiated from a solid surface S can be estimated as:

$$P_S = \int_S \frac{A(\vec{y})}{12\pi\rho_0 c_0^3} \overline{\left(\frac{\partial p}{\partial t}\right)^2} dS \tag{2.26}$$

$A(\vec{y})$ is the correlation area and c_0 is the reference speed of sound. In steady-state CFD simulations, direct time derivatives are not available; hence, complementary formulations based on turbulence quantities may be used to estimate the pressure derivative, such as relating pressure fluctuations to turbulent kinetic energy and wall shear.

2.2.4.2 The Proudman source model

The Proudman model [17] estimates the acoustic power per unit volume generated by isotropic, homogeneous turbulence in free field, under the assumptions of low Mach number and no mean flow. It is particularly appropriate for modelling quadrupole noise from volume turbulence in the absence of solid boundaries. The expression for the acoustic power per unit volume P_V is:

$$P_V = \int_V \alpha_c \rho_0 \frac{U^3}{L} \frac{U^5}{c_0^5} dV \tag{2.27}$$

U and L are the turbulence scales for velocity and length, respectively, and α_c is equal to 0.614. This relation provides an estimate of the volume-integrated acoustic power when applied over the flow domain where turbulence is present.

3

Methodology

The creation of CFD models involves a procedure that starts with the generation of the domain geometry, continues with its spatial discretisation, and progresses with the configuration of the physical properties of the fluid problem and the specification of the numerical settings that drive its solution. The present chapter elaborates on this process, dividing it into key sections: preparation of the geometry, modelling of boundary conditions, computational meshing, selection of physics models, and configuration of solver and convergence settings.

Beyond that, due to the complexity and delicacy of these models, their development often implies presenting a comprehensive methodology consisting of well-founded practices or routines for the determination of a functioning and reliable numerical setup, such as a thorough mesh parametrisation or the definition of robust convergence criteria. This enhances the applicability of the modelling criteria to differing simulation requirements and conditions, in contrast to the idea of rigidly establishing an arbitrary array of computational settings whose extrapolation to other fan model scenarios is not contemplated. Therefore, as outlined in Section 1.2, the current chapter will reflect this aim as well as the intention of building an adaptable simulation setup in view of prospective purposes or users. This is to be accomplished by means of well-thought-out geometry features to efficiently accommodate possible model updates and related modifications involving the numerical settings.

Following the development of the methodology and the definition of an appropriate simulation setup for each fan specimen, i.e. *Fan 1* and *Fan 2*, a representative range of operating speeds (800–3200 RPM) is simulated. The results were subjected to validation to assess the accuracy of the methodology. Comprehensive measurements covering the entire speed range were obtained from an extensive experimental campaign carried out at VGTT's fan test facilities.

The CFD model was built through Siemens' simulation software *Simcenter STAR-CCM+* at its version 2210. The simulations were all run on VCC's HPC cluster and with the double-precision mode set, to ensure capturing significant number of decimals from the tiny values of sound power that small changes in pressure generate. To perform any preliminary work on the 3D CAD model, *ANSA* pre-processor by BETA CAE Systems was used.

3.1 Geometry

The aim of this section is to describe the geometries of the fans and the domain included in the simulation environment, as well as to present the process through which the geometry model has been set up. Hence, the adaptation of the CAD model and its incorporation into the CFD scheme will be showcased.

3.1.1 Domain geometric definition

First of all, the global frame of reference must be made clear to settle geometrical and physical consistency through the whole methodology. The CAD model and the CFD setup were arranged so that the x-axis corresponded to the axial direction defined by the fan, and whose positive orientation is that of the flow running across the test section in the intended direction. The y-axis involves the transversal direction of the test rig (visit Figure 3.6), with its positive direction pointing to the left-hand side looking from the upstream face of the fan. The z-axis corresponds to the gravitational direction, whose positive side points upwards. According to the positioning of the leading and trailing edges, *Fan 1* is designed to spin counterclockwise if seen from the front, i.e. in the negative x-direction; and *Fan 2* is intended to do so clockwise, i.e. in the positive x-direction. This distinction can be noticed in Figure 3.1 if the difference in the orientation of the reference axes is observed:

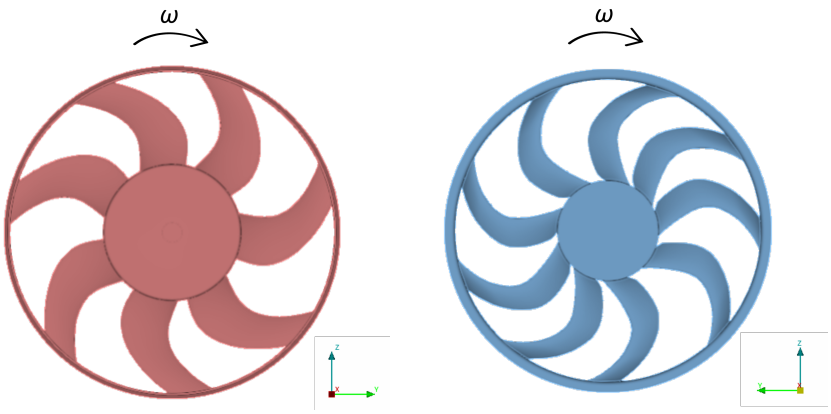

Figure 3.1: Geometry comparison of *Fan 1* (a) and *Fan 2* (b) impellers.

Design-wise, the significant differences between the two fans are that *Fan 1* is slightly larger and has 7 forward-swept blades, whereas *Fan 2*, somewhat smaller, has 9 backward-swept blades and shorter chord. The reason for *Fan 1* being bigger and having larger chord is that the ICE (internal combustion engine) vehicle that it was designed for requires more cooling than the BEV (battery electric vehicle) platform that *Fan 2* operates for. Besides, the spokes that support the hub are located upstream of the impeller in the former, contrary to those of the latter. As can be seen in Figure 3.1, the layout of both impellers is characterised by possessing a simple cylinder-like hub and having all the blade tips connected to a big outer ring

that keeps them structurally robust. A summary of the geometric data for the two fans is presented in Table 3.1, with some details in relative terms for confidentiality reasons:

	Fan 1	***Fan 2***
No. of blades	7	9
Tip radius	x	$0.94x$
Hub radius	y	$0.73y$
Mean chord length	z	$0.69z$
Blade sweep direction	Backwards	Forwards
Spokes location (w.r.t. impeller)	Upstream	Downstream

Table 3.1: Relevant geometric details concerning *Fan 1* and *Fan 2* designs.

According to what is demonstrated in a study by Park, Lee, and Lee, forward sweep can be used to obtain more silent aeroacoustics, while backward-sweep blades provide higher flow rates at the expense of greater noise generation [18]. On the other hand, it could be anticipated that the downstream position of *Fan 2*'s spokes with respect to the impeller might create an unfavourable acoustic drawback in relation to *Fan 1* layout, due to their interference with the wake. However, this could be mitigated thanks to their more streamlined design compared to those of *Fan 1*. With all this information, one can expect *Fan 1* to exhibit characteristics of better cooling performance, but *Fan 2* to provide less noisy operation.

Regarding the fan surroundings, Volvo GTT provided the CAD geometry file of their fan test rig facility, where the actual validation of the model would be done. This is basically two plenum chambers separated by the wall on which the test section is located (an illustration can be visited in Figure 3.6). To determine the extent of test rig to be included in the computational domain, it is relevant to consider that the location of the model boundaries will strongly influence solution feasibility. Since the inlet and outlet boundary conditions in the solver are typically uniform along their associated surface part, it is important to locate these boundaries sufficiently far from the fan vicinity, where predictable and barely disturbed flow is found. On the other hand, reducing the size of the domain moderately is something that can save computational effort by avoiding a significant number of mesh cells. According to this, it might have been unnecessary to include the chambers entirely for the current case, since they are very large and undisturbed flow could be expected at a shorter distance than the total length of the room. Nevertheless, although some grid cells could have been saved, the entire test chambers were finally conserved following the goal of versatility: if further unsteady-flow analysis was desired using the present setup as a baseline, the actual full size could be needed to reliably capture wave propagation phenomena.

3.1.2 CAD model

The full-scale CAD models of the two fans and the experimental test rig were provided in ANSA file format (*.ansa*). In the case of the former, their initial content corresponded to the exact design geometries. Thus, these contained some small sharp details whose effect on the numerical results would not be worth capturing, since it is expected to be negligible and would add unnecessary complexity and computational cost to the simulations. Because of this, some geometry simplification was carried out, and as it can be seen in Figures 3.2 and 3.3, it mostly implied the elimination of structural ribs and air bypass openings from the fan shroud, and of the closure of the motor housing in the hub. Note that the rotor will therefore not be in contact with any part, but just "floating" close to the hub. To conduct all this clean-up, some convenient features from ANSA's section *TOPO* were utilised, like surface deletion, face cutting, vertex projection and edge merging. The simplified fan model files would eventually consist of two distinguishable surface parts: the shroud – including the motor hub and the spokes – and the impeller.

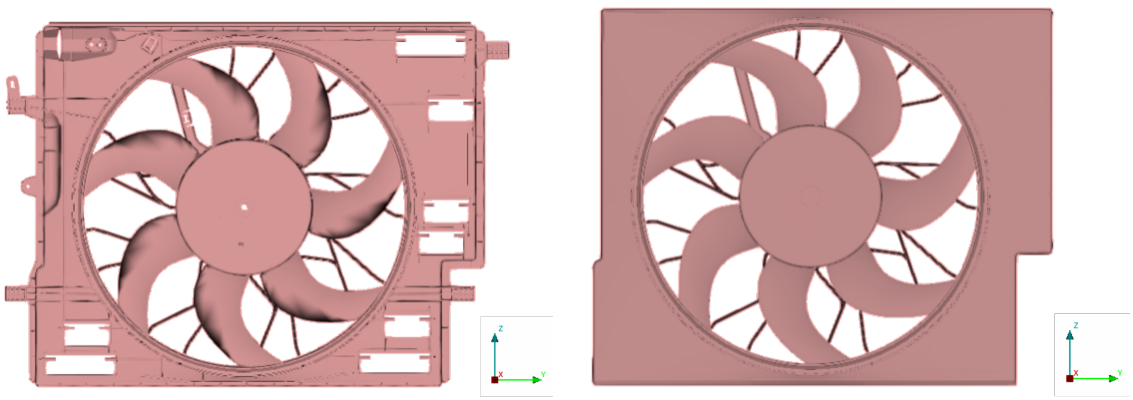

(a) **(b)**

Figure 3.2: *Fan 1*'s CAD model before (a) and after (b) the clean-up.

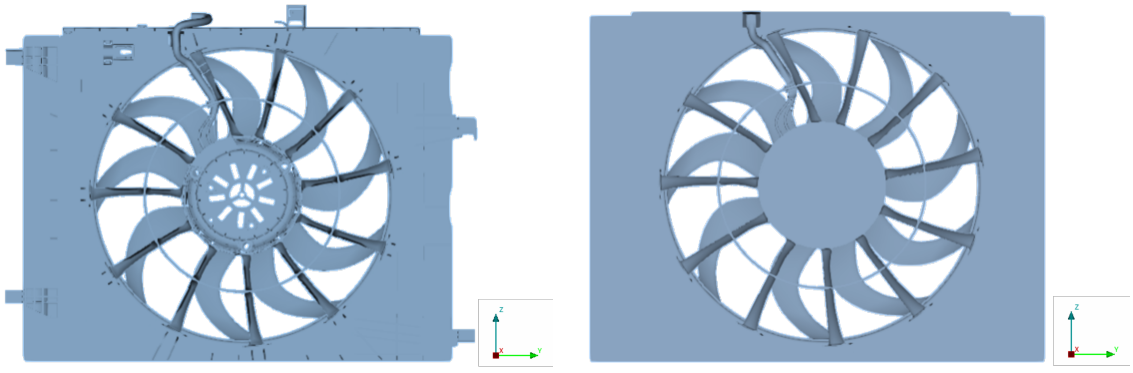

(a) **(b)**

Figure 3.3: *Fan 2*'s CAD model before (a) and after (b) the clean-up.

Regarding the test rig model, it involved the geometry of the inlet and outlet test chambers split by the front and back faces of the separating wall, where a custom test section opening for the fan could be created. The boundary faces through which the air flow enters and leaves the test rooms, i.e. the inlet and outlet of the system, were also defined. This simplified CAD model only contained the elements described by these boundary surfaces, avoiding any irrelevant objects present within the real chambers. As in the actual rig, the original model also included a removable noise-shielding structure located in the outlet room just downstream of the test section, intended to prevent the capturing of acoustic reflection during experimental measurements. However, this was removed from the model to reduce complexity. Moreover, the real locations of the inlet and outlet room pressure probes were well determined. The fact of knowing their coordinates was convenient to place pressure sample points in the STAR-CCM+ set-up accordingly, which would contribute to a more trustworthy comparison between simulation and experimental results.

Once each fan model is joined with a copy of the test rig geometry in a single CAD file, an extra surface entity in the form of a short square-like passage duct is created to connect the shroud with the separating wall. It appears in orange in Figure 3.6 and is a simplified form of the cooling module frame, where the radiator parts would be housed in potential prospective investigation to study their effect. In addition to this, an interface that wraps the impeller has to be generated in each fan's final CAD model, in order to later set up the MRF rotation effect in STAR-CCM+. This interface will allow the CFD program to divide the whole domain into two parts: the virtually rotating region – the area of fluid domain affected by the impeller's moving frame of reference – and the stationary region. The entity was produced as a surface of revolution with a simple cylinder-like shape to avoid geometric complication, but at the same time its boundaries were brought over tightly close to the impeller profile, in order to reduce modelling error as much as possible. In the outer ring, close to the shroud, the interface profile was made so that it followed an approximate meridional mid-line along the blade tip gap, in order to evenly split the radial clearance and to provide the rotor and stator inflation layers with roughly the same margin for eventual stretching. This geometric detail is shown in Figure 3.4b.

With this, each full-geometry file contains three separate groups of continuous surfaces: the *impeller*, the *motion interface* and the *outer boundaries*. The latter, as its name indicates, consists of every surface part that entails any surrounding boundary for the fluid domain. Therefore, this one is likewise subdivided in various parts: *inlet*, *outlet*, *inlet chamber walls*, *outlet chamber walls*, *inlet separating wall*, *outlet separating wall*, *frame passage* and *shroud*. The purpose of this separation is to allow setting distinct boundary conditions at each different boundary surface individually, as well as to enable piecemeal averaging and integration reports of field functions for modular post-processing. In Section 4.2.2 this can be appreciated.

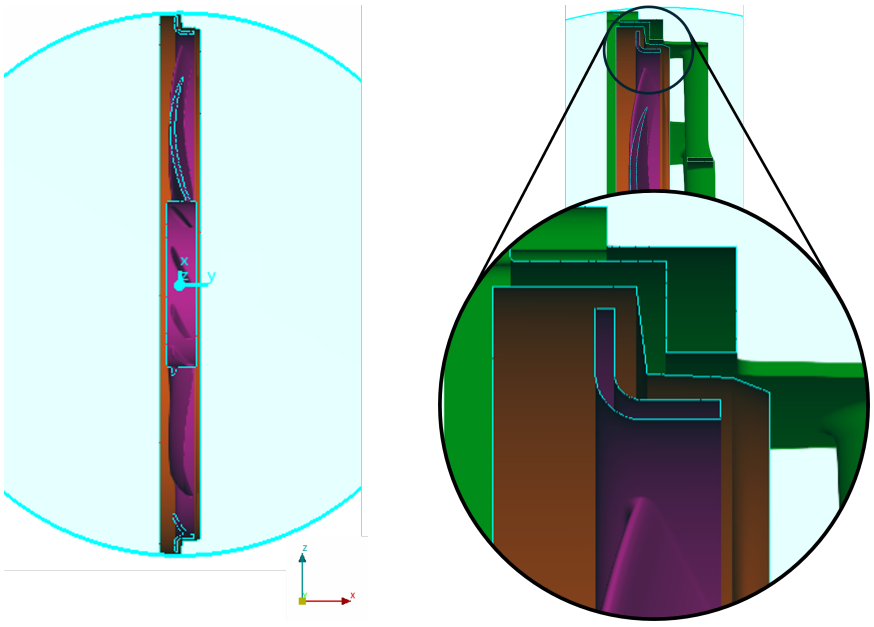

(a) (b)

Figure 3.4: Cross section of *Fan 2* showing the motion interface wrapping the impeller (a) and detailed view of the interface's profile delineating the tip gap between the impeller and the shroud (b).

3.1.3 Geometry model mesh

Once having obtained a CAD model fully defined by analytical surfaces in ANSA, this must be converted to a discretised geometrical entity, so that neutral CAD files for every model part – *impeller*, *motion interface* and *outer boundaries* – can be generated from it and conveniently imported into STAR-CCM+. Each of these will be a grid of points forming "differential" faces, and since their connecting straight lines define the model geometry that the CFD software will use as a baseline to discretise the domain, at least the same resolution as the computational surface mesh is wanted. To create this high-resolution grid in ANSA, two basic steps are carried out through the meshing module: firstly, nodes are spread over all the perimeters of the geometry parts with a convenient spacing for each particular area, depending on the desired resolution; then, when the perimeters are properly discretised, the surfaces are selected and meshed through any of the automatic tools. In general, these actions were conducted by using the *Length* mode within the *Perimeters* section to establish the distance between perimeter nodes, as well as the *CFD* feature within the *Mesh generation* menu to provide a homogeneous hex-type surface meshing. For the leading edges of the impeller and the spokes, the operation gets a bit more critical, since these are very delicate areas that must be smoothly outlined due to their impact on flow and acoustic features. To achieve so, the perimeter tool *Auto STL Spacing* was utilised instead, which in combination with the *STL* meshing provides a nice parameterised way of defining these narrow curves (see Figure 3.5). Once the surface grid is done, there are various formats in which it can be exported: IGES (*.igs*), STEP (*.step*), STL (*.stl*), etc.

