DELFT UNIVERSITY OF TECHNOLOGY

MASTER THESIS (AE5122) - AEROSPACE ENGINEERING

Error and Uncertainty Quantification of a Stochastic Noise Generation and Radiation Method for Predicting Flow-Induced Vibrations from RANS Simulations in Lithography Machines

---

by

Diego Alonso (6071023)

to obtain the degree of Master of Science at TU Delft,
to be defended publicly on

Committee Members:

Prof. Dr. Casalino	- TU Delft	- Chair
Dr. ir. Hulshoff	– TU Delft	- External Examiner
Dr. ir. van Zuijlen	– TU Delft	- Supervisor
Dr. Rezaeiha	– ASML	- Supervisor
Dr. Waterson	– ASML	- Supervisor

Preface

I am very grateful to my parents, whose care and support made it possible for me to move to the Netherlands and to enjoy my life here. I also thank my brother Mario for being as strong and charming as he is.
I am sincerely thankful to my supervisors, Dr. Alexander van Zuijlen at TU Delft, and Dr. Nicholas Waterson and Dr. Rahim Rezaeiha at ASML, for their invaluable guidance and support throughout this study. I also extend my thanks to Gokulakrishnan for introducing me to the Flow Modelling group at ASML and for helping me find accommodation during my stay.

I am truly grateful for the friendships I have built during my internships at ASML. Special thanks to Filip Karaba, Pablo Ruiz, Juan Bedoya, Jeanne Mazen, Yifei, Ben Huang, Kai, Abdul, Paul, Stijn, and many others who made the experience memorable.

I also wish to thank my friends in Delft who made the Netherlands feel like home. In particular, Camillo Podio, Daniel Soler, Pablo Álvarez, Álvaro Lanza, Lorenzo Belli, Sascha, Timo, and everyone I shared a beer with at Bouwpub — you all made this journey unforgettable.

Having studied in seven universities across Europe (including in Spain, Poland, Italy, the Netherlands, Germany, Belgium and Turkey), I have gained a broad academic perspective and a deep appreciation for multicultural environments. These experiences have shaped both my personal and professional growth. With this thesis, I conclude my studies and proudly continue my journey as an engineer at ASML in the Technical Support department.

Diego Alonso
Veldhoven, The Netherlands

Abstract

Accurately predicting flow-induced vibrations (FIV) in high-precision lithography equipment is challenging, as standard Reynolds-Averaged Navier-Stokes (RANS) simulations fail to capture the driving turbulent fluctuations. This work evaluates the accuracy and uncertainty of a Stochastic Noise Generation and Radiation (SNGR) pipeline that reconstructs time-resolved, divergence-free velocity fluctuation fields from RANS statistics to improve in FIV and aero-acoustic assessments. Therefore, helping to ensure nanometer-scale manufacturing precision. The MATLAB based implementation ingests CGNS/HDF5 solver output, assembles modal Fourier fields from a prescribed energy spectrum, enforces incompressibility, applies anisotropic tensor mapping to match Reynolds stresses, and offers both a time-marching (single time-loop) and an ensemble snapshot mode. Validation is performed against canonical Direct Numerical Simulation (DNS) reference data for turbulent channel flow (Lee & Moser, up to $Re_\tau \approx 5200$) and a benchmark backward-facing step, and comparisons are made with RANS (StarCCM+) results. A modular MATLAB pipeline automates diagnostics and produces slice-wise statistics (RMSE, absolute/relative errors), spatial heatmaps, PSDs, and GIF visualizations. Results show that SNGR successfully reproduces spectral content and the spatial distribution of turbulence, yielding instantaneous fields that correlate with RANS TKE maps. However, the analysis reveals that discrepancies - particularly concentrated in near-wall amplitudes, outer-layer overshoots and the smallest resolved scales - are primarily inherited from biases in the initial RANS model. The report quantifies how these RANS biases, domain and boundary conditions choices and numeric resolution propagate into SNGR reconstructions, and it recommends practical diagnostics and sensitivity checks for the robust industrial application of the SNGR method.

Contents

List of Figures

List of Symbols

Table 1: List of Symbols

Symbol	Description
Latin Symbols	
a	Temporal correlation factor
C	Constant in the logarithmic law-of-the-wall
C_f	Skin friction coefficient
$E(\kappa)$	One-dimensional energy spectrum
ER	Expansion Ratio
f_i	Body force per unit volume (N/m^3)
f_τ	Tuning factor for temporal coherence
h	Step height (m)
k	Turbulent kinetic energy (m^2/s^2)
L	Cholesky factor of the Reynolds stress tensor
L_{int}	Integral length scale (m)
L_x, L_y, L_z	Domain sizes in x, y, z directions (m)
l_m	Mixing length (m)
N	Number of Fourier modes
N_x, N_y, N_z	Number of grid points in x, y, z directions
p	Pressure (Pa)
p'	Fluctuating pressure component (Pa)
P_k	Production of turbulent kinetic energy
R_{ij}	Reynolds stress tensor component, $\overline{u_i' u_j'}$ (m^2/s^2)
Re_b	Bulk Reynolds number
Re_h	Reynolds number based on step height
Re_θ	Momentum-thickness Reynolds number
Re_τ	Friction Reynolds number
S_{ij}	Strain rate tensor (1/s)
St	Strouhal number
t	Time (s)
T_{int}	Integral time scale (s)
u'	Fluctuating velocity component (m/s)
u_τ	Friction velocity (m/s)
$\hat{u}_n$	Scalar amplitude for Fourier mode n
U or $\mathbf{U}$	Velocity vector (m/s)
U_0	Inlet free-stream velocity (m/s)
U_b	Bulk velocity (m/s)
U_i	Mean velocity component in direction i (m/s)
$\overline{U}$	Time-averaged velocity vector (m/s)
x_i	Spatial coordinate in tensor notation (m)
y_1	Distance of first grid point from the wall (m)
y^+	Non-dimensional wall distance
Greek Symbols	
δ	Channel half-width (m)
δ_{ij}	Kronecker delta
ϵ	Turbulent dissipation rate (m^2/s^3)
η	Kolmogorov length scale (m)
κ	von Kármán constant
κ_n	Wavenumber vector for mode n (rad/m)
μ	Dynamic viscosity (Pa·s)
ν	Kinematic viscosity (m^2/s)
ρ	Fluid density (kg/m^3)
$\sigma^{(n)}$	Unit direction vector for Fourier mode n
τ_{ij}	Viscous stress tensor (Pa)
τ_w	Wall shear stress (Pa)
Φ_{ij}	Pressure-strain redistribution term

Continued on next page

Table 1 – continued from previous page

Symbol	Description
ψ_n	Random phase for Fourier mode n
ω	Specific dissipation rate (1/s)

List of Abbreviations

Table 2: List of Abbreviations

Abbreviation	**Description**
BFS	Backward-Facing Step
CGNS	CFD General Notation System
CPU	Central Processing Unit
DNS	Direct Numerical Simulation
DUV	Deep Ultraviolet
EUV	Extreme Ultraviolet
FIV	Flow-Induced Vibrations
GIF	Graphics Interchange Format
HPC	High-Performance Computing
LES	Large Eddy Simulation
PSD	Power Spectral Density
RANS	Reynolds-Averaged Navier-Stokes
RMSE	Root Mean Square Error
RSM EB	Reynolds Stress Model Elliptic Blending
SA	Spalart-Allmaras
SNGR	Stochastic Noise Generation and Radiation
SST	Shear Stress Transport
TCF	Turbulent Channel Flow
TDR	Turbulent Dissipation Rate
TKE	Turbulent Kinetic Energy

List of Tables

1 Introduction

Lithography machines produced at ASML consume large amounts of energy (about $1.3MW$ for an EUV machine from 2018, ASML [2024]) to create extremely fine patterns - known as transistors - onto photosensitive silicon wafers. These transistors form the fundamental building blocks of modern computing, acting as binary switches that represent "1s" (open) and "0s" (closed), depending on whether an electric current is flowing. Each silicon disks contain thousands of chips, each composed of millions of integrated circuits (ICs). These wafers are made of hundreds of layers, which need to be perfectly aligned one on top of the other to ensure the proper functioning of the chips.

Integrated circuits, commonly referred to as transistors, are fundamental to the global economy, powering anything that is digitalized, from household appliances to spacecraft. The evolution of semiconductor technology has required decades of research and development M. v. den Brink [2019]. This exponential growth in computational power has largely followed Moore's Law Moore [1965], doubling the number of transistors on a microchip approximately every two years Fang and He [2022]. To sustain this pace of innovation, ASML manufactures the world's most advanced semiconductor lithography machines.

Lithography is the process of transferring patterns from a mask onto a light-sensitive layer of a silicon wafer. ASML´s Deep Ultraviolet Light (DUV) and Extreme Ultra Violet Light (EUV) machines achieve these by using high-frequency light Kaiser [2024]. In EUV machines, it is obtained by shooting small tin droplets with a high-power laser, repeating the process about 50000 times per second Kazazis et al. [2024]. Besides, the laser used in the most recent generation of EUV, known as High-NA-EUV, shoots a tiny drop of tin, increasing massively its temperature. By heating it to approximately 200 000 °C, the material is converted into a hot plasma that emits the extreme ultraviolet (EUV) radiation necessary for etching microchip structures van Schoot [2024].

A schematic view of the inside of an NXE lithography system (part of the EUV line machines manufactured by ASML) is shown in Figure 1.

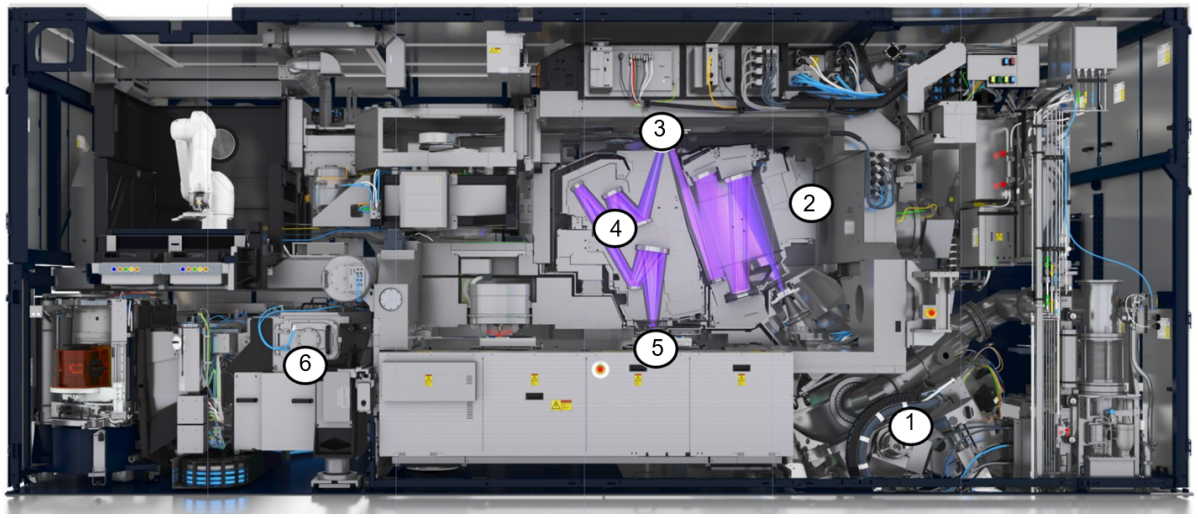

Figure 1: Schematic view of the inside of an ASML NXE lithography system (EUV line). The main subsystems are: (1) **Light source**, where a high-power laser irradiates tin droplets to generate EUV light at a wavelength of 13.5 nm; (2) **Vacuum chamber**, required since EUV light is strongly absorbed by air; (3) **Reticle** (or mask), which contains the circuit pattern to be transferred; (4) **Optical column**, composed of ultra-smooth mirrors manufactured by ZEISS that project and reduce the mask pattern with nanometer precision; (5) **Wafer stage**, which positions the silicon wafer with extreme accuracy during exposure; (6) **Wafer handler**, responsible for loading, unloading, and aligning wafers inside the system.

The emitted light is guided to the wafer through an advanced optical system, using a series of lenses provided by ZEISS, the world's leading manufacturer of ultra-smooth mirrors, crucial to the process as any imperfection may slightly distort the printed pattern. Optical systems enable the use of extremely powerful lasers at the shortest wavelengths required to produce the ultra-fine structures found in modern microchips: between 193 and 248 nanometers, and 13.5 nanometers, for DUV and EUV machines, respectively. Due to the very high accuracy required for these systems, the thermal expansion produced by heat harms the precision of the lithography patterns drawn on the wafers, as it introduces microscopic distortions in the wafer. Neways Electronics [2024].

Thus, lithography tools must remove large amounts of heat while keeping nanometre-scale geometry stable. Cooling is typically routed through complex piping and manifold systems (changes in cross-section, turns, splits and orifices are unavoidable), and those features produce pressure and velocity fluctuations that can excite flow-induced vibrations and acoustic noise which couple into the machine structure. Such disturbances — together with local

heat loads from the source and optics — generate thermal gradients that deform optics and wafers at nanometer scale and therefore directly degrade overlay, focus and critical-dimension control.