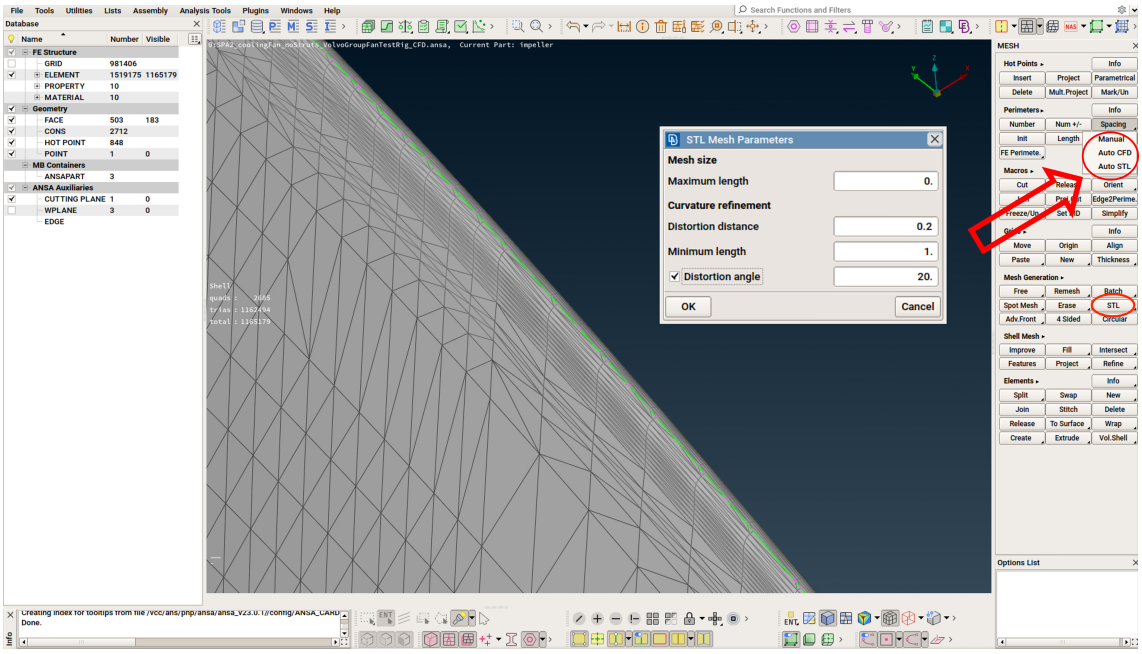

Figure 3.5: ANSA window showing the settings used for the surface mesh around a leading edge and the corresponding result.

It may be useful to note that, if the geometry model mesh is defined by a closed set of surfaces, it can be "filled" and exported as a volume mesh. Playing around with this possibility and depending on what type of format the geometry file is converted into, STAR-CCM+ is able to load the element either as a sheet body or as a solid body. For instance, if the geometry parts are imported as solid bodies with either the IGES or the STEP format, the 3D-CAD Model Manager in the program can recognise them as editable solids, thanks to which some modifications could be performed to build the final geometry model. Otherwise, if the part is imported as a sheet body instead, only minor changes can be made and massive clean-up operations are not feasible.

However, given that the CAD model was cleaned up in ANSA and no further tuning was needed in the simulation software, all the three geometry parts that form each fan model are exported from the pre-processor as surface grids in STL format. With this, they are imported into STAR-CCM+ and recognised by the program as triangulated surface meshes. The above format is chosen because of its robust compatibility and small size, which makes it easy and quick to import. Although the IGES and the STEP formats provide extra-tessellated resolution, they may not be worth it, since later the domain mesher will anyways re-build the imported modelled surface in accordance with the customised cell-size settings, typically targeting smaller surface sizes than the extra-tessellated ones. Apart from that, loading errors are usually encountered because of the large size and higher complexity of these formats.

When the surfaces are imported in STAR-CCM+, they are loaded into the tree menu within the *Parts* subfolder from the *Geometry* main folder. It is important

to check that all the surface meshes have been correctly read. This can be tested by performing a complete error scan through the *Surface Repair* module, and if any improvable area is found, it can be solved with either the manual repairing tools or the automatic one. After that, we still need to tell the simulation program what the volume occupied by fluid is like. To do so, the *Subtract* operation is used in first place, with the purpose of using the three imported parts to determine such volume. In this case, two distinct pieces of volume are required: one that will represent the main core of the domain and another that will be provided with rotating-MRF properties. For the first one, *outer boundaries* and *motion interface* are selected in the *Subtract* settings as input parts, and only *outer boundaries* as the target part. In this way, the volume that *motion interface* would enclose on its own is subtracted from the one surrounded by *outer boundaries*, hence obtaining a new "solid" entity consisting of the volume contained between these two surface parts. Similarly, for the MRF volume, *motion interface* and *impeller* will be the input parts, where *impeller* will be the target, thereby getting a volume with the external shape of the motion interface and an internal gap with the shape of the impeller. In that way, the latter volume will be embedded within the former, and only the impeller-shaped gap in the very inside will be left empty, not being part of the fluid domain. This can be observed through the cross section in Figure 3.14, where the computational mesh is included.

Once the two new volume parts have been obtained, the duplicates of the boundary face that they both have in common – namely the *motion interface* – must be merged together, so that the software interprets them as the same surface. This is done by opening the two volume parts in the *Surface Repair* tool and using the *Imprint parts* feature located within the *Global* tab. With the *Interactive* and *Multi-Part* modes activated, the motion interface of any of the parts is selected as the source and the remaining one as the destination. After the imprint, the geometry model is completely set up in STAR-CCM+' environment.

3.2 Boundary conditions

Prior to the discretisation of the fluid domain, the boundary conditions are determined. This order is usually recommended due to the fact that they will influence to some extent the way in which the computational mesh is built. And before specifying the conditions of the boundaries, the computational domain must be established in STAR-CCM+. That is done by assigning the two generated volume parts to *Regions*, obtaining a *static domain* and a *rotating domain*. After that, the boundary conditions are ready to be configured under these newly created regions, illustrated in Figure 3.6.

The physical representation of the domain boundaries is the element that drives the flow solution principally. It is for this reason that the modelling of boundary conditions is critical when reliable results are desired. Hence, the configuration of those are set according to the freedom of operational adjustment that we have

during real testing. The following are the boundary condition types used in the model, together with their specification settings and the surface parts to which they have been applied (summary in Table 3.2):

- *Inlet*. It is applied to the homonymous surface part and represents the intake through which the flow enters the test rooms in the real facility. Apart from the fan speed, the incoming flow rate can be freely tuned in the present experimental set-up, what could motivate the use of a *massflow inlet*. However, since for the analysis scoped in this work the flow is solely driven by action of the fan, the massflow rate will be an output instead. Subsequently, *stagnation inlet* becomes the chosen specification for this boundary. This condition requests for the values of total temperature and total pressure at the inlet, i.e. the static temperature and pressure that undisturbed fluid would virtually have in an upstream point sufficiently far away, resembling the effect of an environmental reservoir. In this way, regular atmospheric values are set for the total magnitudes, and these get automatically distributed among their static and dynamic components in accordance with the flow state that the fan induces. The flow direction specification is constrained as *Boundary-normal*, given that the circulation system of the rig is made to deliver properly aligned flow at the inlet. Also, the specifications of turbulent intensity and turbulent viscosity ratio were set as 0.085 and 10, respectively.
- *Outlet*. It is also applied to the surface part with the same name. As it can be guessed, it depicts the hatch through which the flow comes out of the test rooms. Since the test rig is controlled so that the conditions in the back chamber are always environmental, the specification is set as *pressure outlet* and the required value of static pressure is set to atmospheric. The flow direction specification is set as *Extrapolated*, given that the stream is able to be evacuated in a free, inhomogeneous way, and non-normal flow or even backflow can happen in some areas of the outlet. Besides, the same values of turbulent intensity and turbulent viscosity ratio as in the inlet were set.
- *Wall*. In general, this boundary condition sets no-penetration across the subject plane, i.e. zero wall-normal velocity. In addition, the tangential components can be constrained in two different ways through the following specifications, depending on whether the shear stress between the boundary and the fluid is of interest or not:
  - *No-slip*. Likewise, this setting simply establishes the tangential velocity components as zero, thus enabling non-zero wall-normal gradients and subsequent boundary layer build-up, as a regular wall would do. This phenomenon has to be certainly taken into account in the vicinity of the test section, as the substantial turbulence produced close to it and in the the fan itself is translated into large acoustic output. Therefore, the no-slip condition is applied to the *impeller*, the *shroud*, the *frame passage*, and also to the front face of the separating wall, on which the incoming flow impinges and whose generated turbulence might influence the entrainment effect and thereby the way the stream heads to the fan.
  - *Slip*. This method makes the values of the wall-normal gradients for the tangential velocity components equal to zero. Since that would avoid the

formation of boundary layer, it would involve some computational savings. This is convenient as long as it is used in areas where the effect of boundary layer is not decisive, namely the floor, roof and lateral walls of the two chambers (*inlet chamber walls* and *outlet chamber walls*), as well as the back face of the separating wall (*outlet separating wall*). Because of this, no inflation layer is generated on these surface parts.

Beyond the "regular" boundary conditions above, a physical specification must be set to model the fan rotation. As it is known, this is conducted by applying a rotating reference frame to the domain region in close contact with the fan, i.e. the so-called *rotating domain*. This is numerically equivalent to adding the corresponding tangential component to the velocity field of this region. To make this work properly, the two domain parts – static and rotating – must be physics-wise communicated, avoiding the interruption of the conservation equations across their linking boundary *motion interface*. To achieve this, the two duplicates of such boundary are selected together from the program's tree menu and the *Create Interface* feature is executed, generating an extra boundary item in the dropdown folder under each of the two regions. Additionally, it will also be found as a new element within the *Interfaces* folder, where the *internal interface* condition type must be set. Furthermore, within the *Tools* folder, the moving reference frame that will implement the rotational flow is to be created. Since the fan rotation axis is oriented in the x direction and passes through the origin of the global reference frame, the requested axis properties are set as (1,0,0) and (0,0,0), as per the Cartesian coordinate system. The rotation rate is entered according to the fan speed for the desired operating condition. Once the interface and the rotating MRF are configured, the motion specification of the *rotating domain* within the *Physics Values* folder is linked to the newly created frame of reference.

Surface part	**Boundary condition**
Inlet	Stagnation inlet
Inlet chamber walls	Slip wall
Inlet chamber separating wall	No-slip wall
Frame passage	No-slip wall
Shroud	No-slip wall
Motion interface	Internal interface
Impeller	No-slip wall
Outlet chamber separating wall	Slip wall
Outlet chamber walls	Slip wall
Outlet	Pressure outlet

Table 3.2: Overview of the boundary conditions set in STAR-CCM+ for each surface part.

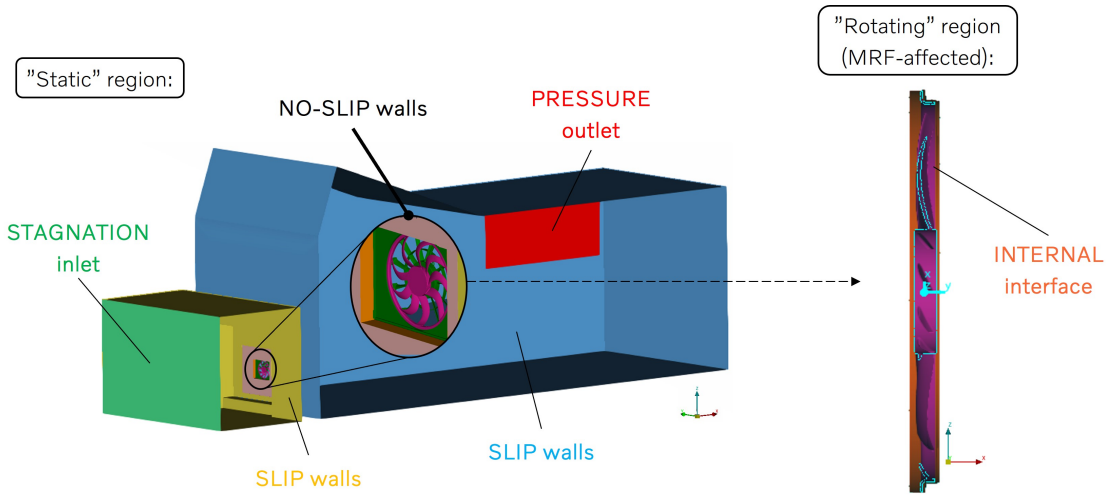

Figure 3.6: Visual representation of the domain boundaries and the physical conditions applied, where the differentiation between the static and rotating regions is highlighted.

3.3 Physics models

In order to realistically emulate the behaviour of the flow within the domain, the proper physics equations must be applied. As expected, the current time-independent CFD model will be based on a three-dimensional steady-state formulation. By default, the feature of solution data interpolation is also enabled for adaptation of calculated results after mesh regeneration. These general time and space settings are specified in STAR-CCM+' simulation setup through the model selection list:

- *Steady*
- *Three Dimensional*
- *Solution Interpolation*

Further on, three distinguishable sets of mathematical apparatus are considered for the full construction of the physics model: the core thermo-fluid dynamics, the modelling of turbulence, and the acoustic source models to extract noise output.

3.3.1 Fluid flow

For a fluid continuum, as seen in Section 2, the principal governing laws are the Navier-Stokes equations. In this case, they will operate in their steady-state RANS form, which implies Reynolds averaging and cancellation of the temporal term. Mainly consisting of the equations of mass continuity, equilibrium of momentum and energy conservation, the Navier-Stokes formulation is complemented by an equation of state. The energy and state equations can usually be omitted in fan applications due to negligible changes in density being expected, according to the discussion in Chapter 2.1. On this occasion, however, an initial test simulation with a non-optimised mesh was run at high with enabled compressibility, to foresee whether local density variations larger than 5% should be expected somewhere in the do-

main. This sanity check was motivated by the simplified hand calculation giving a relatively high blade tip Mach number for *Fan 1*:

$$M_{tip} = \frac{v_{tip}}{a} = \frac{2\pi n r_{tip}}{\sqrt{\gamma RT}} = 0.23 \tag{3.1}$$

If the result of the test simulation made this case deserve to be regarded as compressible, variable density would be used for the study. Although this would slightly increase the computational cost of the model, the fan performance would be better mimicked and the flow field variables more precisely represented in the region of interest, what may also generate more accurate aeroacosutic sources. Otherwise, if the aforementioned local variations are not that high, the constant-density specification could be set instead for the whole simulation study, making it a priori somewhat cheaper. Regardless of the density setting, the simulation is run using the segregated-flow method to resolve the pressure-velocity coupling, as justified in Section 3.5.1.

Therefore, with air media working under calorically perfect conditions, the ideal gas law (Eq. 2.5) is applied to the test simulation for computing local density, thus closing the system of governing equations. In addition, to assess the compressibility assumption, the maximum positive and negative percentage changes in density along the domain were reported through STAR-CCM+:

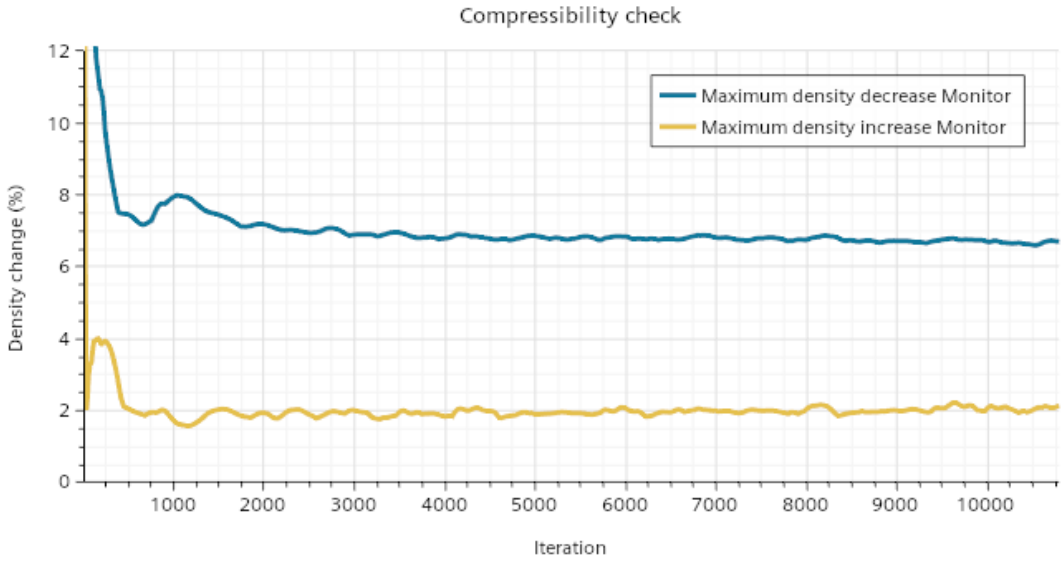

Figure 3.7: STAR-CCM+ monitor plot showing maximum density increase and decrease for the simulation test with *Fan 1* at 3200 RPM using the ideal gas model.

As can be observed in Figure 3.7, a maximum density increase of only 2% occurs for *Fan 1* due to flow compression, whereas a decrease of 6.7% is found in expansion regions. Although the latter datum indicates that taking compressibility into account might be necessary for cases approaching 3200 RPM, the flow field revealed that the minimum density is found as a tiny spot likely to have minor impact

(Figure 3.8). Even so, a comparison of the computation time per iteration with that of an equivalent constant-density case showed a difference of only 0.05 seconds longer for the variable-density model, approximately 13%. In fact, the simulation with constant density proved to be less stable, requiring 6000 more iterations to converge (see the established convergence criteria in Section 3.5.4). Subsequently, the total computation time spent was of the same order of magnitude for both simulations, even 23% lower for the compressible model.

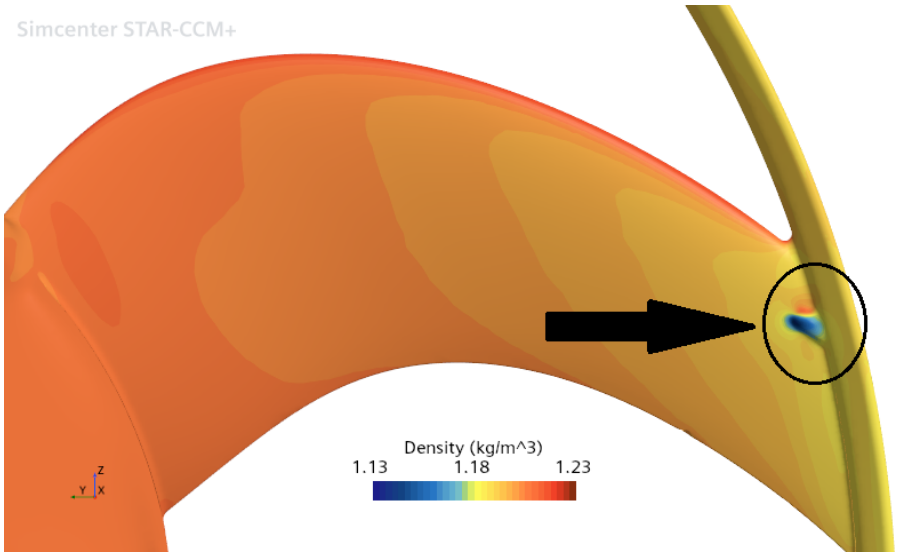

Figure 3.8: Converged flow field solution over one of *Fan 1*'s blade showing a maximum density decrease of 6.7% at 3200 RPM.