This has an impact on the perturbation of the flow, which will cause disturbances that will be transferred to the interior of the machine. Estimating how these disturbances are generated and propagated is crucial to improve the accuracy of the fine patterns being drawn onto the wafers. This ensures that the transfer of hydraulic and thermal disturbances into the optical path is limited so that EUV and DUV scanners can maintain the nanometer (and sub-nanometer) placement accuracy required for advanced nodes. The whole chain is validated with CFD and flow induced vibration (FIV) analysis and experiments.

1.1 Motivation - Why SNGR?

Flow-Induced Vibrations (FIV) are a critical concern in high-precision systems such as those developed by ASML. These vibrations are driven by turbulent pressure and velocity fluctuations, which must be accurately modeled to predict and mitigate their impact. They may be transferred to the inner core of the machine, where the source shoots a powerful laser onto the scanner. Any particles and vibrations at this stage will decrease the accuracy of the patterns drawn in the silicon wafers. Prior studies emphasize that FIV in piping networks arises from complex interactions between turbulence, geometry (e.g., bends, junctions, orifices), and structural response, often leading to fatigue or degraded system performance if not well understood and controlled Siba et al. [2021].

Reynolds Averaged Navier-Stokes (RANS) methods, such as the Reynolds Stress Model with Elliptic Blending (RSM-EB), are widely used due to their computational efficiency. They solve for the mean flow field and provide statistical quantities such as the turbulent kinetic energy (k) and Reynolds stress tensor ($R_{ij} = \overline{u_i' u_j'}$). However, RANS does not resolve the instantaneous velocity fluctuations $u_i'(t)$ that are directly responsible for FIV. This makes RANS insufficient for capturing the dynamic behavior of turbulence needed for vibration analysis Pope [2000] Menter [2009]. RANS methods are explained in detail in Section 2.1.3.

Direct Numerical Simulation (DNS) resolves all scales of turbulence and provides full spatiotemporal fidelity, including instantaneous fluctuations. While ideal in terms of accuracy, DNS is computationally prohibitive for the complex geometries and fast turnaround times required in industrial design workflows Moin and Mahesh [1998b] Sagaut [2006], where the design of the machine may change faster than the time it takes to fully run a DNS analysis.

Stochastic Noise Generation and Radiation (SNGR) offers a practical compromise between RANS and DNS, bridging the gap between both approaches. It uses statistical data from RANS (e.g., R_{ij} and k) to generate synthetic velocity fluctuation fields that match the desired turbulence statistics Schoder [2023]. These fields can be applied to arbitrary geometries and are computationally efficient Weitz et al. [2019]. Thus, providing a solution with higher accuracy than RANS and inferior compared to DNS, while taking less time than performing a DNS study, and longer than just RANS itself.

By reconstructing realistic fluctuations, SNGR enables the estimation of key quantities such as:

- Instantaneous pressure fluctuations (C_p')
- Turbulent kinetic energy (k)
- Reynolds stress distributions (R_{ij})

These outputs are essential for understanding and mitigating FIV in ASML systems, making SNGR a valuable tool in the design and validation process.

1.2 Goals of the project

The project aims to quantify the error and uncertainty in the Stochastic Noise Generation and Radiation (SNGR) method, specifically measuring how far a single time-loop or ensemble SNGR realization may deviate from a trusted DNS baseline. This leads to the first research question:

1. *How can the quantified uncertainties guide improvements to the SNGR code for internal flow applications?*

   Validate SNGR against high-fidelity Direct Numerical Simulations (DNS) based on refined Computational Fluid Dynamics (CFD) simulations. This is done by showing three- and four-way comparisons (CFD, $SNGR_{time}$, $SNGR_{ens}$, DNS) of Reynolds stresses and velocity profiles, confirming that SNGR reconstructs both the inner-layer scaling (U^+ vs. y^+) and the outer-layer recovery (U/U_0 vs. y/h) across six downstream slices. This leads to the second research question:

2. *How accurately does the SNGR code replicate DNS reference data for Backward Facing Step (BFS) flows?*

   Dynamically visualize turbulent structures by analyzing a backward-facing step flow, generating GIFs of instantaneous velocity fluctuations on the xy, xz and yz planes. These visualizations illustrate how eddies evolve, dissipate, and interact with the step. Besides, Power Spectral Density (PSD) is mapped at selected points to demonstrate correct inertial-range scaling($-5/3$) and high-frequency roll-off. This leads to the third research question:

3. *To what extent do temporal correlation and convective transport enhance the physical fidelity of SNGR, and what are the statistical convergence constraints associated with these mechanisms?*

The core objectives that guide the achievement of these goals are the following:

- **Embed physically consistent turbulent fluctuations on top of CFD mean fields.** Since RANS simulations only provide mean quantities, an essential step of SNGR is the reconstruction of synthetic velocity fluctuations that reproduce the correct turbulence statistics. This ensures that the reconstructed flow field captures the instantaneous dynamics responsible for flow-induced vibrations and aeroacoustic sources.
- **Validate SNGR against DNS for key statistics.** Direct Numerical Simulation (DNS) provides a reference with full spatiotemporal resolution. Comparing SNGR-generated fluctuations against DNS allows for quantitative assessment of accuracy in terms of Reynolds stresses, turbulent kinetic energy (TKE), and pressure fluctuations. This step establishes how closely SNGR can replicate physical reality while remaining computationally efficient.
- **Develop a modular MATLAB pipeline for implementation and analysis.** A robust workflow is required to handle the large amount of data and repeated testing involved in this study. The pipeline integrates modules for CFD data ingestion, SNGR setup, time-loop integration, and post-processing (e.g., statistical analysis, spectral maps, and visualization). This modularity enables flexible testing of different flow cases and facilitates improvements to the SNGR methodology.
- **Comprehensive post-processing and diagnostics.** Having applied the SNGR method using the input data extracted from StarCCM+, a structured post-processing approach is performed, converting raw instantaneous snapshots into the quantitative products required for validation and interpretation. Time-averaged statistics (from a single time-loop where eddies are convected and temporal correlation is present) and ensemble statistics (TKE and the full Reynolds-stress tensor) are computed. Slice-wise profiles are extracted and interpolated onto a common vertical grid for direct comparisons, and automated per-slice diagnostics (RMSE, maximum and mean absolute error, relative RMSE, etc) are produced and saved as CSV summaries. Aditionally, spatial diagnostics (combined difference plots and heatmaps clipped to physically relevant y/h ranges), and temporal outputs (GIFs of instantaneous fluctuation fields and PSDs at probe points), are also generated. Furthermore, before and after reports when anisotropy corrections or tensor repairs are applied.

1.3 Stochastic Noise Generation and Radiation (SNGR)

The Stochastic Noise Generation and Radiation nmethod method was originally proposed by Billson Billson et al. [2004], reconstructs time-resolved, spatially correlated velocity fluctuations starting from steady RANS statistics. SNGR is designed to produce realizations of $u_i'(\mathbf{x}, t)$ that reproduce prescribed second-order statistics (Reynolds stresses, turbulent kinetic energy and energy spectrum), satisfy incompressibility conditions, and possess prescribed temporal correlation and convection properties so that the synthetic fluctuations can be used in downstream unsteady or aeroacoustic computations Schoder [2023], Weitz et al. [2019].

The SNGR method used in this study follows the approach proposed by Billson Billson et al. [2004], generating turbulent velocity fluctuations using a Fourier series, whose argument is influenced by certain random variables and whose amplitude is derived from the energy wave spectrum. After these turbulent fluctuations have been produced, they are transported by the mean velocity field, ensuring that the synthetic turbulence moves in the right physical direction and the turbulence structures are convected by the primary flow. Following this convection process, an exponential time correlation is applied to model the temporal evolution of the turbulent fluctuations. Finally, an anisotropic scaling transformation is performed using the Reynolds Stress Tensor obtained from the RANS solution (with or without Wilcox correction), ensuring that the turbulence structures accurately represent the physical anisotropy that characterizes turbulent flow. Bechara et al. Bechara et al. [1994] were the first to apply this stochastic approach for noise computation in the context of turbulent flows, marking a significant advancement in the field by integrating synthetic turbulence generation with noise prediction methodologies. Their study was similar to the approach followed by Karweit et al. Karweit et al. [1991].

In this work, the method is implemented as described in the original paper by Billson Billson et al. [2004], with the goal of validating its applicability to flow-induced vibration prediction in lithography machines. A different geometry - backward-facing step, which is highly relevant for estimation of flow induced vibrations - is evaluated. Furthermore, a set of distinct validation metrics is employed to assess the performance of the method.

The synthetic turbulence is compared against target Reynolds stress tensors to verify that the anisotropic scaling transformation is correctly applied. Mean velocity profiles are analyzed to ensure that the background flow field used for the convection of synthetic fluctuations is physically consistent and properly developed upstream of the step. These profiles are also evaluated in wall units to assess near-wall behavior and confirm consistency with physical turbulence structures.

Additionally, Power Spectral Density (PSD) comparisons are performed to validate that the energy distribution across frequencies aligns with expected turbulence characteristics.The turbulent kinetic energy is evaluated along the streamwise midplane and spanwise-averaged, under different inlet conditions and mesh resolutions, to assess the robustness of the SNGR method across varying flow regimes.

Finally, a time-marching analysis of turbulent fluctuations is conducted, comparing low and high turbulence planes between CFD and SNGR-generated fields. This provides insight into the temporal evolution and spatial coherence of the synthetic turbulence, further validating its suitability for downstream vibration and acoustic analyses.

The conceptual stages of the SNGR method — including mode sampling, isotropic Fourier synthesis, convection and temporal correlation, and anisotropic scaling — are briefly introduced in Section 3.1 to provide an overview of the methodology. These steps form the foundation of the synthetic turbulence generation process and are essential for understanding the subsequent analysis.

1.4 Comparison Between Time-Marching and Ensemble-Averaged SNGR Approaches

The Stochastic Noise Generation and Radiation (SNGR) turbulence generation method can be implemented in two distinct ways: a **time-marching (single-loop)** approach and an **ensemble-averaged** approach. Each method serves different purposes and exhibits unique characteristics in terms of temporal coherence, statistical convergence, and applicability to dynamic load analysis.

1.4.1 Time-Marching Approach

In the time-marching implementation, SNGR generates a single, continuous realization of the turbulent fluctuation field, $u'(x, t)$, that evolves over time. Each new field is constructed by convecting and decorrelating the previous one, thereby preserving temporal coherence. This results in physically plausible eddy evolution—eddies persist, grow, and dissipate—allowing for the extraction of meaningful spectral content from unsteady loads.

Main strengths from this approach is that it captures temporal dynamics of turbulence - enabling analysis of force spectra and vibration responses - and it is suitable for aeroelastic and vibroacoustic studies, where time-resolved loading is critical. However, it requires long simulation times to achieve statistical convergence, and it is computationally more expensive due to time-stepping and memory requirements.

The process to advance from one timestep to the next consists of four main operations, ensuring that the anisotropy transformation is applied as the final step to a consistently isotropic field:

1. **Numerical Advection:** The *isotropic* fluctuation field from the previous timestep, $u'_{iso}(t)$, is transported forward in time using the steady, local mean velocity field, $\overline{U}(x)$, from the RANS simulation. This is performed using an Eulerian advection solver (First-Order Upwind in this work). This step produces a convected, but still isotropic, field, $u'_{conv,iso}(t + \Delta t)$.

2. **New Fluctuation Generation:** A completely new, statistically independent, isotropic fluctuation field, $u'_{new,iso}(t + \Delta t)$, is generated at each timestep using the Fourier synthesis method described in Section 3.3.

3. **Temporal Correlation:** To ensure that eddies decorrelate over a physical timescale rather than disappearing abruptly, the two isotropic fields are blended. The result is the new blended isotropic field for the current timestep, $u'_{iso}(t + \Delta t)$. This is achieved using an exponential correlation factor, $a = \exp(-f_\tau \Delta t/\tau)$, where τ is the local turbulence timescale (k/ϵ) and f_τ is a tuning parameter. The blended isotropic field is computed as:

$$u'_{iso}(t + \Delta t) = a \cdot u'_{conv,iso}(t + \Delta t) + \sqrt{1 - a^2} \cdot u'_{new,iso}(t + \Delta t) \tag{1}$$

4. **Apply Anisotropy:** The blended isotropic field, $u'_{iso}(t + \Delta t)$, is transformed into the final, *anisotropic* output field, $u'(t + \Delta t)$. This is done using a linear mapping based on the Cholesky decomposition of the local Reynolds stress tensor (R_{ij}). This final step introduces the correct directional properties, ensuring the resulting field matches the target anisotropic turbulence statistics. The field $u'(t + \Delta t)$ is stored for analysis, while $u'_{iso}(t + \Delta t)$ is carried forward to the next time step.