Concluding that the computational expense does not pose any big concern, the ideal gas model was maintained. It was indeed kept for simulating all the operating points, thus meeting simulation consistency between cases. On top of that, as it can be appreciated in Section 3.5, an extra reason to use non-constant density and non-isothermal flow is that the arranged configuration of solver settings could be kept for a hypothetical study with heat addition, reinforcing the versatility of the simulation setup.

In conclusion, the following items are selected in STAR-CCM+' model menu to solve the flow thermodynamics:

- *Gas*
- *Segregated Flow*
- *Ideal Gas*
- *Gradients*
- *Segregated Fluid Temperature*

3.3.2 Turbulence

As part of a good CFD analysis, an adequate turbulence model must be set up to get reliable and efficient simulations, especially in a clearly turbulent case like the current one, in which high Reynolds numbers and widespread regions of steep

velocity gradients are found due to the great interaction of rotational flow with the solid walls of the fan body. In line with the scope of the project, which involves sticking to simple and economical analysis, the use of a Reynolds-averaged turbulence approach from the eddy-viscosity type is preferred. Thus, the governing RANS formulation previously presented brings in two additional equations that will model the Reynolds-stress term and provide the turbulent quantities.

The SST k-ω turbulence model is theoretically more reliable for capturing flow phenomena in turbomachinery-based applications, compared to other eddy-viscosity models. This stems from its ability to accurately predict turbulent kinetic energy production in flows governed by adverse pressure gradients [14]. After all, this model is the most complete and thorough within the inexpensive eddy-viscosity family.

However, research such as Fares'[19] reveals that in some cases, the computationally lighter k-ε model can produce results that match or even outperform SST k-ω in accuracy. For this reason, while the primary focus of this study is on SST k-ω, results from standard k-ε will also be extracted. In this way, comparison between the most comprehensive model in the family and its most economical counterpart will be obtained, getting a reference on the potential margin of improvement that our methodology could have in terms of computational efficiency. Furthermore, it should be noted that both k-ε and SST k-ω were kept with STAR-CCM+' default settings and correlation constants.

Since the k-ω turbulence model is designed for regions of low Reynolds number while the k-ε model is better suited for high-Reynolds flows, each requires a specific near-wall configuration to perform properly. This involves separate inflation-layer meshing, as elaborated in Section 3.4.2. Besides, a wall treatment specification must be set accordingly, in order to condition the turbulence model to the generated near-wall mesh. Hence, the SST k-ω simulations are set up with low-y^+ wall treatment and the k-ε simulations will utilise high-y^+ wall treatment.

Alternatively, STAR-CCM+ offers the all-y^+ wall treatment, a more versatile option that is capable of fairly adapting both types of turbulence model to a reasonable inflation layer, regardless of their Reynolds-number specification. This would ease the modelling work and could even enable the use of a less demanding inflation mesh, contributing to reduce the computational cost. Nonetheless, the study at hand prioritises a more thorough boundary-layer prediction, as it is essential for accurately capturing dipole-type acoustic sources, which are key for the present purpose. Because of this, a separate near-wall configuration for each turbulence model is preferred in this case.

As a conclusion, the described turbulence configurations are specified in STAR-CCM+ with the following model selection:

- *Turbulent*
- *Reynolds-Averaged Navier-Stokes*
- Either k-ω *Turbulence* or k-ε *Turbulence*, as per the simulation case.
- Either *SST (Menter)* k-ω or *Standard* k-ε, as per the simulation case.
- *Wall Distance*

- *Low-y+ Wall Treatment* for k-ω; *High-y+ Wall Treatment* for k-ε.

3.3.3 Acoustics

In the last place, the aeroacoustics models must be activated to enable the computation of irradiated acoustic power based on the resulting RANS flow field. The Curle and Proudman broadband-noise source models, derived from Lighthill's acoustic analogy, are selected to fulfil this task. Hence, the computation of the acoustic output is activated in STAR-CCM+ by picking the following items in the model selection list:

- *Aeroacoustics*
- *Broadband Noise Sources*
- *Noise Source Models*
- *Curle*
- *Proudman*

The Curle model aims to predict the acoustic power associated with disturbances of dipole character, caused by turbulent pressure fluctuations arising from the interaction of the boundary layer with wall surfaces. As is seen in Eq. 2.26, its integration over the fan surface mesh gives the total surface acoustic power.

In contrast, the Proudman model accounts for the effect of quadrupole-character phenomena across the volume mesh, extracting the acoustic power generated by isotropic turbulence from swirling regions such as wakes or the outer levels of the boundary layer. Its total contribution across the domain is represented by Eq. 2.27.

The combined total of surface and volume acoustic power represents the overall acoustic output considered in these models, which can be expressed as sound wattage level in decibels. However, because of the low Mach number in the present application, the Proudman power contribution is negligible compared to Curle's. Even so, it will indeed be taken into account, in order to analyse the model's ability to discern relative variations in the quadrupole contribution across the different fans, turbulence models, and operating conditions.

3.4 Computational mesh

A very important stage in the modelling process is the mesh construction. The computational domain must be discretised with sufficient refinement to achieve the desired level of resolution, ensuring consistent and trustworthy results for the intended purpose. Nevertheless, that will also depend on the solver settings, the initialisation values, and the continuum model used.

On the other hand, a second reason that motivates a proper mesh generation is naturally the search for computational efficiency. Although a multi-million-cell sim-

ulation of this model can be solved relatively quickly by means of HPC, it is always good engineering practice to create an economical mesh. This ensures feasibility in view of potentially running a massive set of simulations, like a sweep through a range of operating conditions, or in case of needing to use the setup with more limited computational resources.

In the present work, given the highly turbulent character of the flow and the interest in the acoustic production, the greatest effort in terms of meshing is concentrated on the inflation layer construction and on the refinement of the wake region, especially the former. Thus, after establishing suitable values for the general meshing parameters, mesh independence studies for these two aspects will be conducted to provide the methodology with reliable and efficient computational domains.

3.4.1 General mesh settings

For the smooth functioning of the MRF-based fan rotation, a conformal mesh transition must be achieved between the rotating region and the static region. That means that the cell faces from the rotating region that are in contact with the MRF interface must be coincident with the cell faces that conform the interface from the side of the static region, as can be observed in Figure 3.14. This contributes to a more desirable discretisation, avoiding numerical error in the computation of face fluxes across the motion interface and contributing to a reliable propagation of the physical quantities. Conformal arrangement is achieved if both rotating and static regions are selected to be discretised under the same *Automated Mesh* operation.

As often in similar CFD applications, the following automated meshing tools were enabled as the basis for the generation of the grid:

- *Automatic Surface Repair*. It solves surface malformation of the imported geometry that could eventually lead to mesh errors or calculation issues.
- *Surface Remesher*. It regenerates the surface mesh of the imported geometry with a higher quality to facilitate the formation of the volume mesh.
- *Polyhedral Mesher*. It establishes the use of polyhedrons to create the core volume mesh.
- *Prism Layer Mesher*. It enables the generation of a prismatic inflation mesh for near-wall regions to assist with boundary layer modelling.

The main parameters of the automated mesh operation, accessible within the *Default Controls* subfolder, were adjusted such that the non-critical regions of the domain were provided with coarse grid and guaranteeing a smooth transition in element size. Their values were established as shown below after some trial-and-error process:

- *Target Surface Size*: 0.1 m
- *Minimum Surface Size*: 0.001 m
- *Surface Growth Rate*: 1.3
- *Volume Growth Rate*: 1.2

- *Maximum Tetrahedral Size*: 0.1 m

Detailed information on how STAR-CCM+ defines the meshing parameters can be found in its User's Guide [20]. The rest of the general mesh settings were left at their default values.

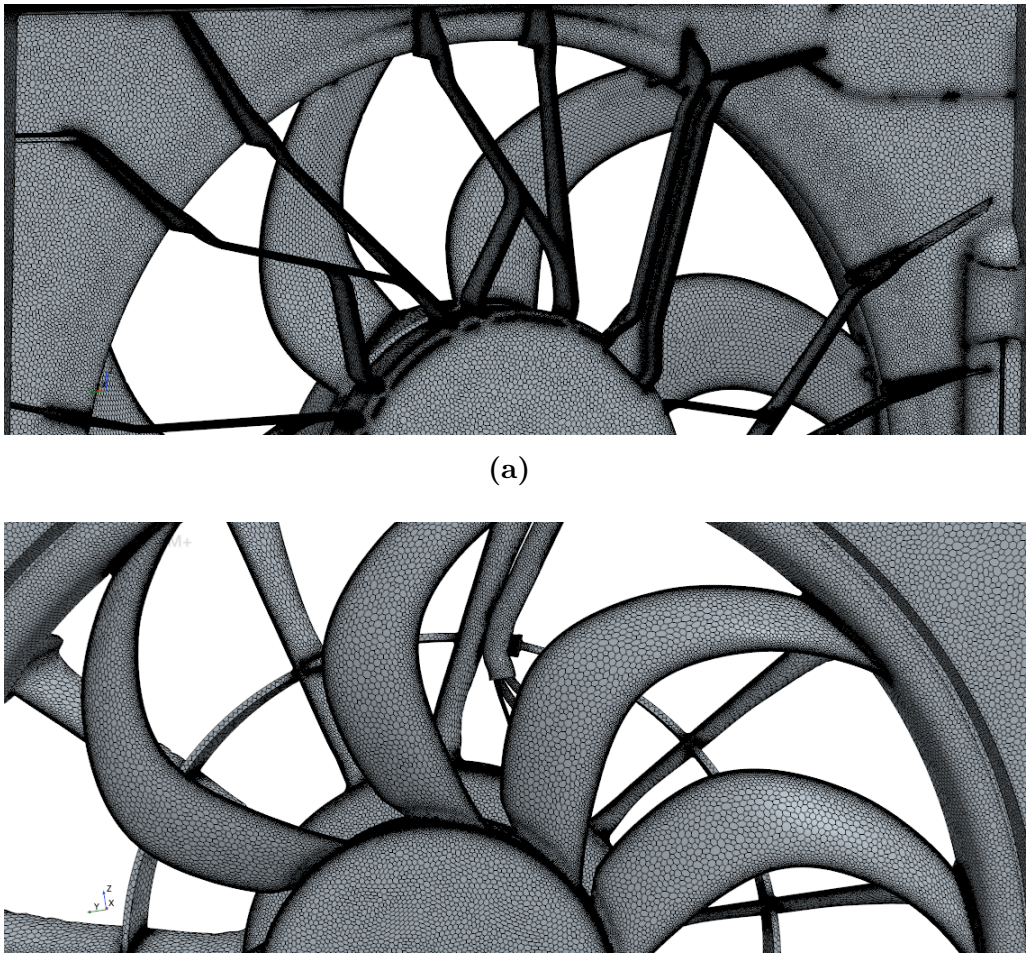

(a)

(b)

Figure 3.9: Details of the obtained surface meshes on *Fan 1*'s back side (a) and *Fan 2*'s front side (b).

In addition, a more thorough refinement was performed in the proximities of the test section, where areas of geometric complexity and no-slip wall boundaries are found. This was carried out through surface *Custom Controls* that included a reduction of target and minimum surface sizes, which was applied to the surface parts identified as *impeller*, *frame passage*, *shroud*, as well as *motion interface* on both its rotating and static sides, ensuring correct generation of conformal mesh transition. After some meshing attempts, the following were found to be refinement values that provided good enough surface resolution:

- *Target Surface Size*: 0.003 m

- *Minimum Target Size*: 0.0001 m

In qualitative terms, these settings enabled visibly proper definition of relevant surface details like the sharpness of trailing edges and the curvature of leading edges and spokes, all of which are certainly decisive for the capture of significant flow gradients and thereby acoustic output. This level of refinement can be observed in Figures 3.9 and 3.10.

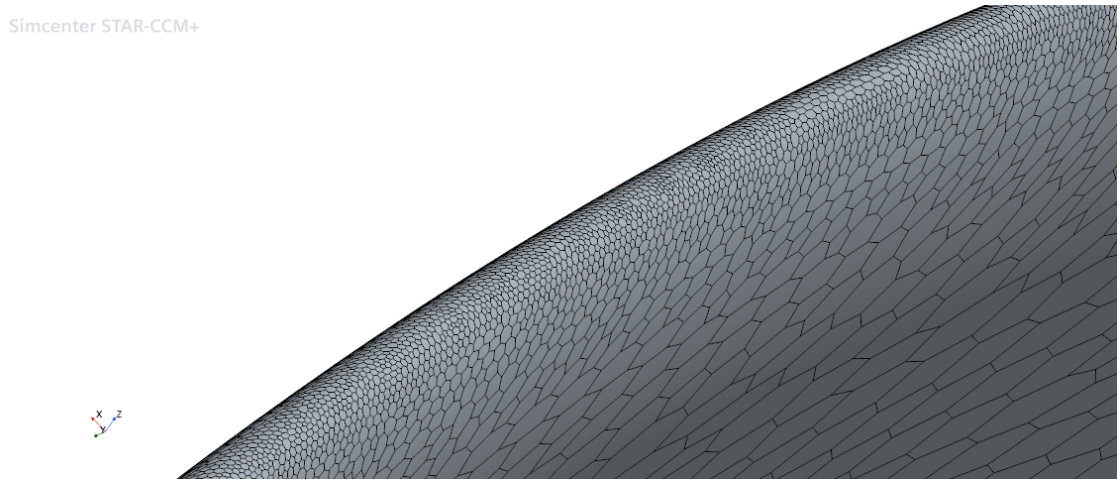

Figure 3.10: Close-up view of the surface mesh at one of the leading edges.

Even so, to obtain a more objective idea of the overall suitability of the grid, mesh quality reports were generated and regularly checked. These included distributions of values of aspect ratio, skewness angle, face validity, and volume change over the cell population, as illustrated in Figure 3.11. The ranges of target values to consider the mesh acceptable were established on the basis of STAR-CCM+' recommended guidelines [20]:

- *Cell Aspect Ratio*: 0.8–1
- *Skewness Angle*: 0–85 deg
- *Face Validity*: 0.95–1
- *Volume Change*: 0.01–1

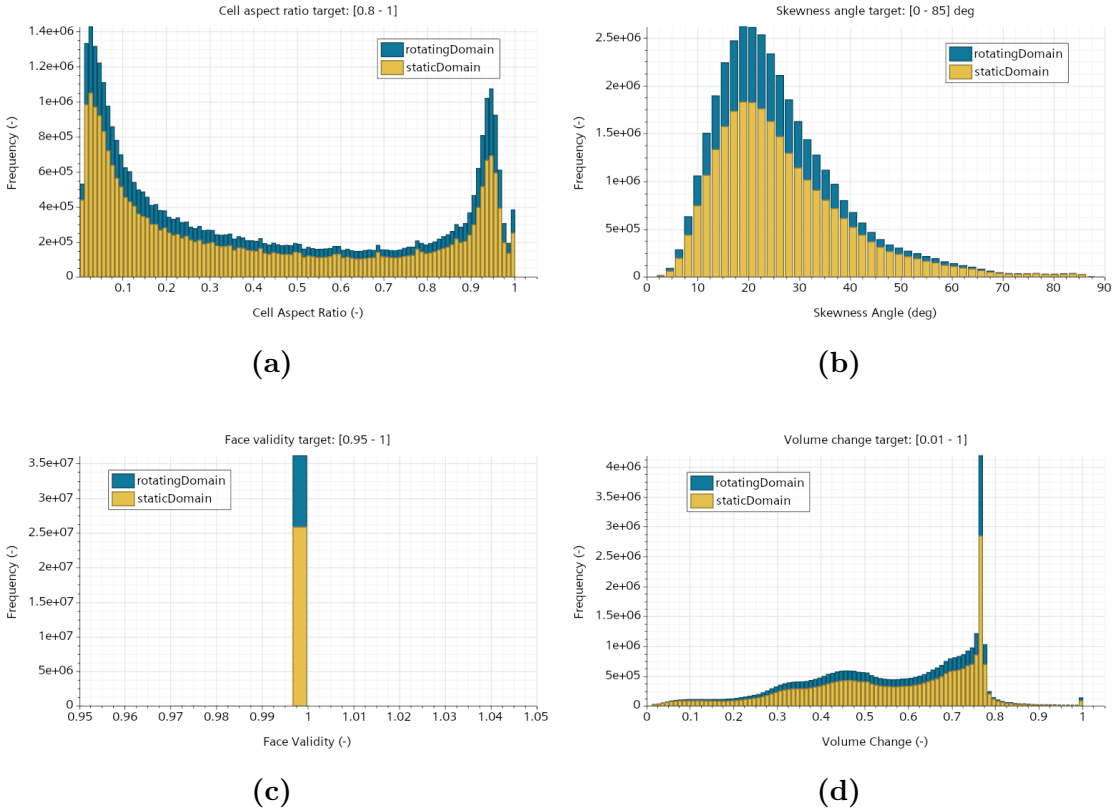

Figure 3.11: Plots of mesh quality report showing the distribution of each geometric parameter over the cell population in an SST k-ω grid, including aspect ratio (a) skewness angle (b), face validity (c), and volume change (d).

3.4.2 Inflation mesh

As is known, a correct prediction of the boundary layer is key for successful flow simulation. To achieve so, the critical near-wall areas of the volume mesh are generated as an inflation layer that fits in with the characteristics of the employed turbulence model and wall treatment. At the same time, the set of prism sublayers forming the inflation region must be configured in accordance with the needed level of near-wall discretisation, also aligning with the desired computational affordability. In the present case, the boundary layer significantly influences the acoustic output, particularly in Curle's noise prediction from dipole-character sources. Because of this, a mesh independence study will be a key cornerstone of the methodology, determining the best-performing inflation mesh sizing. On the basis of Pietroniro's method [22, 21], the approach will examine the sensitivity of the simulation results to the total wall-normal extent of the inflation layer.

When constructing a prism-layer mesh, there are normally three main parameters that can be tuned to fully define it in the wall-normal direction: the number of prism layers, the stretching factor and the total thickness. The approach is to run a set of simulations varying the number of prism layers with a fixed stretching factor, while the height of the layer closest to the wall is also kept approximately constant.

The latter is achieved by parallelly adjusting the total thickness with each change in the number of prism layers, so that the wall-y^+ values, which are proportional to the first-layer height, remain unchanged. In that way, we will produce stacking of layers in a sequential manner, keeping the wall distance of previously added layers mostly constant, as seen in Figure 3.12.

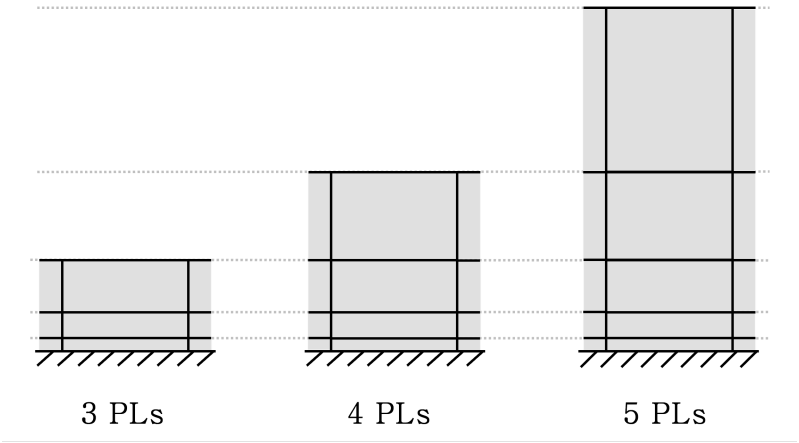

Figure 3.12: Simplified sketch showing the construction approach adopted to study the effect of different prism layer (PL) extents on simulation results.

The stretching factor will determine the distribution of prism layers along the total inflation height, which is important since the steepest gradients and most intense turbulent phenomena occur in the inner part of the boundary layer. Following Pietroniro's work [22], a moderate value of 1.3 was fixed for the current study, which means that each prism layer is 30% taller than the one beneath it. It may be worth pointing out that, although the influence of stretching factor variation could have been part of the study, doing so would have overcomplicated the study. Also, the simulation software offers more advanced settings to manage the prism-layer mesh in greater detail across the domain's geometry, which were not utilised in this case due to the commitment to a simple modelling process.

The wall-y^+ values were arranged to meet STAR-CCM+' guidelines on wall treatment for proper turbulence modelling. According to these, the target must be to keep the values fully below 1 when using the low-y^+ wall treatment and above 30 in the critically turbulent areas when using the high-y^+ wall treatment [20]. The total prism-layer thickness needed to obtain the desired wall y^+ was initially estimated through the theoretical boundary layer formulas for flat-plate application, which consider the fan speed and flow properties (see Section 2.1). Since the wall-y^+ distribution is inherently dependent on the flow field, a simulation must be completed to judge whether the total thickness is adequate. Therefore, iterative re-adjustments are made until the wall-y^+ values over the fan surfaces fall within the target range, as demonstrated in Figures 3.13a and 3.13b.

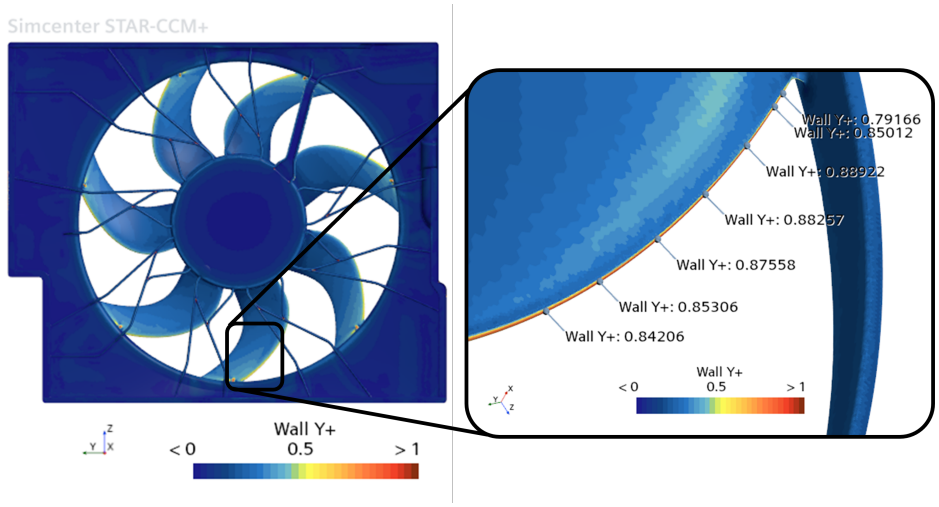

(a)

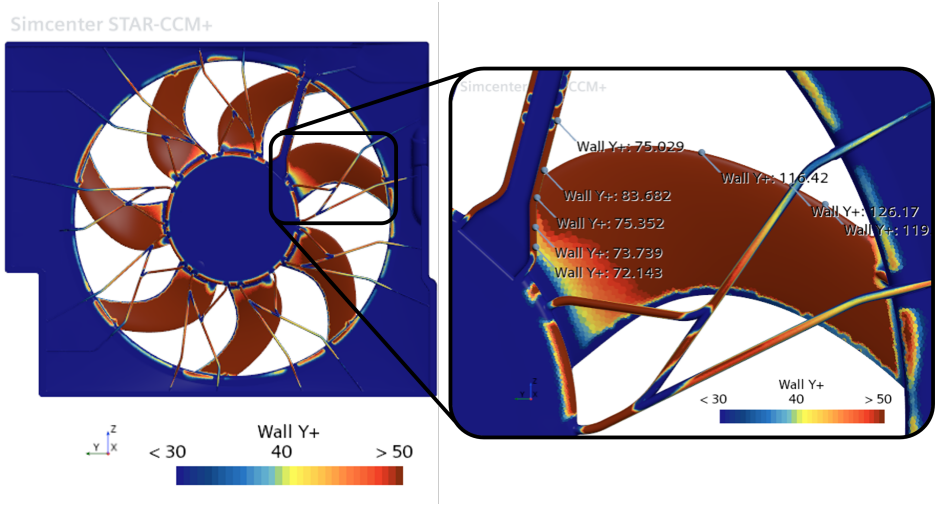

(b)

Figure 3.13: Example case of obtained simulation results featuring acceptable wall-y^+ surface fields for SST k-ω (a) and k-ε (b) models on *Fan 1*.

This process will be meaningful for the SST k-ω setup, whose mesh allows a wide range of numbers of prism layers to be tested. This is because a low-y^+ turbulence model requires several thin layers accumulated in the lowest region of the inflation zone to perform adequately. In fact, the number of prism layers will be varied from 8 to 22 for this analysis, as seen in Figure 3.14a. Conversely, such an extensive sweep of layers is not so well-suited to a k-ε setup, whose approach based on high-y^+ inflation cells results in the boundary layer being fully covered by just a few thick prism layers, typically 2–5 depending on the geometry. A greater number would imply an unnecessarily thick inflation mesh and could even lead to meshing issues in narrow areas of the domain. For the present high level of geometry complexity and the chosen stretching factor, only 2 and 3 layers provided a meaningful inflation mesh without causing the automatic mesher to trigger an error. Thus, a more limited evaluation was performed for the k-ε model. An example of a 3-layer inflation mesh is shown in Figure 3.14b.

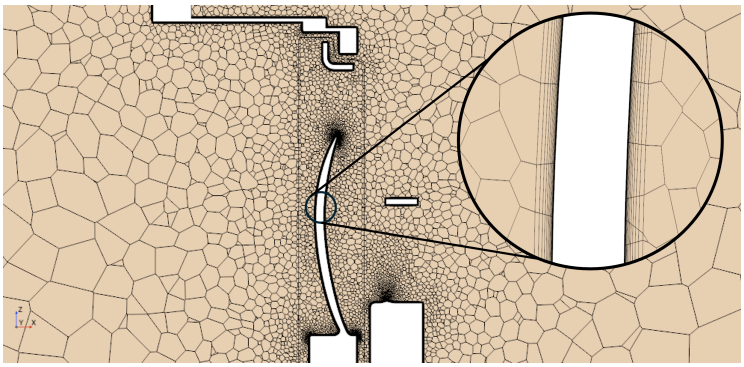

(a)

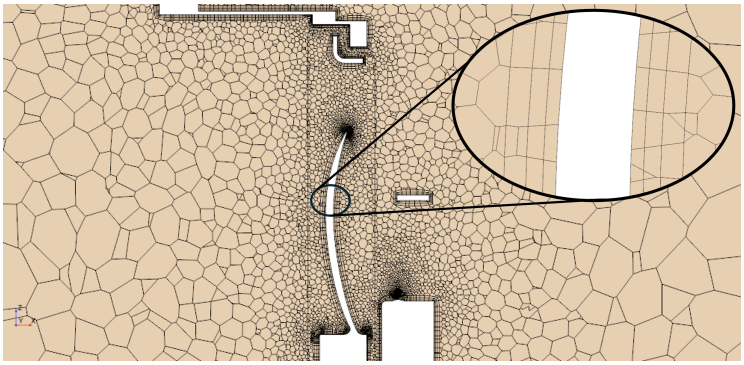

(b)

Figure 3.14: Cross section view of two grids obtained for *Fan 2* close to the impeller, highlighting the difference between an inflation mesh adapted for the SST k-ω model (a) and another for the k-ε model (b).

For each fan specimen and turbulence model, the analysis was moreover conducted at two contrasting rotational speeds, with the aim of ensuring the robustness of the boundary layer prediction across the addressed working range. In representation of a lower speed case, the inflation mesh study for both *Fan 1* and *Fan 2* was conducted running the simulations at 800 RPM. For the case at higher speed, *Fan 1* was run at 2400 and *Fan 2* at 3025 RPM, as experimental data from the fan suppliers for these operating points were available beforehand, serving as a sort of reliable reference or pre-validation. Given that the boundary layer enlarges at lower flow velocities and shrinks at higher ones, the adjustment of the total inflation mesh thickness as a function of the number of prism layers was carried out independently for each rotational speed, resulting in different first-cell heights with the wall-y^+ values being kept within the targeted span. Similarly, this adaptation of the total thickness to the operating velocity is later applied to every speed within the examined range during the final validation simulations.

Once the effect of varying the number of prism layers is analysed at the two different speeds, the configuration that performs most conveniently at both will be adopted as the recommended one to conduct conclusive simulations, such as the validation of the model across the aforementioned working range. The desirable

configuration will correspond to the minimum number of prism layers beyond which further increase no longer produces substantial differences in the simulation results, i.e., when the variation in the output variables with respect to the addition of prism layers exhibits an asymptotic trend. The output variables observed in this evaluation include Curle and Proudman sound levels and performance metrics such as volume flow rate, total pressure rise, and fan torque. Together with the minimisation of cell count, the impact on the total number of iterations is also monitored in the search of rapid and stable convergence, enhancing computational efficiency.

3.4.3 Wake refinement

The mesh independence study for the wake region aims to assess the potential of its refinement in capturing volume-based acoustic power generated by quadrupole-type turbulence in the flow trail downstream of the fan. The Proudman source model is particularly relevant in this area, where highly disturbed flow and homogeneous turbulence dominate. Although contributions of quadrupole character are expected to be minimal compared to dipole-type sources due to low Mach numbers [23], this reasonably quick analysis on wake resolution is found to be additionally convenient for studying the numerical stability of the simulation.

The resolution of the wake flow field will be analysed by testing different levels of refinement, corresponding to reference cell-size values of 0.1, 0.05, 0.025, and 0.01 metres, with the first one serving as the baseline and matching the *Target Surface Size* core setting (see Section 3.4.1), i.e., no refinement. This sweep is performed in both the k-ε and SST k-ω setups, and only after completing the inflation-mesh independence study and deciding the number of prism layers, to prevent any overlap between the effects of the two meshing analyses. The simulations for this study are run at rotational speeds of 2400 RPM for *Fan 1* and 3025 RPM for *Fan 2*, not only due to the availability of supplier data for those working points, but also because of the proximity to the usual upper operating limit, at which conditions of highly chaotic flow provide a preferable worst-case scenario to test the wake resolution. The impact of varying wake cell sizes on simulation results was analysed in relation to total cell count and number of iterations. The most desirable cell size was identified using a criterion similar to that for the inflation mesh study and was established as the recommended choice for conclusive analysis, thereby selecting it to perform the model validation across the working range.

To execute this type of refinement, STAR-CCM+ provides an in-built tool within the meshing module on purpose for that. It appears as a surface custom control labelled *Wake Refinement*. By means of this tool, the programme will generate a virtual conic volume arising from the selected surface in the chosen direction. This body, which mimics the shape of the wake, can be arranged in length and apex angle according to the expected size and shape of the wake. The desired cell size is also set, and when the mesher is run, it will execute the specified refinement operation in the domain's volume region delimited by the defined conic body. These details can be seen in Figure 3.15.

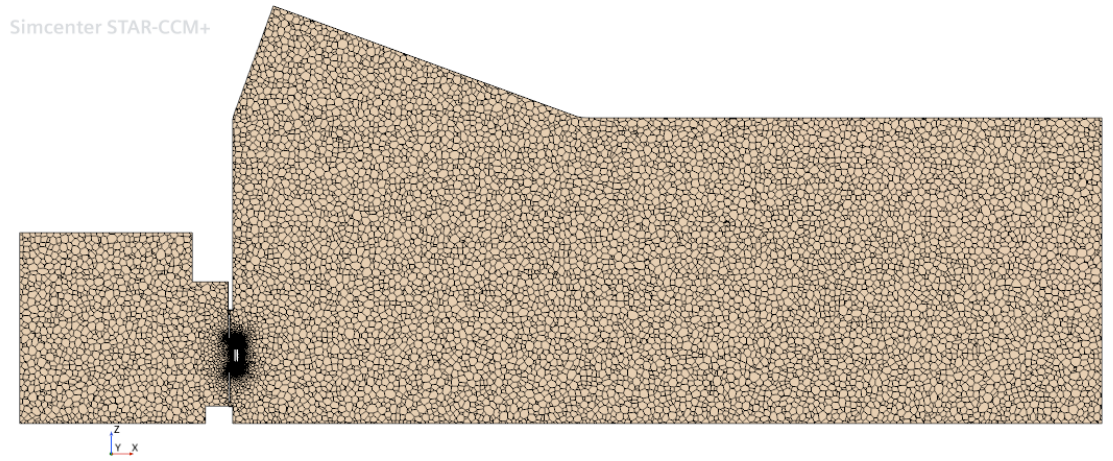

(a)

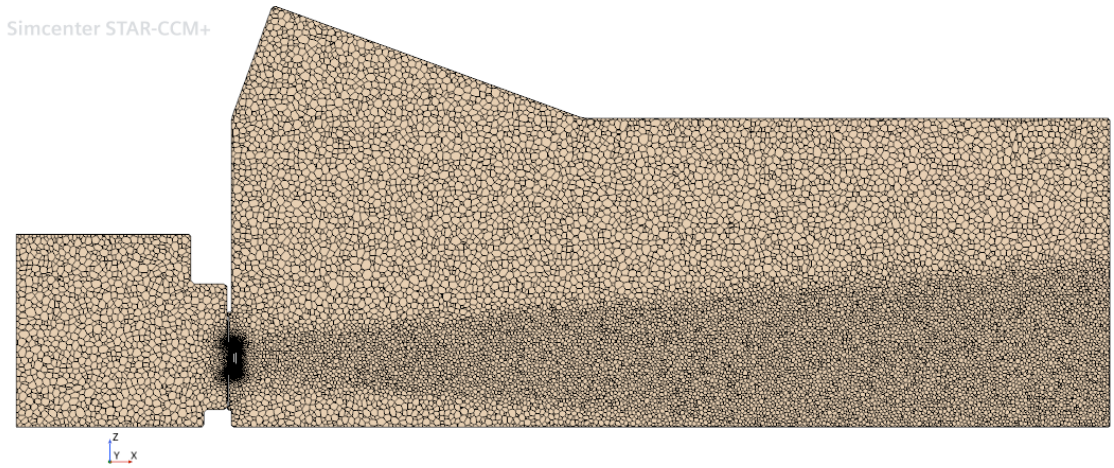

(b)

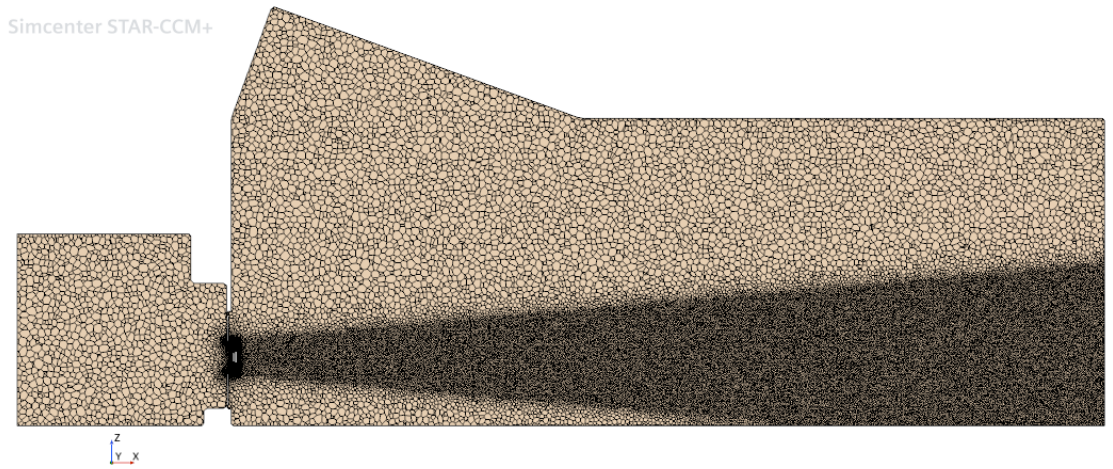

(c)

Figure 3.15: Longitudinal cross section of the meshed domain showing different levels of wake refinement: cell sizes of 0.1 (a), 0.05 (b) and 0.025 (c) metres.

3.5 Solver and solution convergence

To produce a solution field with minimal error and in acceptable computation time, apart from fine-tuning the mesh, the calculation process through which the set of model equations are solved must be optimised. This is intended by applying careful criteria on the adjustment of the solver, discretisation scheme, initial field values, and criteria of convergence.

3.5.1 Solver settings

The tuning of the numerical solver, along with the mesh construction, is a relevant factor for the efficient solving of the flow field. On this occasion, since the methodology development focused on other aspects, optimisation of the solver was not studied exhaustively. Instead, after configuring its pivotal features according to the problem characteristics, a selection of settings that ensured reasonably good computational performance was identified through experimentation, subsequently maintaining it for all simulations.

As specified during the selection of physics models, the solver was configured to handle pressure-velocity coupling using the segregated approach, given its higher efficiency for cases where shocks and discontinuities are not expected. The *Enable Enhanced Stability Treatment* feature was additionally activated for the segregated flow model, with the intention of reducing potential oscillations in the solution trend stemming from numerically demanding areas of the mesh.

Similarly, the energy equation was solved using the segregated temperature mode, as outlined in Section 3.3.1. Although temperature effects in this case are minimal, configuring the solver with settings tailored for an ideal-gas model ensures consistency and facilitates future work involving heat exchangers, where temperature variations become more significant.