This three-step process, repeated in a time loop, results in a physically plausible evolution where eddies are transported by the mean flow while gradually decaying and being replaced by new structures. This temporal coherence is essential for analyzing frequency-dependent phenomena like force spectra and vibration responses, making the approach suitable for aeroelastic and vibroacoustic studies.

However, this method is computationally expensive. The primary reason for the high memory requirementis the need to store the entire fluctuation field at every timestep. The output variable `vFluctHistory` is pre-allocated as a three-dimensional array of size [`nSteps` × `nCells` ×3]. For a typical simulation with millions of cells and thousands of timesteps, this array can easily require hundreds of gigabytes of RAM, making it demanding on computational resources. Furthermore, the iterative, step-by-step nature of the calculation is more time-consuming than generating independent snapshots.

The complete iterative workflow of the time-marching approach is visualized in Figure 2.

Figure 2: Workflow of the SNGR Time-Marching (Single-Loop) Approach.

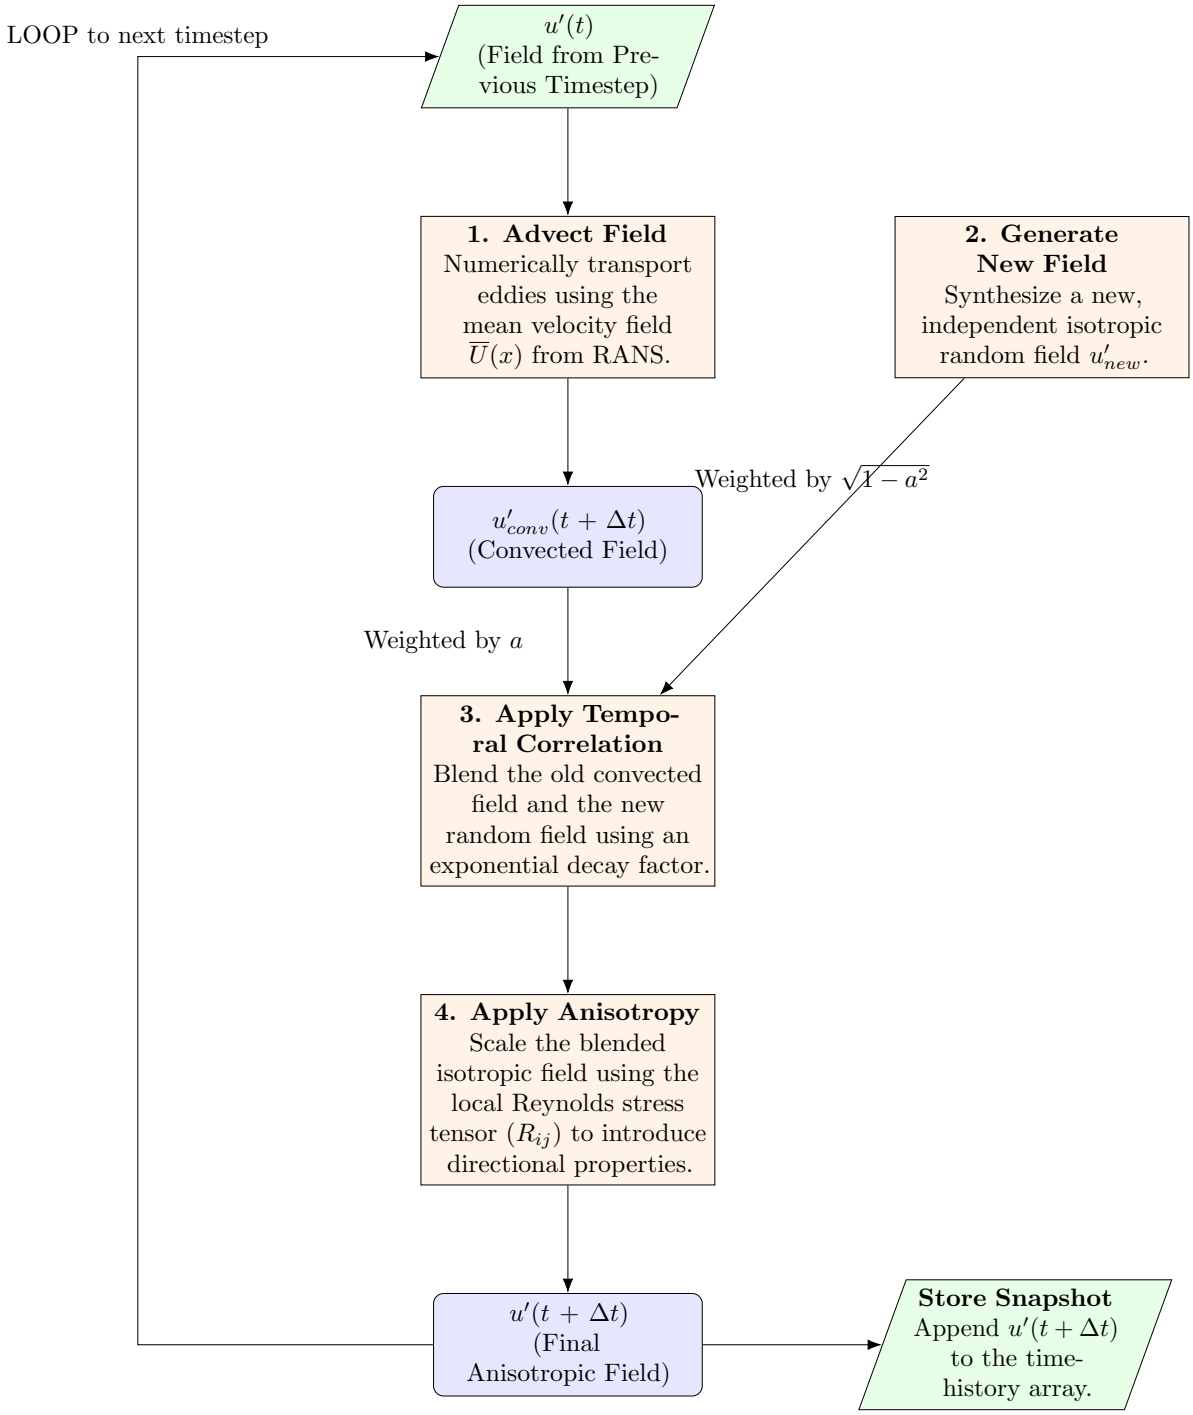

1.4.2 Ensemble-Averaged Approach

The ensemble method generates a set of M statistically independent turbulence snapshots, $\{u'_1, u'_2, ..., u'_M\}$. Each individual snapshot is a unique, random realization of the turbulent field. A single snapshot, by itself, does not and cannot match the prescribed statistical properties (e.g., the Reynolds stress tensor, R_{ij}). The statistical properties are only recovered when averaging over a large number of these independent snapshots. This is implemented in MATLAB by looping from $m = 1$ to M and accumulating the statistics, then dividing by M, as expressed in Equation (2).

$$\langle R_{ij} \rangle_{ens} = \frac{1}{M} \sum_{m=1}^{M} u'_{i,m} u'_{j,m} \tag{2}$$

Therefore, it is the resulting **ensemble-averaged field** that is constructed to match the input RANS statistics. This approach is ideal for validating the statistical fidelity of the fluctuation generator itself, as it directly verifies that the underlying synthesis method is correct without the complexities of time-stepping.