With this, the following values for the under-relaxation factors were used to control the algorithm that couples the flow variables at each iteration, providing a reasonable convergence stability:

- Velocity: 0.3
- Pressure: 0.4
- Energy: 0.5
- k-ω turbulence: 0.5; k-ε turbulence: 0.8
- k-ω turbulent viscosity: 0.8; k-ε turbulent viscosity: 1

The resulting system of equations is computed iteratively using the algebraic multigrid (AMG) linear solver, which is the default solver offered by STAR-CCM+ in this case. It is based on an adaptative process that cyclically updates the resolution of the numerical grid, delivering an efficient solution convergence.

3.5.2 Discretisation scheme

The selection of the discretization scheme is an important factor in ensuring accurate flux computation across cell faces and maintaining reliable propagation of flow variables throughout the domain.

As usual for problems that involve highly convective flow, an upwind scheme is preferred here. In particular, second-order discretisation is selected for pressure, velocity, temperature, and turbulence fields, in contrast to first-order, which would introduce excessive numerical diffusion, or to the third-order MUSCL scheme, which is designed for extremely steep gradients and would suppose a computational cost unlikely to be justified for the present application.

3.5.3 Initialisation values

The established initial state of the field values, at which the calculation starts, also influences the solution trend. If they are set according to a right estimation of the final solution, it can help to fulfil a rapid convergence, since the flow field will have been approached to the correct result. In this case, these were the initial values set:

- Pressure: 0 Pa. Since most part of the domain is at reservoir conditions, the pressure is approached as atmospheric.
- Static temperature: 300 K. Similarly as for the pressure, the fluid is mostly present in static conditions, so the total temperature is estimated to be equal to the environmental static temperature.
- Turbulence intensity: 0.085. This corresponds to the usual value for fully-developed turbulent pipe flow conditions, an assumption that we take for this case study.
- Turbulent velocity scale: 1 m/s. Since not much focus was put on the configuration of some advanced settings, the turbulent velocity scale was left as default.
- Turbulent viscosity ratio: 10. The default value that was also kept for the specifications of inlet and outlet boundary conditions was also used for initialisation.
- Velocity: (0,0,0) m/s. As happens for pressure and temperature, given that the flow is close to stagnated conditions for most of the included volume, null velocity is set for any direction at the initial iteration.

3.5.4 Convergence criteria

The residuals of the RANS equations are commonly the most preferred and reasonable metric to track solution progress, as they indicate the compliance with the governing physics. In this case, the residuals of the turbulence quantities, i.e. TKE (turbulent kinetic energy) and SDR (specific dissipation rate), are of particular interest, given the highly turbulent character of the showcased scenarios.

In addition, relevant performance parameters such as flow rate or shaft torque are in most cases convenient as well, and therefore monitored in the present study. Three different definitions of pressure rise were also computed and considered within the convergence criteria: the static pressure rise between the inlet and outlet boundaries; the pressure difference measured according to the locations of the real probes in the inlet and outlet chambers of the rig; and the total pressure jump across the test section, obtained by calculating the rise between the area-averaged values of two virtual discs located just upstream and downstream of the fan.

In the last place, local values of turbulent kinetic energy at three different distances downstream of the fan were also computed to monitor wake stabilisation, as shown in Figure 3.16. One point was located 0.1 metres behind the fan, another at 0.5 metres, and the last one at 1 metre, all of them along the axis of rotation. Since this is a region of high physical unsteadiness, fully steady trends are not expected for all simulation conditions. Instead, the objective was to achieve at least a controlled oscillation.

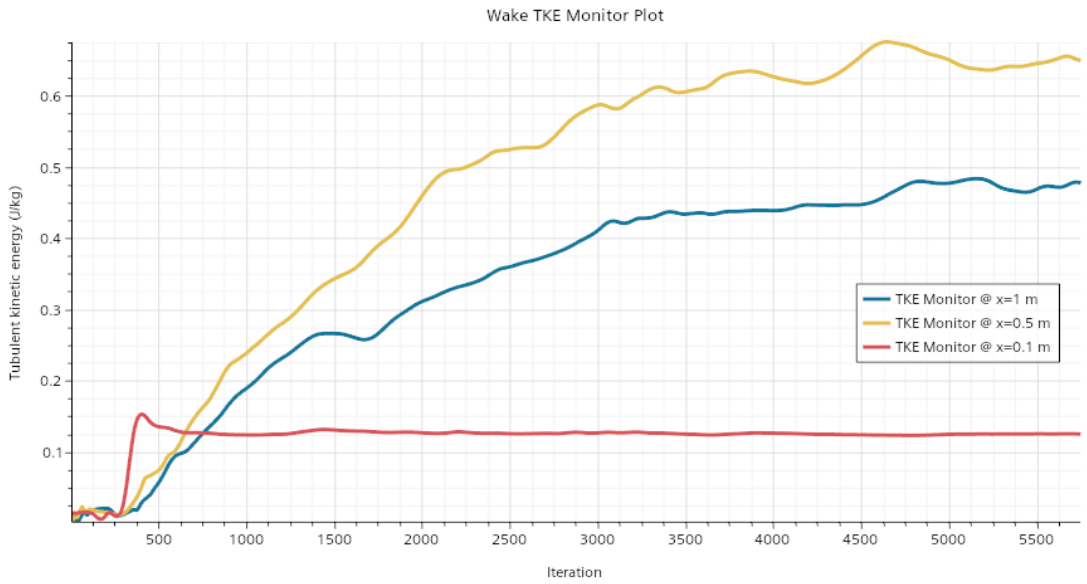

Figure 3.16: Monitoring plot showing the solution convergence of turbulent kinetic energy at three different points of the wake.

Among all these metrics, only the parameters that in general exhibited noteworthy sensitivity during simulation solving were employed to trigger the automatic stop of the simulations. Thus, the stopping criteria were finally programmed through the following conditions and cut-off values, which were conceived to accept a reasonable tolerance aligned with the unsteadiness level of each variable's convergence trend:

- Minimum TKE residual of 1%.
- Minimum SDR residual of 1%.

- Asymptotic inlet-to-outlet static pressure rise – maximum variation of 0.05 Pa within 2000 consecutive iterations.
- Asymptotic probe static pressure rise – maximum variation of 0.05 Pa within 2000 consecutive iterations.
- Asymptotic total-to-total pressure rise – maximum variation of 1 Pa within 2000 consecutive iterations.
- Asymptotic near-wake TKE – maximum variation of 30% within 2000 consecutive iterations.
- Asymptotic mid-wake TKE – maximum variation of 10% within 2000 consecutive iterations.
- Asymptotic far-wake TKE – maximum variation of 10% within 2000 consecutive iterations.

As shown in the example case of Figure 3.17, the probe static pressure rise and the inlet-to-outlet static pressure rise are expected to give values close to zero for all the conducted simulations, due to the utilised boundary conditions. Thus, although monitoring these parameters does not deliver meaningful information in terms of fan performance, it is useful for evaluating convergence, as they can serve as indicators of a well-resolved flow field that meets the established inlet and outlet boundary conditions.

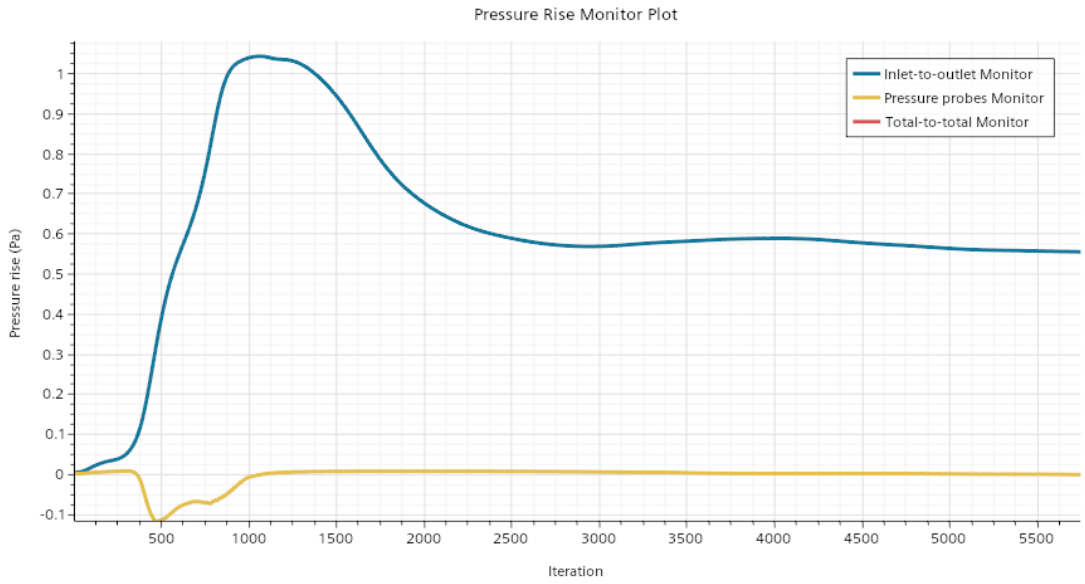

Figure 3.17: Sample monitoring plot showing solution convergence of the inlet-to-outlet static pressure rise and probe static pressure rise.

4
Analysis of results

Following the presentation of the proposed methodology, this chapter focuses on showing and analysing the obtained results. First, the outcome of the mesh independence studies is evaluated, and the selection of the final meshing setup is justified. The next section presents the simulation results across the defined RPM range computed by the final meshing and fully settled models, showcasing performance and acoustic comparisons between the two fan designs and the two turbulence setups. Finally, the last section compares the simulation data with the measurements from the experimental campaign carried out at VGTT's test rig for model validation.

Note that the plotted numerical results from both simulations and experimental measurements, such as flow rate, sound level, pressure rise, or torque, were mathematically normalised to keep the absolute values confidential, thereby complying with the agreed non-disclosure policies established by Volvo Cars and Volvo Group. For the sake of simplicity, this normalisation was applied by dividing all the values by the maximum one within each plot's dataset.

4.1 Mesh independence studies

Having established the physics setup on which the model is configured – namely, compressible flow with temperature effects and boundary conditions that induce a null static pressure rise – the definitive mesh parametrisation is now to be settled. This represents the final step of the modelling process, and the resulting mesh configuration will be considered specifically optimised for such conditions. Thus, the fully established models will be used for subsequent operational range analyses.

In the inflation mesh independence study, the sensitivity of the simulation outcome to the number of prism layers is assessed, while the impact of overall cell size is analysed in the wake refinement study. The values of Curle and Proudman sound levels, flow rate, pressure rise, and fan torque are monitored against the variation in the mesh setting to observe this sensitivity. In addition, the total iteration count is plotted to track the simulation cost and stability for each mesh configuration, allowing for an evaluation of computational efficiency.

4.1.1 Inflation mesh study

A simulation sweep spanning from 8 to 22 prism layers with fixed stretching factor was performed, evaluating the influence of inflation mesh extent. It was done only for SST k-ω turbulence, as justified in Section 3.4.2. The analysis was conducted for both fans, at low and high rotational speeds each, adapting the total inflation mesh thickness correspondingly. Graphical representations of the results for *Fan 1* model are shown in the plots of Figure 4.1 and for *Fan 2* in the plots of Figure 4.2:

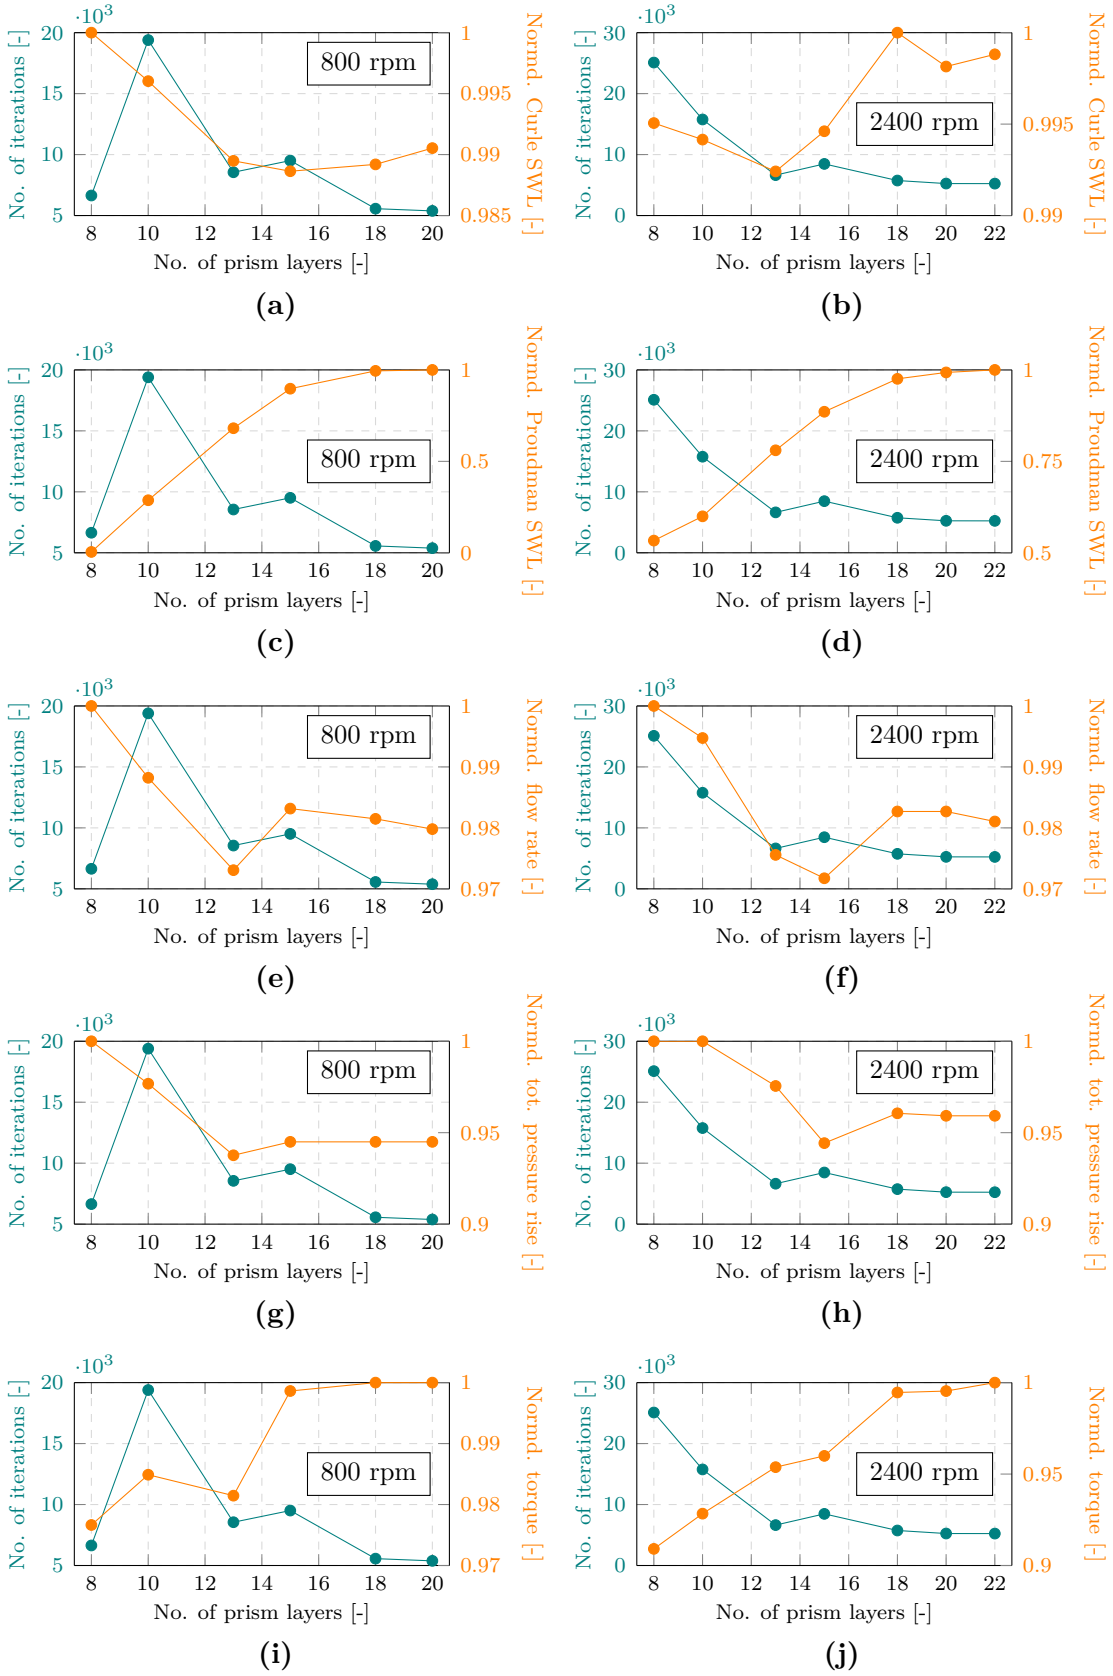

Figure 4.1: Sensitivity of acoustic and performance output to different inflation layer extents for *Fan 1*, together with the impact on solver iterations.

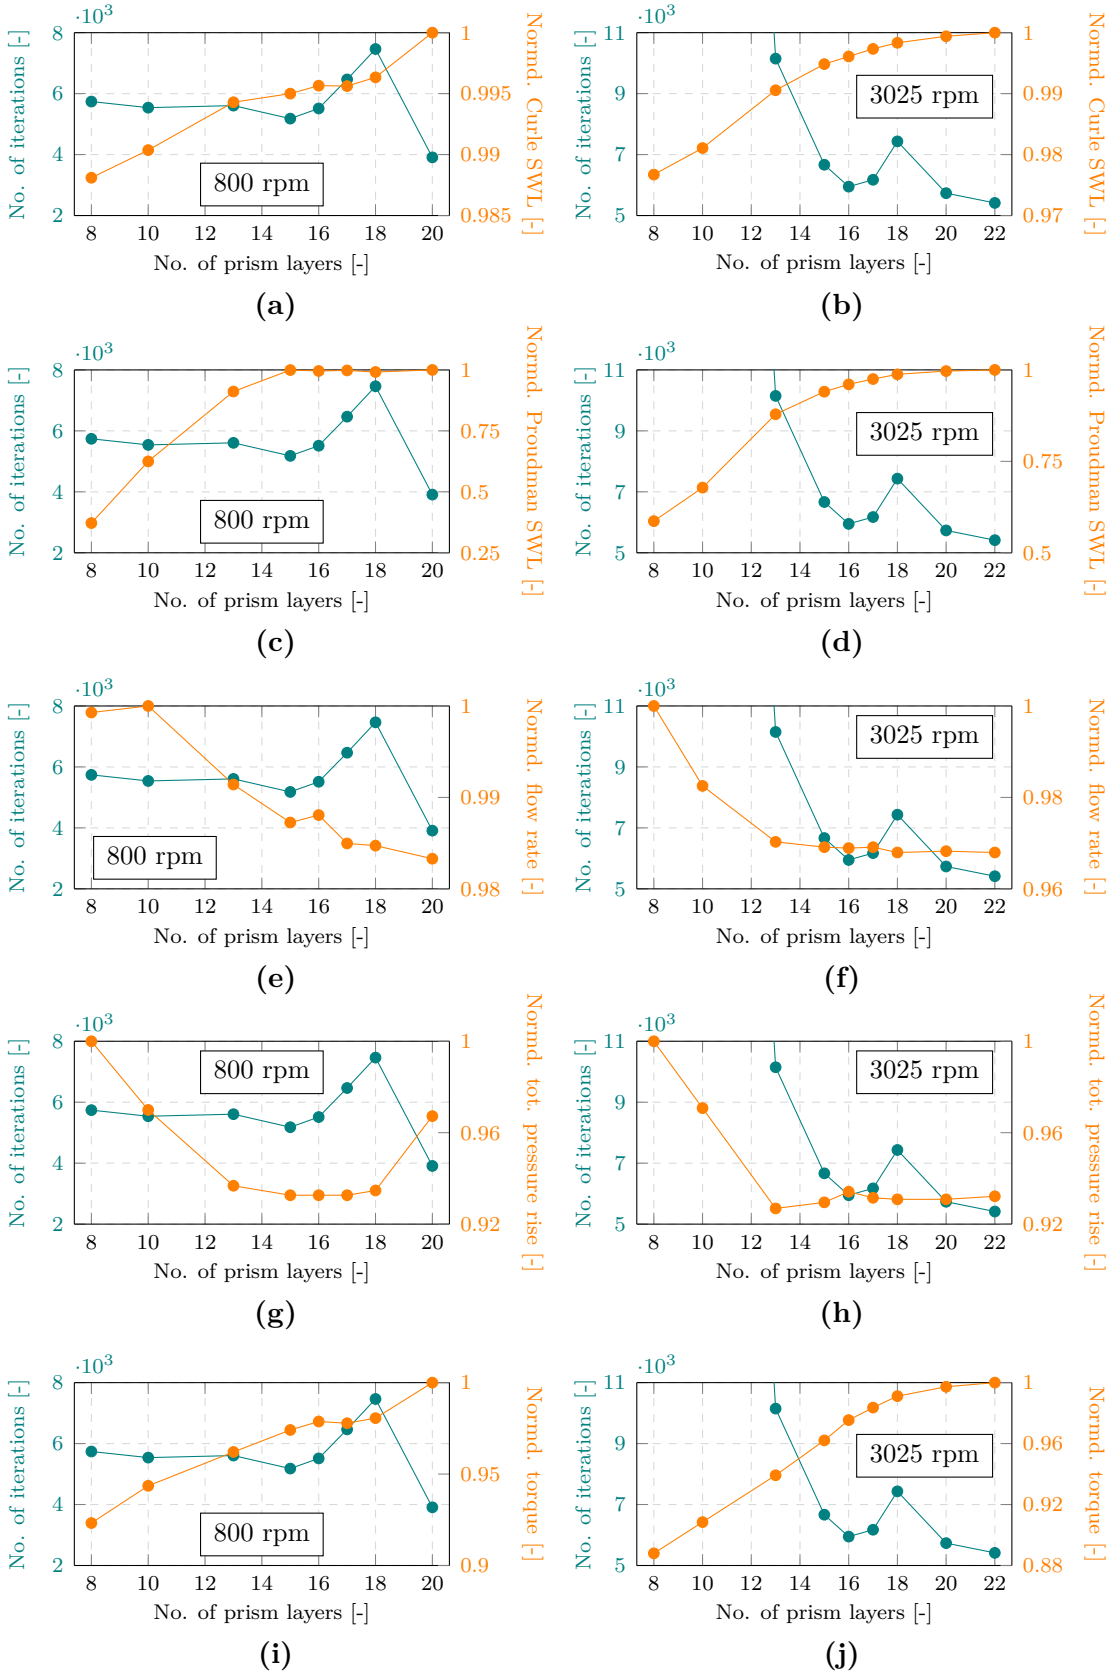

Figure 4.2: Sensitivity of acoustic and performance output to different inflation layer extents for *Fan 2*, together with the impact on solver iterations.

As observed, although some of the analysed variables, such as flow rate and Curle sound level, exhibit a maximum variation of only 0.5 to 3% for the evaluated layer adjustment range, others, like the total pressure rise and the fan torque, may present uncertainties of 5 to 10% if the inflation mesh is not sufficiently extended. This percentage is high enough to warrant further refinement, particularly in this case, where the variables of interest require fine-tuning of the model and its validation is a critical step. In fact, a particularly large variation is found in the Proudman sound level, indicating that with too few prism layers, only 0 to 50% of its contribution may be captured. Although the Proudman noise level does not play a dominant role in this study due to the low acoustic impact of quadrupole sources in the low-Mach-number regimes, properly resolving this magnitude can serve as a useful indicator of the model's convergence robustness, helping to reliably identify the precise computation of turbulent phenomena.

The objective of the mesh independence study is to determine the number of layers beyond which simulation results no longer exhibit significant sensitivity. In other words, the goal is to identify a point on the plotted orange curve where the dependent variable stabilizes. In *Fan 1*, this plateau begins between 13 and 15 prism layers for the low-speed case and around 18 for the higher-speed case, the latter being more conservative and thus a more usable choice. Moreover, beyond 18 layers, the number of iterations is almost halved, making this option particularly advantageous. In *Fan 2*, consistent results are observed between 15 and 18 layers, with 20 layers introducing a questionable shift. Apart from that, avoiding the peak in iteration count at 18 layers would save approximately 1,000 to 2,000 iterations, which makes 17 layers a compelling choice.

Thus, it can be said that, for flow conditions of null pressure buildup, 18 prism layers in the SST k-ω inflation mesh are optimal for *Fan 1* and 17 seem reasonable for *Fan 2*. These numbers of layers are therefore adopted and fixed for the simulation of the rotational speed sweep. For simulations with the k-ε model, a 3-prism-layer setup was established after observing an acceptable capture of results and no significant difference in computational expense compared to other numbers of layers.

4.1.2 Wake refinement study

In this case, the mesh study involves sweeping through 0.1, 0.05, 0.025 and 0.01 metres of target cell size in the wake, to evaluate the influence of its refinement level. The analysis was conducted for both fans at high rotational speed and for both turbulence models, setting the corresponding number of prism layers as established in the inflation mesh study. Graphical representations of the results for *Fan 1* are shown in the plots of Figure 4.3 and for *Fan 2* in the plots of Figure 4.4:

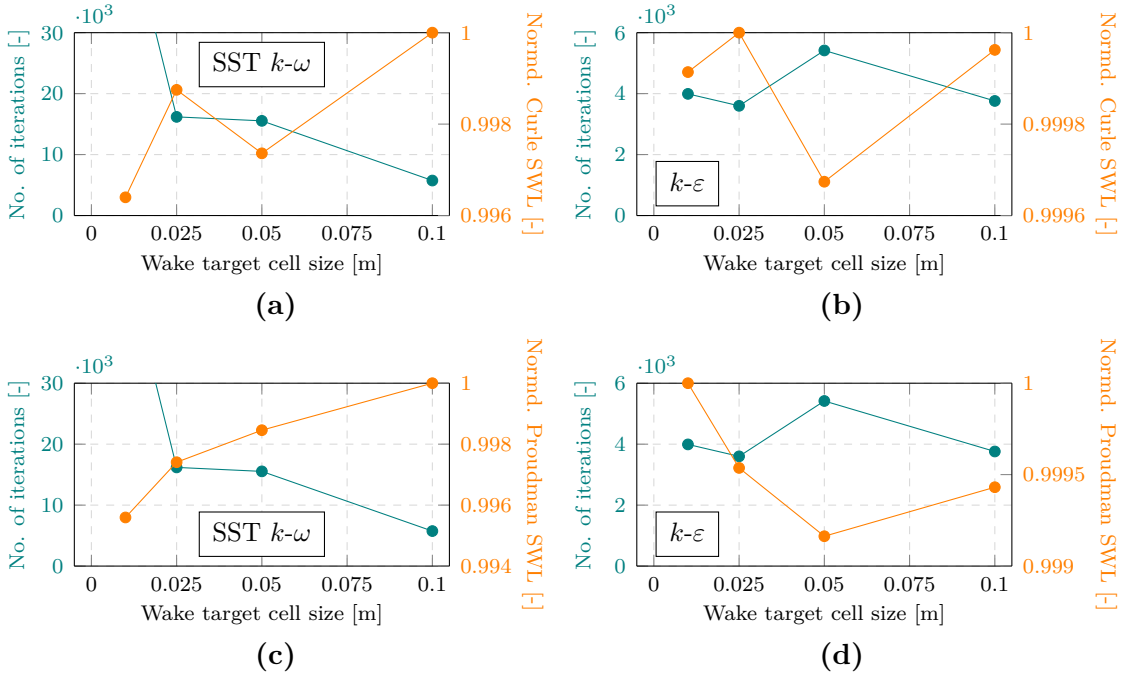

Figure 4.3: Sensitivity of the acoustic output to wake refinement for *Fan 1*, together with its impact on solver iterations.

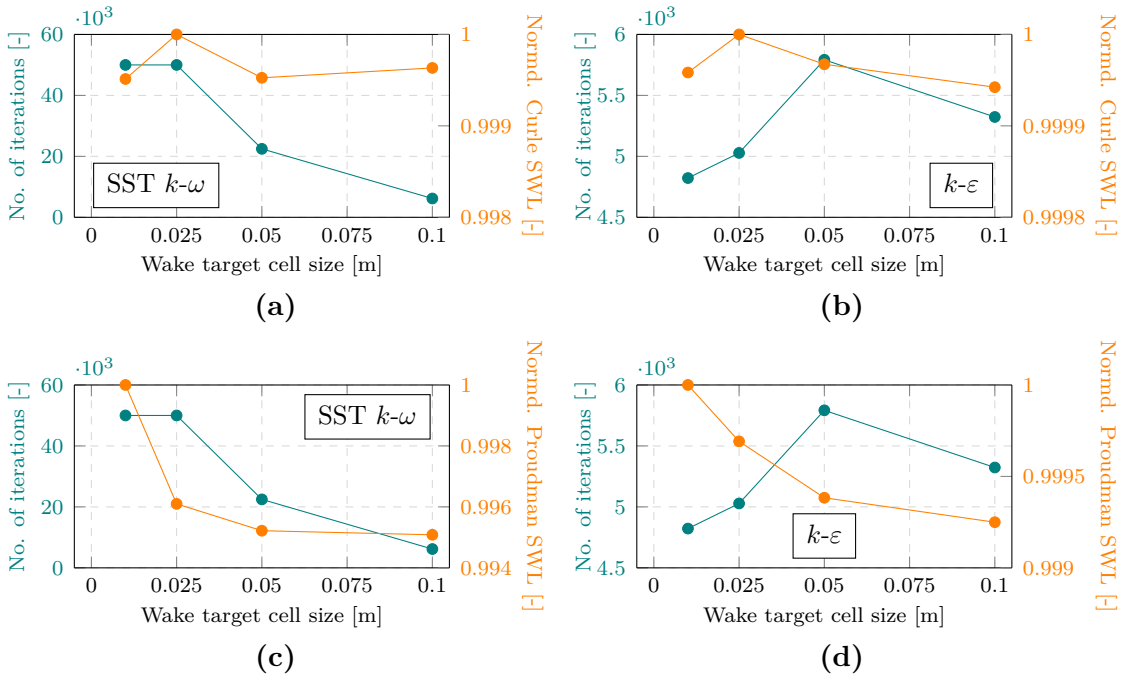

Figure 4.4: Sensitivity of the acoustic output to wake refinement for *Fan 2*, together with its impact on solver iterations.

As expected, there is no strong correlation between wake refinement and variations in Curle acoustic power for either the k-ε or SST k-ω models. This is because

the Curle model computes surface-generated acoustic power, whereas the present mesh refinement primarily affects the core volume. Instead, the expected correlation is observed for the Proudman sound level, as this model accounts for quadrupole-type sources originating from far-wall turbulent regions. However, as shown in Figures 4.3c, 4.3d, 4.4c, and 4.4d, an asymptotic trend in the Proudman level is observed for larger cell sizes rather than with finer adjustment.

More importantly, when simulations are executed in the SST k-ω scenario, the low-Reynolds-number approach makes it significantly more challenging to achieve a convergent solution trend as the wake region becomes more refined. This results in a substantially higher iteration count – up to one order of magnitude greater than that for the k-ε model. The increased computational effort is not justified for a difference in results of only 0.01 to 0.001%. In contrast, as expected from a turbulence formulation based on high Reynolds number, the k-ε model succeeds in delivering a solution in many fewer iterations compared to SST k-ω, as the refinement demand is concentrated in a high-y^+ region. However, the improvement in capturing acoustic power would remain too small to be worthwhile.

In view of all these observations, it seems that performing wake refinement does not provide the simulation results with any valuable insight for the present application, and the implementation of such a mesh independence study in the modelling methodology could therefore be avoided.

4.2 CFD prediction

After freezing the models with the aforementioned mesh settings, simulations across the range of operation are conducted for both fans. The results in terms of performance and acoustic power are analysed in this section.

4.2.1 Performance metrics

As noted previously, the monitored performance variables were volume flow rate, total pressure rise, fan torque, and total-to-total efficiency. The flow rate and the fan torque are easily calculated in STAR-CCM+ – the former is taken from the inlet boundary, and the latter is similarly calculated by selecting the surfaces to consider for the integration of normal and shear stresses and specifying the axis of rotation about which the moments are computed. The total-to-total pressure rise is calculated through the difference between the surface-averaged total pressures in virtual circular sections located just upstream and downstream of the fan. The efficiency is computed as the total-to-total pressure rise times the volume flow rate over the torque and the rotational speed.

In Figure 4.5, the obtained simulation results are plotted against fan speed, distinguishing the different fan variants and the two evaluated turbulence models. As expected, the flow rate as a function of rotational speed, shown in Figure 4.5a, follows a linear trend, resulting in a gain per 1000 RPM that is approximately 15%

higher for *Fan 1* compared to *Fan 2*. This difference in slope is reasonable given the larger diameter of *Fan 1*, which provides a greater flow-passing cross-sectional area. The lower number of blades and the higher angle of attack may also contribute to that. Additionally, these latter two characteristics significantly enhance *Fan 1*'s ability to generate a pressure rise. As shown in Figure 4.5b, both fans exhibit the typical quadratic relationship for total pressure rise, yet *Fan 1* delivers, on average, 60% higher values than *Fan 2*. Due to *Fan 1*'s overall superior capability to transfer momentum, it also implies a greater torque, as seen in Figure 4.5c. However, the amount of torque obtained, with a 40% higher than for *Fan 2*, is relatively low if we consider *Fan 1*'s remarkable superiority in total pressure generation, which indicates a more aerodynamically optimised blade design. As a result, *Fan 1* achieves a much higher total-to-total efficiency. As observed in Figure 4.5d, *Fan 1*'s efficiency reaches up to 40%, while *Fan 2* barely exceeds 25%, according to the SST k-ω model.

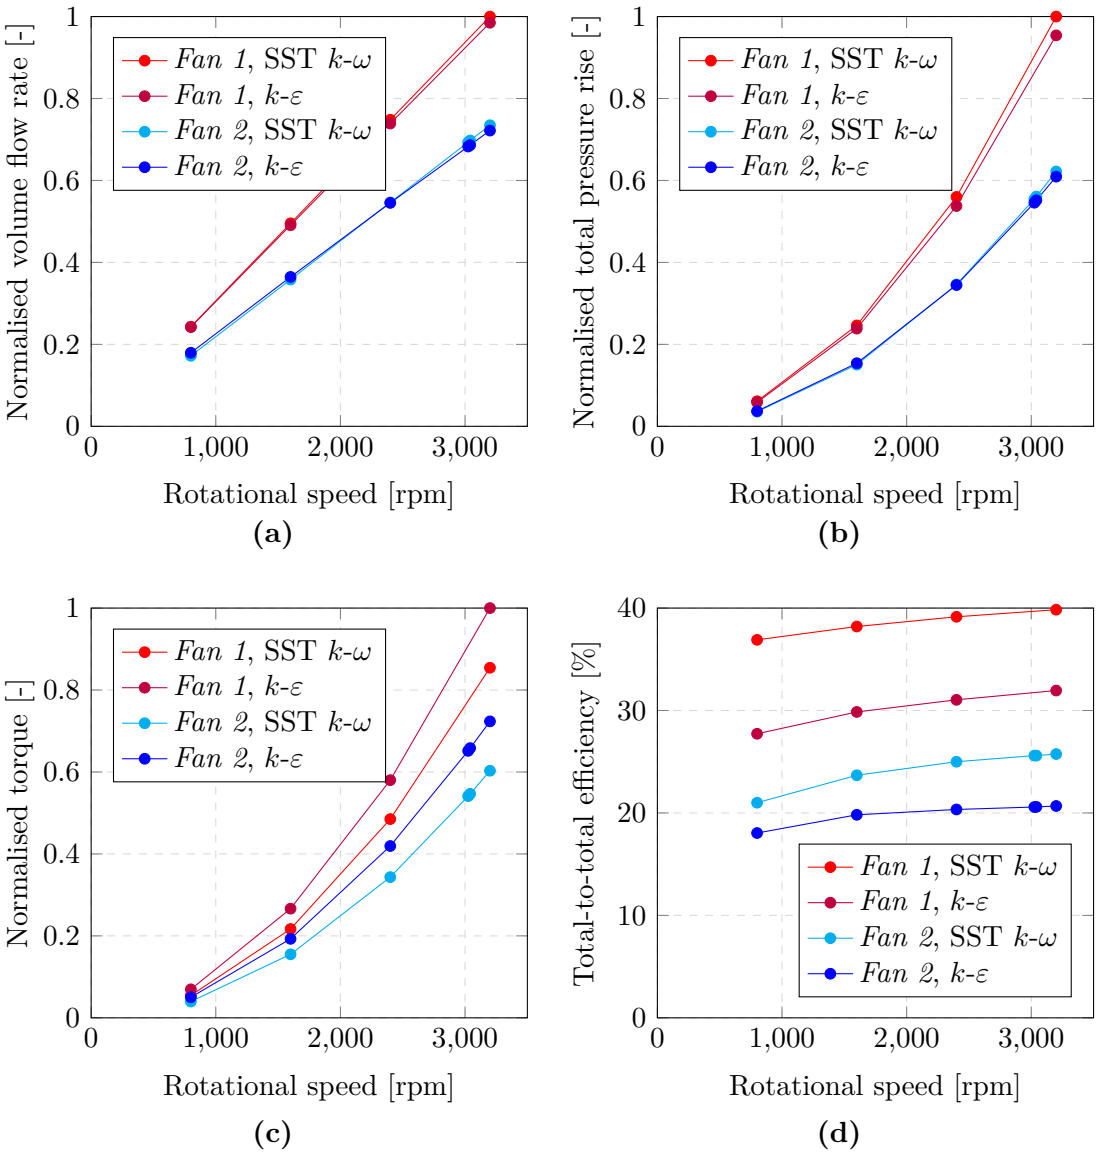

Figure 4.5: Simulation output of performance parameters for the speed working range, where the comparison between fans and turbulence models is observed.