The most relevant advantages are that it enables direct verification of the SNGR generator's statistics against the RANS/CFD inputs and facilitates quantification of the number of samples (M) required for statistical convergence. For instance, by comparing the ensemble statistics to the Reynolds stresses imported from STAR-CCM+, it can be confirmed whether the SNGR method accurately reproduces the prescribed turbulence characteristics without the need for long time integrations.

Nonetheless, as there is no temporal correlation between snapshots, this method cannot be used to compute frequency spectra or dynamic loads, making it inapplicable for vibration studies with dynamic analyses.

This parallel process of generation and averaging is summarized in the workflow diagram shown in Figure 3.

Figure 3: Workflow of the SNGR Ensemble-Averaged Approach.

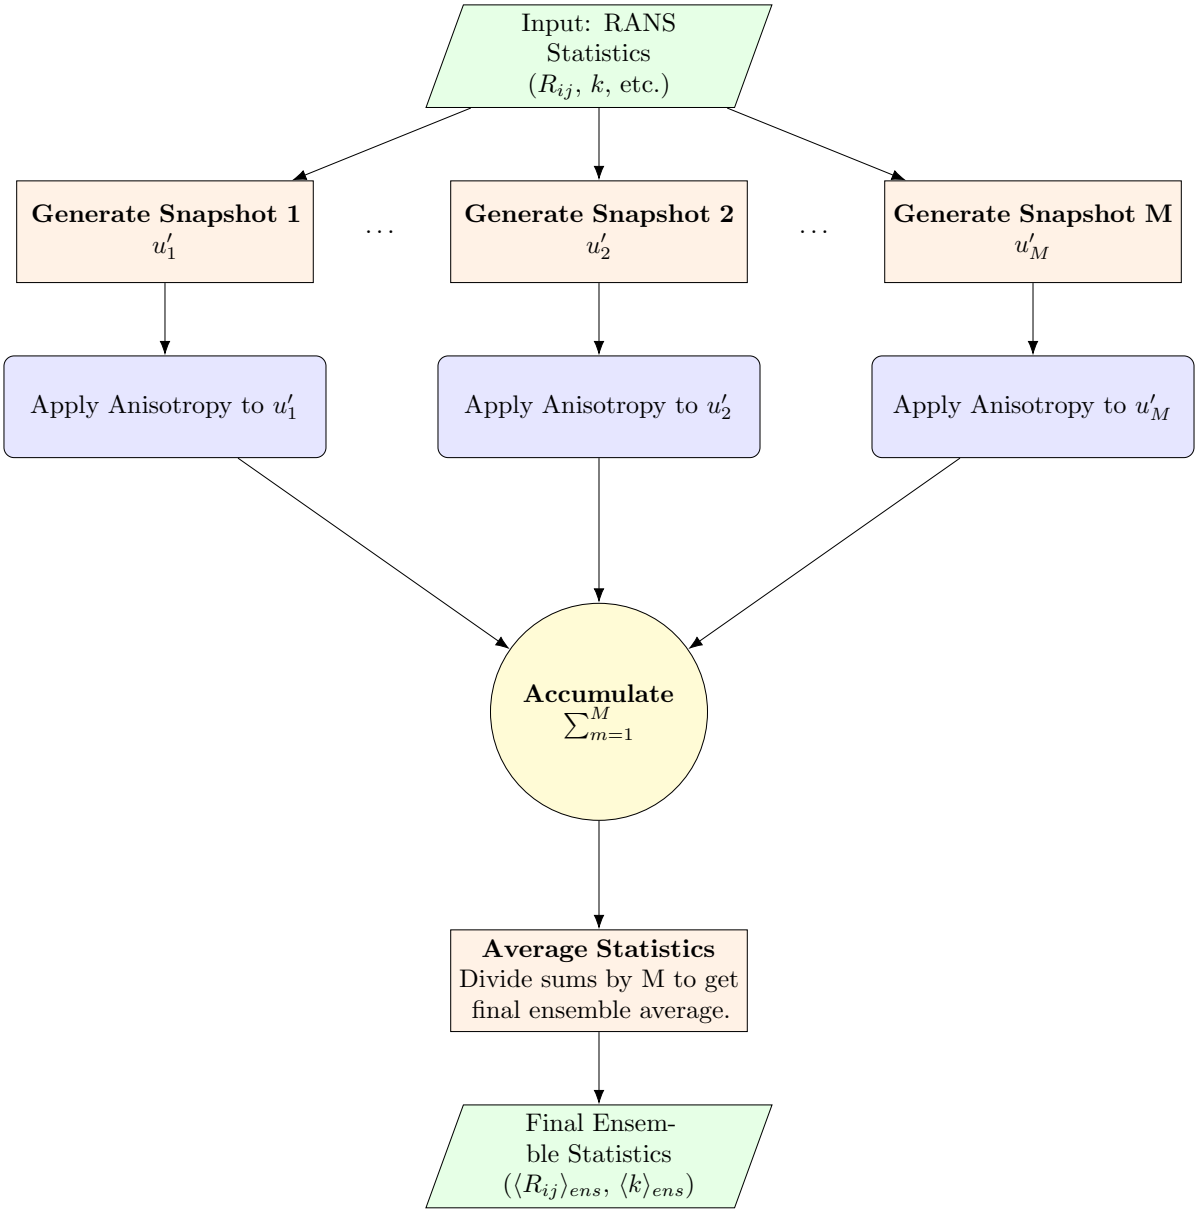

2 Literature Review

2.1 Turbulent Flows

The concept of isotropic turbulence was originally introduced by Taylor Taylor [1935]. Early works considered this assumption to be true, ignoring anisotropic effects at sufficiently small scales. One of the most influential contributions to this idea is Kolmogorov's classical turbulence theory, which suggests that in the inertial subrange, turbulence becomes statistically isotropic. Jiménez [2004] This implies that any large-scale anisotropies introduced by boundary conditions or external forces are gradually dissipated as the turbulent cascade develops, leading to a loss of directional dependence in the turbulent structures. As discussed by Biferale and Procaccia Biferale and Procaccia [2005], see p. 46, in all realistic flows, there always exists some anisotropy at all scales. This is because either the driving mechanism or the boundary geometry affects the velocity field's statistical characteristics, which are never rotationally invariant. Tavoularis and Corrsin [1981] Thus, the synthetic turbulent flow must include anisotropic elements to accurately represent the turbulent structures as they are physically occurring.