Regarding the difference between the turbulence models, it can be seen in Figures 4.5a and 4.5b that the k-ε model will not significantly underperform SST k-ω in predicting flow rate and total pressure rise, as both yield nearly identical results in those aspects. The main difference lies in the torque computation: k-ε consistently predicts torque values about 20% higher than SST k-ω for both fans. This discrepancy is expected and can be attributed to the k-ε model's tendency to overpredict turbulent kinetic energy, particularly near walls and in regions of flow separation, both of which strongly influence the generated torque. As with the fan comparison, the torque disparity leads to an amplified mismatch when translated into efficiency. For a case study involving complex flow patterns, this observation partly justifies the preference for SST k-ω as the primary turbulence model. However, experimental torque measurements would ideally be required to validate that.

4.2.2 Acoustic output

To analyse the broadband aeroacoustic output, values of both Curle and Proudman SWL components are presented. Note that a direct comparison between Curle and Proudman levels is not meaningful, as the difference between their respective contributions in terms of acoustic power is enormous. Therefore, the purpose of Figure 4.6's layout is solely to compare the two studied fans and turbulence models.

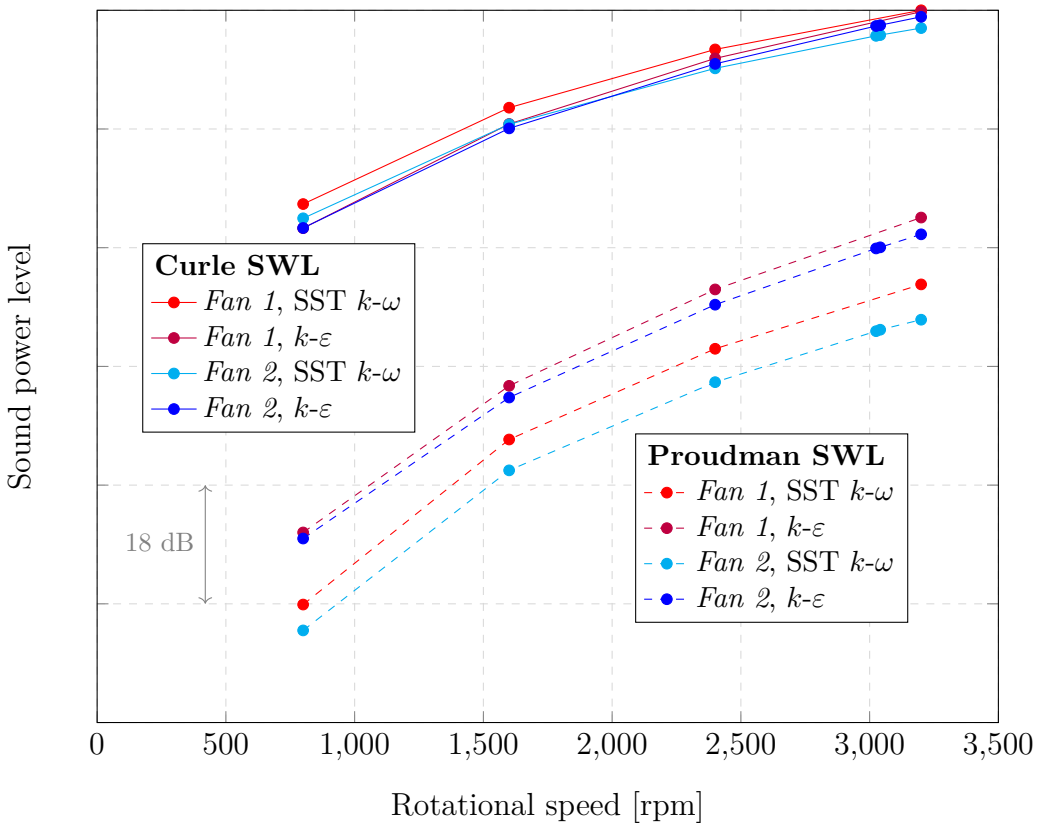

Figure 4.6: Sound level broken down into Curle and Proudman sources along the working range. Both fan specimens and turbulence models are compared.

Despite the significant difference in aerodynamic performance, *Fan 1* is only slightly noisier – by around 3.5–4% across the range – compared to *Fan 2*, based on the results of the source models. These simulations do not capture sound generated by blade volume displacement, and thus represent only a partial estimate of the total acoustic load. The slightly higher broadband noise observed in *Fan 1* is mostly attributed to more prevalent boundary-layer and separation turbulence, given its larger wetted surface area, which involves increased wall friction, and the more aggressive blade angles of attack, which promote separation regions. In fact, flow separation is more likely in *Fan 1* due to the wider pitch spacing between blades. Conversely, *Fan 2* features milder blade angles and narrower pitch, which generally favour flow attachment and, consequently, lower turbulence-induced noise.

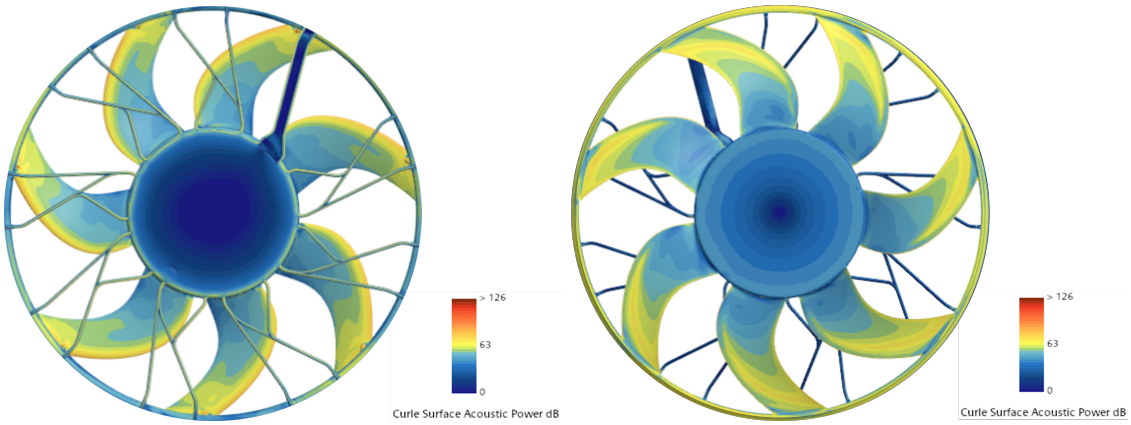

(a) **(b)**

Figure 4.7: Field of Curle acoustic production over *Fan 1*'s front (a) and back (b) surfaces at 800 RPM.

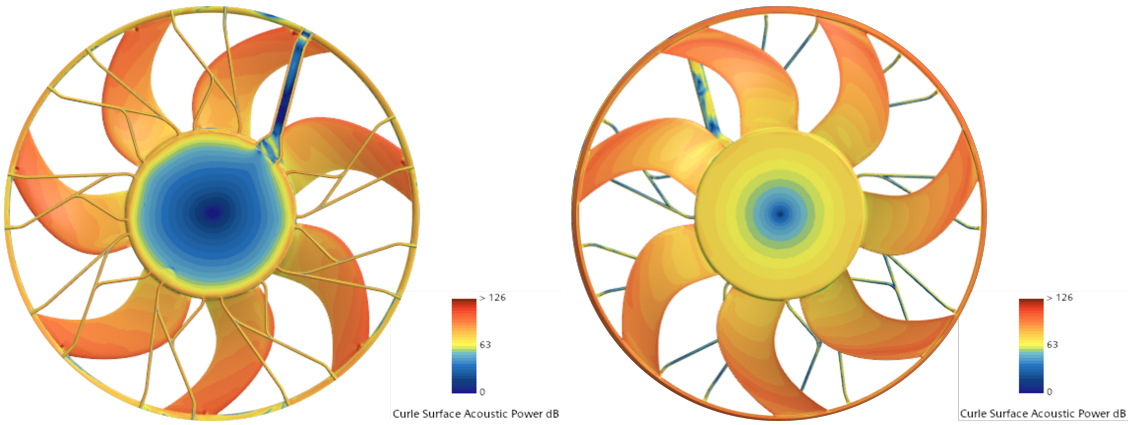

(a) **(b)**

Figure 4.8: Field of Curle acoustic production over *Fan 1*'s front (a) and back (b) surfaces at 3200 RPM.

The SST k-ω model successfully captures a difference in Curle levels between the two fan designs, whereas the k-ε model barely shows sensitivity, as observed in

Figure 4.6. Only a difference in Proudman levels is noticed, where *Fan 1* values are slightly higher than those of *Fan 2*, and with k-ε expectedly exceeding k-ω. Even so, this contribution remains mostly negligible for the current application. In view of the limitation of k-ε, results only from the SST k-ω simulations were used to extract further aeroacoustic analysis of the Curle-level source distribution via surface field visualisation. This is illustrated in Figures 4.7 to 4.10.

A critical region for the aeroacoustic analysis of airfoils is the tip of the wing or blade, where the pressure difference between the suction and pressure sides gives rise to strong vortices. Phenomena like this become especially important in turbomachinery applications such as the present fans, where continuous blade–vortex interaction caused by the spinning of the rotor leads to significant acoustic fluctuations. This usually represents an additional contributor not only to narrowband noise but also to part of the broadband content. Although vortical structures cannot be directly reconstructed in RANS simulations, their presence can be inferred from the modelled turbulence fields. This is observed for *Fan 1* in Figures 4.7 and 4.8. Especially in the former, high acoustic emission is concentrated not only at the leading edges, where stagnation of turbulent flow occurs, but also along the blade's mid to high span. As demonstrated by Park, Lee, and Lee, this behaviour is reinforced by a backward-sweep layout [18], which introduces a radial component to the flow, effectively influencing its direction and intensifying the blade loading near the tip. As a result, a distinct gradient of noise production in that direction is visible on *Fan 1*'s blades, showing a relevant impact also on the pressure surface, as clearly seen in the back view provided in Figure 4.7b.

By contrast, Figures 4.9 and 4.10 demonstrate that the aforementioned effect is notably diminished for *Fan 2*. The forward sweep at high spans counteracts the tipward flow tendency seen in *Fan 1*, partially mitigating the outwash stream and limiting broadband noise generation mainly to the leading edges. Another notable observation is that, despite their streamlined geometry, the spokes of this fan make a noticeable acoustic impact at 3200 RPM.

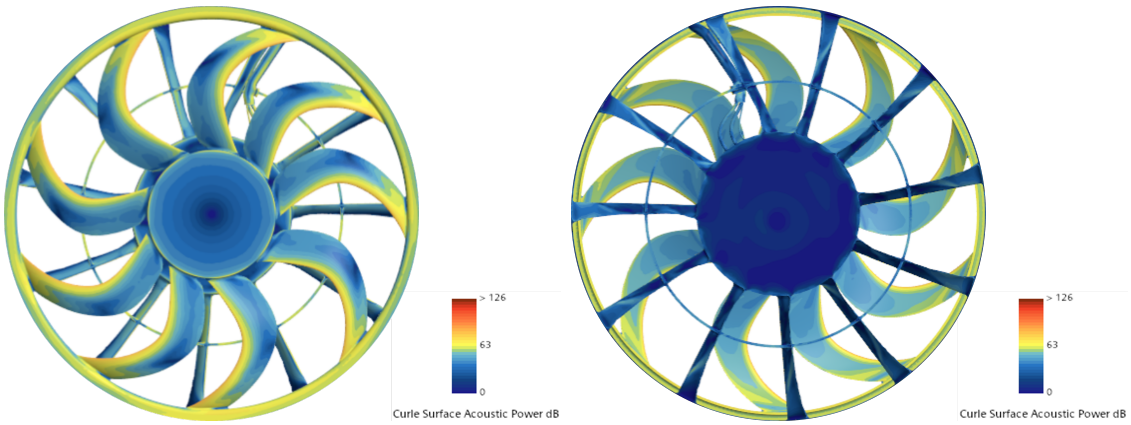

(a) **(b)**

Figure 4.9: Field of Curle acoustic production over *Fan 2*'s front (a) and back (b) surfaces at 800 RPM.

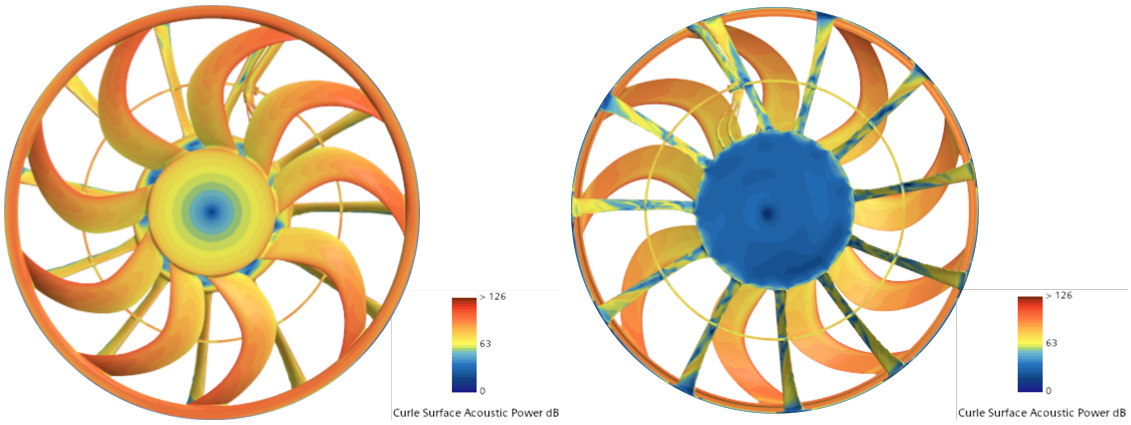

(a) **(b)**

Figure 4.10: Field of Curle acoustic production over *Fan 2*'s front (a) and back (b) surfaces at 3200 RPM.

To further investigate the spokes' acoustic footprint, an isolated integration of the Curle acoustic power was extracted only involving the stator surfaces. Although this cannot be validated by experimental data, an analysis approach of the comparative estimate is presented. As anticipated above, Figure 4.11 confirms that the sound level contribution from the spokes in relation to the total is greater in *Fan 2* than in *Fan 1*. These proportions remain relatively consistent across the rotational speed range, and are derived from the SST k-ω results, as the k-ε model fails to capture any significant distinction between the designs.

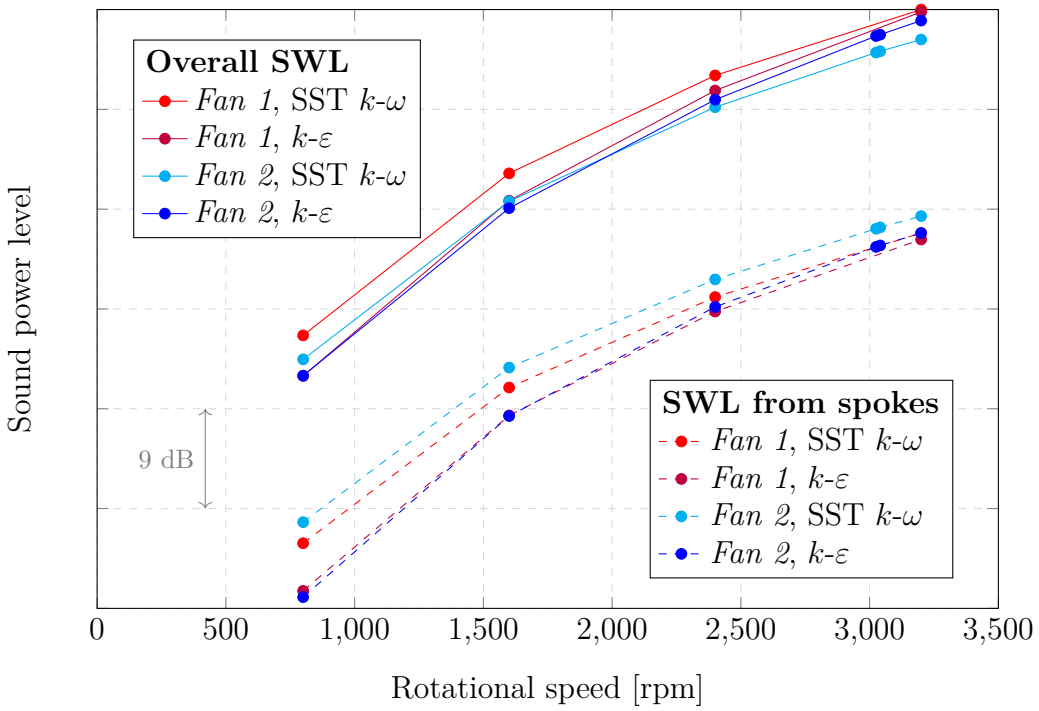

Figure 4.11: Comparison between overall and spoke-irradiated sound level for both fans and turbulence models along the speed range.

Figure 4.12 plots the percentage of total acoustic power originating from the spokes. Consistent with earlier results, the stator layout in *Fan 2* definitely makes a larger aeroacoustic impact than that of *Fan 1*. Besides, this contribution decreases with increasing rotational speed, underscoring the role of flow impingement into rotating parts as a key driver of noise generation in the featured scenarios. Again, this information can only be strictly stated from the SST k-ω model, as the k-ε model does not yield interpretable patterns across designs or operating conditions.