Taking into account the fact that compressibility has no impact on the specification of the turbulent length scales of the flow (as suggested by Morkovin Morkovin [1962]), turbulence structures remain largely unaffected by compressibility effects, considering small density fluctuations. This explains the use of incompressible turbulence models with minor adjustments for compressible regimes (which are mostly neglected in this study).

2.1.1 Governing equations of Turbulent Flow

Turbulent flows are governed by the Navier–Stokes equations, which describe the conservation of mass and momentum in a fluid. These equations form the basis for both experimental and numerical studies of turbulence. The continuity equation - written in the general form - ensures mass conservation:

$$\frac{\partial \rho}{\partial t} + \frac{\partial (\rho u_i)}{\partial x_i} = 0 \tag{3}$$

The momentum equation accounts for pressure, viscous stresses, and body forces:

$$\frac{\partial (\rho u_j)}{\partial t} + \frac{\partial (\rho u_i u_j)}{\partial x_i} = -\frac{\partial p}{\partial x_j} + \frac{\partial \tau_{ij}}{\partial x_i} + \rho f_j \tag{4}$$

Here, the viscous stress tensor τ_{ij} is defined as:

$$\tau_{ij} = \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \frac{2}{3} \delta_{ij} \frac{\partial u_k}{\partial x_k} \right) \tag{5}$$

Where ρ is the fluid density [kg/m^3], u_i the velocity components [m/s] in directions x, y, z, x_i the spatial coordinates in tensor notation, t the time [s], p represents pressure [N/m^2], μ the dynamic viscosity [Ns/m^2], δ_{ij} the Kronecker delta (1 if $i = j$, 0 otherwise), f_j is the body force per unit volume [N], τ_{ij} means the viscous stress tensor [N/m^2].

2.1.2 Compressibility Assumption

The above equations describe compressible flow, where density ρ varies with space and time. However, for this study, compressibility effects on the flow field are small and the incompressible approximation is commonly used, due to low Mach number conditions (Ma $\ll 0.3$), allowing simplification of the governing equations.

Under constant density ($\rho = \text{const}$) the continuity equation simplifies to the divergence-free condition

$$\frac{\partial u_i}{\partial x_i} = 0, \tag{6}$$

and the momentum equation reduces to the incompressible Navier–Stokes form used throughout this work, commonly written in the divergence-free form when assuming incompressible flow:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\frac{1}{\rho}\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}, \tag{7}$$

$$\nabla \cdot \mathbf{u} = 0 \tag{8}$$

Numerical solution algorithms typically enforce $\nabla \cdot \mathbf{u} = 0$ using projection or pressure-Poisson methods Chorin [1968]. Besides, numerical analyses of compressible solvers in the low-Mach regime highlight additional caveats (e.g. the need for preconditioning or specially designed schemes) if a compressible code is used near the incompressible limit Guillard and Viozat [1999].

Although Ma $\ll 1$ justifies incompressibility for momentum transport, thermal effects (local heating, thermal expansion) may still be important in precision equipment. For small density variations driven by temperature and

where only buoyancy is relevant, the Boussinesq approximation can be used to include buoyancy while keeping an incompressible momentum formulation. If density variations or acoustic effects become significant, a compressible formulation must be retained. The approach followed in this project neglects thermal changes.

2.1.3 Reynolds decomposition and Reynolds-Averaged Navier-Stokes (RANS) Equations

RANS models approximate the turbulent effects by solving the governing equations of the flow considering mean values of their parameters, modeling all turbulent structures present in the flow while providing a reasonably accurate solution that requires significantly less computational resources than methods like Large Eddy Simulation (LES) or Direct Numerical Simulation (DNS).

The fundamental formulation originates from the Reynolds decomposition of flow variables into mean and fluctuating components. By substituting these decomposed variables into the Navier-Stokes equations and applying averaging, the closure problem (discussed in Section 2.2) is obtained, which requires modeling of the Reynolds stress tensor RAN [2009].

Reynolds decomposition is applied to both velocity and pressure to derive equations for mean quantities:

$$u_i(\mathbf{x},t) = \overline{u}_i(\mathbf{x},t) + u_i'(\mathbf{x},t) \tag{9}$$

$$p(\mathbf{x},t) = \overline{p}(\mathbf{x},t) + p'(\mathbf{x},t) \tag{10}$$

where the overbar denotes an appropriate averaging operator (time, ensemble or spatial averaging depending on context) Tennekes and Lumley [1972b], Pope [2000]. Inserting these decomposed fields into the incompressible Navier–Stokes equations and averaging yields the Reynolds-averaged Navier–Stokes (RANS) momentum equation:

$$\frac{\partial \overline{u}_j}{\partial t} + \overline{u}_i \frac{\partial \overline{u}_j}{\partial x_i} = -\frac{1}{\rho}\frac{\partial \overline{p}}{\partial x_j} + \nu \frac{\partial^2 \overline{u}_j}{\partial x_i \partial x_i} - \frac{\partial \overline{u_i' u_j'}}{\partial x_i} + \overline{f}_j, \tag{11}$$

where the term $R_{ij} \equiv \overline{u_i' u_j'}$ is the Reynolds stress tensor. The Reynolds stresses represent momentum transport by turbulent fluctuations and introduce the classical *closure problem*: Eqs. (11) are not closed unless a model is provided for R_{ij}. This is the fundamental reason why RANS solves only mean flow and low-order statistics (e.g. turbulent kinetic energy k and dissipation ε) but does not provide instantaneous fluctuations $u_i'(t)$ that directly drive flow-induced vibration (FIV) Pope [2000], Wilcox [1998].

Having discussed the concept of RANS equations and the procedure followed for their obtantion, the three primary categories of RANS turbulence models are described:

- **Zero-equation model**: Also known as *Prandtl's mixing length model* Prandtl [1949], this approach estimates eddy viscosity as exhibited in Equation (12):

$$\nu_t = l_m^2 \left| \frac{dU}{dy} \right| \tag{12}$$

  where ν_t is the turbulent eddy viscosity, l_m is the mixing length, and $\frac{dU}{dy}$ is the velocity gradient in the shear direction. This model is inexpensive but cannot capture complex non-equilibrium effects.

- **One-equation models**: One-equation models evolve a single transported scalar that is then used to compute ν_t. Therefore, introducing a single transport equation for estimating turbulent velocity scales. A widely used option in aerospace applications is the Spalart–Allmaras (SA) model, whose generic structure is presented in Equation (13).

$$\frac{\partial \tilde{\nu}}{\partial t} + \overline{u}_j \frac{\partial \tilde{\nu}}{\partial x_j} = P_{\tilde{\nu}}(\tilde{\nu}, \overline{\mathbf{U}}) - D_{\tilde{\nu}}(\tilde{\nu}, d) + \frac{1}{\sigma}\frac{\partial}{\partial x_j}\Big[(\nu + \tilde{\nu})\frac{\partial \tilde{\nu}}{\partial x_j}\Big], \tag{13}$$

  and the turbulent viscosity is obtained from a nonlinear function of $\tilde{\nu}$, typically $\nu_t = f(\tilde{\nu})$ (SA uses several limiter/near-wall functions). The SA model is compact and robust for attached and mildly separated flows Spalart and Allmaras [1992].

- **Two-equation models**: Two-equation models solve transport equations for two turbulence scalars (commonly k and a dissipation or frequency variable) from which ν_t is formed Menter [1994]. They are the most widely used class in CFD due to their balance between accuracy and computational cost. Key implementations include:

  - **k-ϵ model**: Developed by Jones and Launder [1972a], this solves two transport equations for turbulent kinetic energy (TKE, k), and turbulent dissipation rate (TDR, ϵ). It is effective primarily for fully

turbulent flows Inc. [2006]. The standard form of this family of turbulence models is described in Equation (14) and Equation (15).

$$\frac{\partial k}{\partial t} + \overline{u}_j \frac{\partial k}{\partial x_j} = P_k - \varepsilon + \frac{\partial}{\partial x_j}\left[(\nu + \nu_t/\sigma_k)\frac{\partial k}{\partial x_j}\right], \tag{14}$$

$$\frac{\partial \varepsilon}{\partial t} + \overline{u}_j \frac{\partial \varepsilon}{\partial x_j} = C_{1\varepsilon}\frac{\varepsilon}{k}P_k - C_{2\varepsilon}\frac{\varepsilon^2}{k} + \frac{\partial}{\partial x_j}\left[(\nu + \nu_t/\sigma_\varepsilon)\frac{\partial \varepsilon}{\partial x_j}\right], \tag{15}$$

where $P_k = 2\nu_t S_{ij} S_{ij}$ is the production term. A common set of model constants is $C_\mu = 0.09$, $C_{1\varepsilon} = 1.44$, $C_{2\varepsilon} = 1.92$, $\sigma_k = 1.0$, $\sigma_\varepsilon = 1.3$ Jones and Launder [1972b], Launder and Spalding [1974].

– k-ω **model (and SST variants):** Particularly effective for boundary-layer studies due to its superior near-wall behaviour, the k–ω family solves transport equations for turbulent kinetic energy and the specific dissipation (or frequency) ω. A common form of the transport equations used in practical implementations is described in Equation (16) and Equation (17).

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho \overline{u}_j k)}{\partial x_j} = P_k - \beta^* \rho k\omega + \frac{\partial}{\partial x_j}\Big[(\mu + \sigma_k \mu_t)\frac{\partial k}{\partial x_j}\Big], \tag{16}$$

$$\frac{\partial(\rho \omega)}{\partial t} + \frac{\partial(\rho \overline{u}_j \omega)}{\partial x_j} = \gamma\frac{P_k}{\mu_t} - \beta\rho\omega^2 + \frac{\partial}{\partial x_j}\Big[(\mu + \sigma_\omega \mu_t)\frac{\partial \omega}{\partial x_j}\Big] + (\text{cross-diffusion}). \tag{17}$$

Here P_k denotes the production of k, μ_t the eddy viscosity and the model constants $(\beta, \beta^*, \gamma, \sigma_k, \sigma_\omega)$ depend on the chosen variant (Wilcox, SST blending values, etc.). As with other eddy-viscosity closures the modeled Reynolds stresses use the Boussinesq relation:

$$R_{ij} \approx -\frac{2}{\rho}\mu_t S_{ij} + \frac{2}{3}k\,\delta_{ij}, \qquad S_{ij} = \tfrac{1}{2}\left(\partial_j \overline{u}_i + \partial_i \overline{u}_j\right). \tag{18}$$

Menter's Shear-Stress-Transport (SST) model is a hybrid that blends a robust near-wall k–ω formulation with k–ε behaviour in the free stream to avoid freestream sensitivity while retaining near-wall accuracy. The blending is achieved using functions (commonly denoted F_1, F_2) so that coefficients and certain terms smoothly transition from the near-wall (ϕ_1) to the outer-region (ϕ_2) values, as expressed in Equation (27).

$$\phi = F_1\,\phi_1 + (1 - F_1)\,\phi_2. \tag{19}$$

SST also introduces a stress-limiter (eddy-viscosity limiter) that reduces excessive μ_t production in regions of strong adverse pressure gradient. A convenient operational form for the SST eddy viscosity is expressed in Equation (20).

$$\mu_t \;=\; \rho\,\frac{a_1 k}{\max\left(a_1\omega,\; SF_2\right)}, \tag{20}$$

where $S = \sqrt{2S_{ij}S_{ij}}$ is the strain-rate magnitude and a_1 a model constant; the $\max(\cdot)$ limiter is key to SST's improved performance in separated flows. The SST model is widely recommended for BFS-type separated flows because it can be integrated down to the wall without damping functions, blends to a k–ε behaviour away from walls (avoiding freestream sensitivity), and contains mechanisms that improve prediction of separation and reattachment locations compared with simpler two-equation eddy-viscosity models Menter [1994], Wilcox [1998].

– **Reynolds-stress transport (RSM) and elliptic-blending variants:** When flow physics produce strong anisotropy (significant normal stress differences, strong curvature, or highly separated wakes), the Boussinesq assumption underlying eddy-viscosity models is inadequate. Reynolds-stress models (RSM) address this by solving transport equations for the Reynolds-stress tensor components themselves, as formulated in Equation (21).

$$\frac{\partial R_{ij}}{\partial t} + \overline{u}_k \frac{\partial R_{ij}}{\partial x_k} = P_{ij} + \Phi_{ij} - \varepsilon_{ij} + D_{ij}, \tag{21}$$

where P_{ij} is production, Φ_{ij} the pressure–strain (redistribution) term, ε_{ij} the dissipation, and D_{ij} the diffusive transport. The pressure