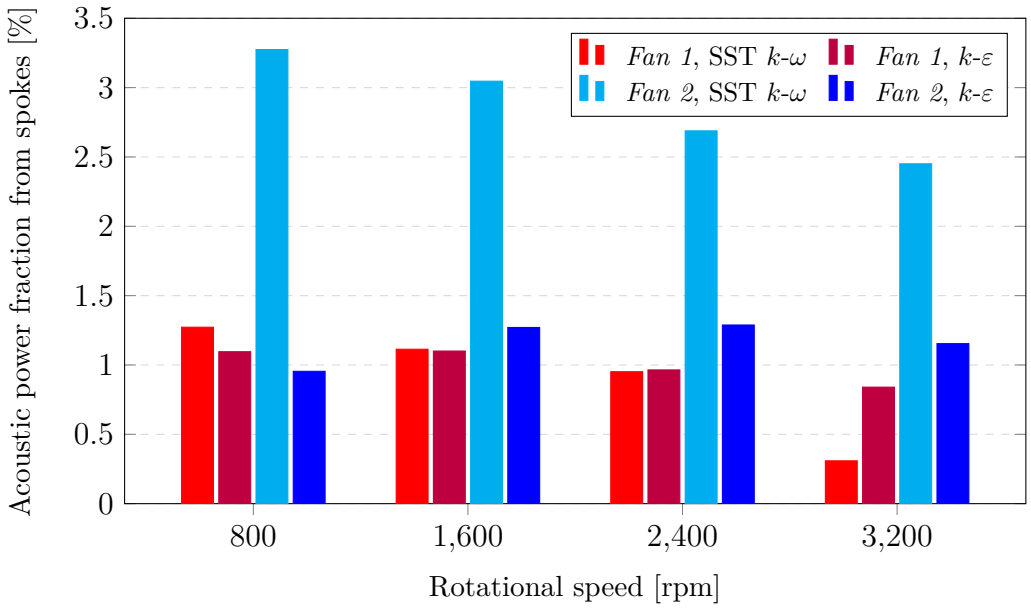

Figure 4.12: Percentages of acoustic power produced by the spokes out of the overall amount. Data for the two fans and turbulence models classified by speeds.

4.3 Experimental validation

Model validation in terms of volume flow rate and sound power level was conducted through the experimental measurements taken in Volvo GTT's fan test rig. The rig features a closed-loop airflow circuit, where the stream is driven by a master fan and regulated via a set of control nozzles. These components are located beyond the test chambers and are adjusted to achieve the desired flow rate and pressure rise across the test section. The procedure to set working points for their measurement consisted of iteratively driving the tested fan while tuning the flow delivery, until achieving a combination of measured RPM and volume flow rate that produced a null pressure jump across the pressure probes. At each measurement point, proper stabilisation of the monitored variables was aimed, in order to capture a consistent and reliable trend across the simulated operating range.

Noise data were collected through acoustic pressure measurements using a hemispherical array of microphones positioned at a fixed radius upstream of the test section, oriented toward the fan. The raw data were processed to account for the distance from the source, allowing for the estimation of power-based sound levels.

4.3.1 Volumetric flow rate

Figure 4.13 shows that the simulation results for *Fan 1* closely replicate the measured flow rate trend with respect to rotational speed. However, a slight mismatch in the slope is observed, corresponding to an increasing underprediction of the flow rate that amounts to approximately 2.5% per 1000 RPM. A similar behaviour is observed for *Fan 2*, though with a larger deviation that increases by 14% per 1000 RPM, resulting in an absolute error of around 28% at 1200 RPM. These discrepancies are considered acceptable in light of the achieved computational efficiency. They stem from the cumulative effect of various sources: potentially overly rigid boundary conditions, the approximations of the turbulence modelling, and uncertainty inherent to the measurement instruments. Nevertheless, the most significant error contributor might be that associated with the MRF rotation, a modelling approach that simplifies the actual mechanism by which the fan imparts momentum to the flow. Compared to solid-body motion, the MRF method introduces limitations in precisely capturing the real dynamic interaction between the blades and the fluid, leading to inaccuracies in the resulting flow field. In fact, as observed in this case, the MRF method can lead to slightly different modelling mismatches between fans. This is due to the fact that the two impellers at hand possess significantly different design features related to blade loading. These features, such as number of blades, angle of attack, or camber, strongly determine the distribution and propagation of injected momentum over the volume swept by the rotor, which is precisely the main aspect affected by the MRF simplification. Any further, more detailed study based on rigorous flow field analysis is left for future stages of this research front.

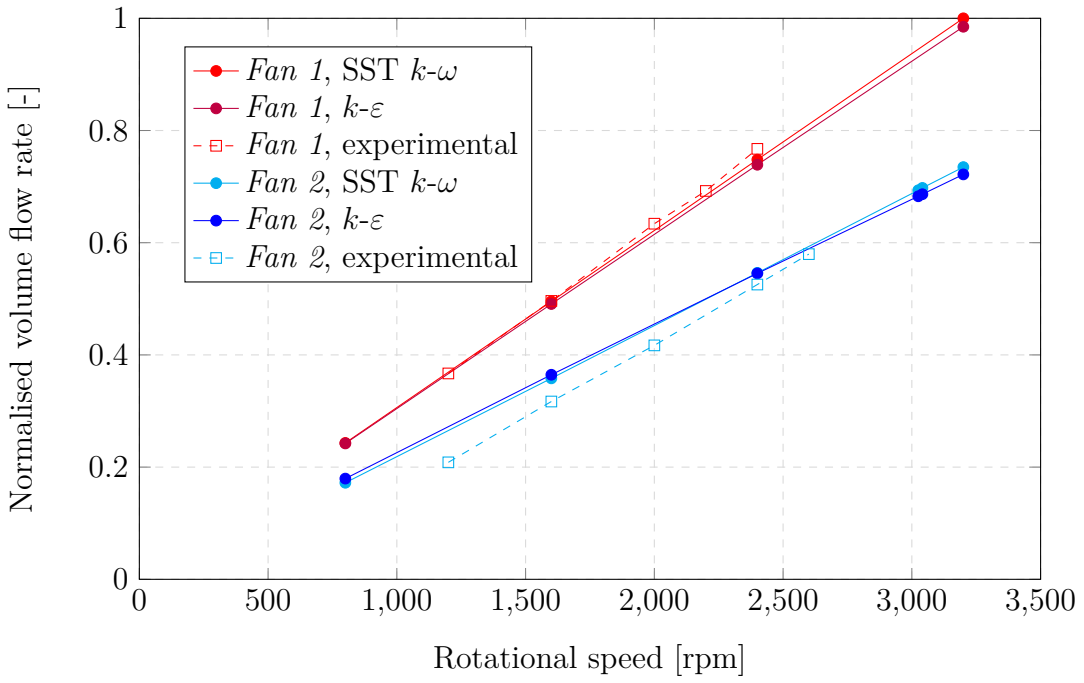

Figure 4.13: Comparison of delivered volume flow rate between simulation and experimental data along the observed working range for both fan specimens.

4.3.2 Sound power level

At first glance, it is evident from Figure 4.14 that while the overall trends in sound power level are reasonably well predicted, a significant offset is present, especially for *Fan 1* at increased speeds. This underprediction was indeed expected, given that the adopted steady-state modelling approach does not capture pressure field fluctuations caused by the blade-passing events. In this way, only the broadband sound power contribution arising from the modelled turbulence fields is accounted for.

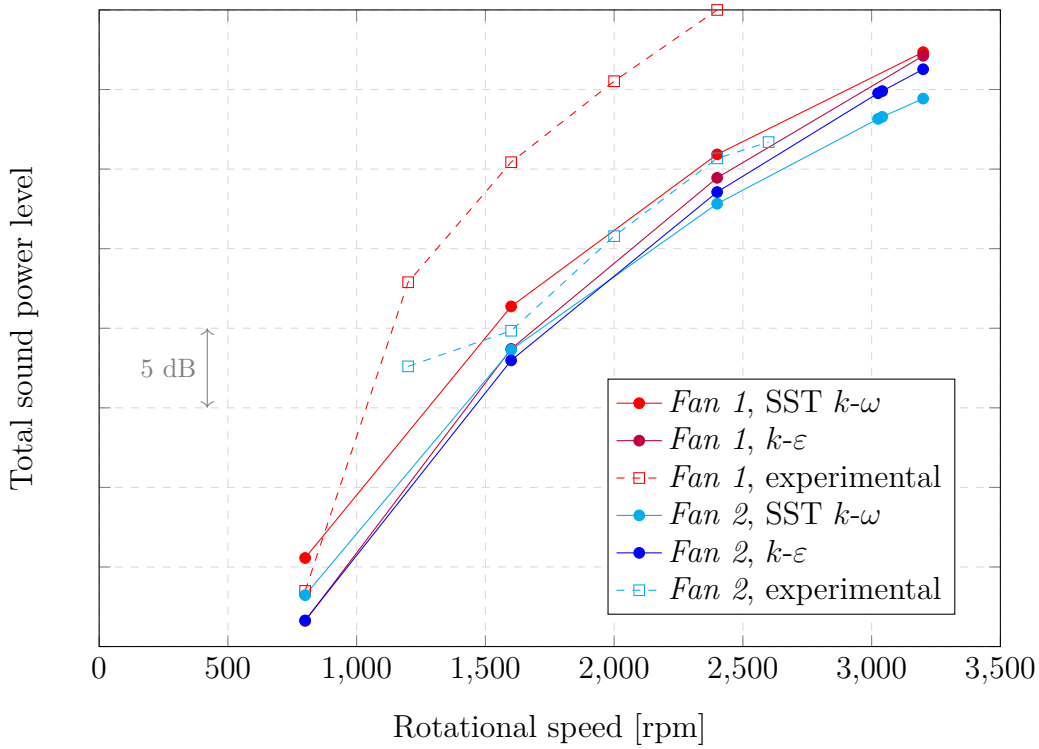

Figure 4.14: Comparison of acoustic levels between the simulation and the experimental data along the observed working range for both fan specimens.

For *Fan 1*, a noticeable underprediction of approximately 10% emerges from 1000 RPM onward. At 800 RPM, however, the error drops to around 4%, suggesting a more accurate representation at lower speeds. The sudden increase in measured noise at higher speeds is likely due to the onset of strongly unsteady phenomena, such as the formation of large-scale vortical structures, which are not resolved in a steady-state framework. Furthermore, no clear preference between the SST k-ω and k-ε turbulence models was adopted, as the results for both exhibited similar levels of mismatch in SWL predictions.

On the other hand, *Fan 2* shows a more moderate underprediction of 2–2.5% that additionally appears to be more uniform, with only a slight deviation from the trend at 1200 RPM. Moreover, the offset of the two turbulence models with respect to the validation data is similar on average, but SST k-ω mimics more accurately

the trend of the experimental curve.

The experimental results indicate then that *Fan 1* is actually around 12% noisier than *Fan 2* across most of the operating range, in contrast to the smaller difference suggested by the simulations. The speed-sensitive discrepancy between model and measurements, along with the analysis of the computed flow fields, reveals an important insight into the relationship between the geometric characteristics of the fans and the nature of the noise content they provide: assuming that the simulation's underprediction for *Fan 1* is primarily due to undetected tonal or narrowband noise components, the significantly smaller discrepancy observed for *Fan 2* suggests that its acoustic improvement largely stems from the suppression of such components. With this, beyond developing models that reasonably well predict broadband noise, the investigation has qualitatively identified that the backward sweep angle applied to the blades generates a substantial portion of the noise associated with large-scale turbulent interactions, while the forward sweep helps to mitigate it.

5

Conclusions and discussion

The key conclusions drawn regarding the RANS-based modelling methodology and its validation are outlined below:

- The mesh independence study demonstrated that an adequate number of prism layers improved simulation certainty and convergence, avoiding an information loss of 1–2.5% in SWL output with a reasonable increase in cell count. For the present application and operating regime, 18 and 17 prism layers under the SST k-ω model were found to be optimal for *Fan 1* and *Fan 2*, respectively.

- Wake refinement was found to be computationally expensive and ultimately unjustified for the current models. It had minimal impact on Curle and Proudman level predictions and significantly increased iteration count, particularly for the SST k-ω model. While potentially useful for setups involving downstream components or higher pressure rise scenarios, it was not beneficial for the present configuration.

- Although Proudman noise contribution is negligible in low-Mach regimes, its sensitivity to changes in the inflation mesh extent proves it as a good indicator of boundary layer robustness and solution stability.

- Model predictions consistently agree with experimental measurements that *Fan 1* outperforms *Fan 2* in terms of flow rate and pressure rise, a difference primarily attributed to its higher blade angle of attack and slightly larger diameter. Only a minor overpredicting tendency for low RPM is observed in *Fan 2*'s flow rate, likely caused by limitations of the MRF approach or slight discrepancies in the outlet boundary condition setup.

- The models provide a generally good prediction of sound level trends across the speed range for broadband noise comparison between the fans. However, a consistent underprediction is observed, particularly for *Fan 1* beyond 1000 RPM. Apart from containing the blade-passing tonal component, this mismatch is partially attributed to narrowband noise components generated by large-scale turbulent phenomena – such as blade–vortex interactions – that are not captured by the steady RANS modelling approach. Flow field analysis supports this interpretation, revealing high-span patterns that associate the excess noise measured for *Fan 1* with its backward blade sweep, and *Fan 2*'s quieter behaviour with its forward sweep. This reflects that the models are sufficiently sensitive to capture geometry-driven acoustic differences.

- Simulations indicate that the spokes in *Fan 2* contribute a greater share of the

total sound power compared to those in *Fan 1*, despite their more aerodynamically streamlined design. This difference stems from their position downstream of the impeller, where interaction with turbulent wake structures leads to increased noise generation.

- The observations outlined above apply primarily to the SST k-omega turbulence model, which outperformed k-ε in this study. While the latter produced very similar results in terms of flow rate and pressure rise with lower computational cost, it showed limited acoustic output sensitivity to differences in fan geometry and spoke configuration, making it less reliable for comparative design assessment.

Thus, the modelling study has yielded valuable insight into the noise prediction capability of RANS simulations for different fan layouts. In summary, while some limitations were observed due to the steady-state nature of the models, especially in capturing transient acoustic phenomena, the results support the proposed methodology for comparative acoustic analysis in early-stage design.

Bibliography

[1] International Energy Agency, "Global EV Outlook 2024", April 2024. Available online: https://www.iea.org/reports/global-ev-outlook-2024

[2] Eurostat, "House or flat: where do you live?", May 2021. Available online: https://ec.europa.eu/eurostat/web/products-eurostat-news/-/ddn-20210521-1

[3] World Health Organization, "Environmental Noise Guidelines for the European Region" (2018). Available online: https://www.who.int/europe/publications/i/item/9789289053563

[4] W. Jung, S. Song, M. Park, D. Lee, "The noise prediction of automotive axial fan with different blade sweep angle using the unsteady CFD analysis", *Fan 2015: International Conference on Fan Noise, Aerodynamics, Applications and Systems*, Lyon (France), April 2015, p. 5. Available online: https://fan2025.org/archives/fan2015/papers/fp-pdf-8-JUNG.pdf

[5] M. Granlöf, "Shape optimization of axial cooling fan via 3D CFD simulation and surrogate modeling", Master's Thesis from KTH Royal Institute of Technology in collaboration with Volvo Car Group, 2021. Report No. 2021:354.

[6] D. Ghosh, "Aerodynamic Design and Installation Effects of Automotive Electric Cooling Fans", Licentiate Thesis from Chalmers University of Technology in collaboration with Volvo Car Group and Volvo Group Trucks Technology, 2023. Report No. 2023:08.

[7] F. Zenger, C. Junger, M. Kaltenbacher, S. Becker, "A Benchmark Case for Aerodynamics and Aeroacoustics of a Low Pressure Axial Fan", SAE Technical Paper 2016-01-1805, 2016. Available online: https://doi.org/10.4271/2016-01-1805.

[8] D. Ghosh, N. Andersson, S. Etemad, "Computational Aeroacoustics of Inlet Geometry on Tip Noise for Low Pressure Axial Fans", AIAA SciTech Forum and Exposition, 2024, AIAA 2025-2166. Available online: https://doi.org/10.2514/6.2025-2166.

[9] M. Vourakis, E. Zea, M. Karlsson et al, "Installation Effects on Axial Fans: Combined Aeroacoustic and Psychoacoustic Perspective", Applied Acoustics journal Vol. 240 (110872), 2025. Available online: http://dx.doi.org/10.1016/j.apacoust.2025.110872.

[10] D. Ghosh, M. Vourakis, N. Andersson, S. Etemad, "Computational Aeroacoustics of Low-Pressure Axial Fans Installed in Parallel", Journal of

Fluids Engineering Vol. 147 Issue 2 (021204), 2025. Available online: https://doi.org/10.1115/1.4066752.

[11] P. Gullberg, L. Löfdahl, S. Adelman, P. Nilsson, "An Investigation and Correction Method of Stationary Fan CFD MRF Simulations", SAE Technical Paper 2009-01-3067, 2009. Available online: https://doi.org/10.4271/2009-01-3067.

[12] T. von Kármán, "Mechanical similarity and turbulence", NACA Technical Memorandum, 1930.

[13] B. E. Launder, D. B. Spalding, "The numerical computation of turbulent flows", Computer Methods in Applied Mechanics and Engineering, 1974, p. 269-289.

[14] F. R. Menter, "Two-equation eddy-viscosity turbulence models for engineering applications", AIAA Journal Vol. 32 Issue 8, 1994. Available online: https://doi.org/10.2514/3.12149.

[15] M. J. Lighthill, "On sound generated aerodynamically I. General theory", Proceedings of the Royal Society of London, Series A, Mathematical and Physical Sciences Vol. 211 Issue 1107, 1952. Available online: https://doi.org/10.1098/rspa.1952.0060.

[16] N. Curle, "The influence of solid boundaries upon aerodynamic sound", Proceedings of the Royal Society of London, Series A, Mathematical and Physical Sciences Vol. 231 Issue 1187, 1955. DOI: 10.1098/rspa.1955.0191

[17] I. Proudman, "The generation of noise by isotropic turbulence", Proceedings of the Royal Society of London, Series A, Mathematical and Physical Sciences Vol. 214 Issue 1116, 1952. DOI: 10.1098/rspa.1952.0060

[18] M. Park, D. Lee, H. Lee, "Experimental and computational investigation of the effect of blade sweep on acoustic characteristics of axial fan", Publication in Applied Acoustics journal Vol. 189 (108613), 2022. Available online: https://doi.org/10.1016/j.apacoust.2021.108613.

[19] O. Fares, “Numerical Investigation of Noise Generation by Automotive Cooling Fans”, Master’s Thesis from Chalmers University of Technology, 2019. Report No. 2019:90.

[20] Siemens AG, Simcenter STAR-CCM+ Documentation, Version 2402.

[21] A. G. Pietroniro, R. Kabral, Z. Huang, J. Backman et al., "HVAC Blower: A Steady State RANS Noise Prediction Method", SAE Technical Paper 2024-01-2937, 2024. DOI: 10.4271/2024-01-2937.

[22] A. G. Pietroniro, E. Trigell, S. Jacob, M. Mihaescu et al., "Effects of Boundary Layer and Local Volumetric Cells Refinements on Compressor Direct Noise Computation", SAE Technical Paper 2022-01-0934, 2022. DOI: 10.4271/2022-01-0934.

[23] M. Åbom, "An Introduction to Flow Acoustics" (4th edition), 2010, eBook for course lectures at KTH Royal Institute of Technology, ISBN: ISRN/KTH/AVE/N-06/04-SE. Section 2.41: "Low Mach number jets". Available online: https://www.researchgate.net/publication/296691390

DEPARTMENT OF MECHANICS AND MARITIME SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY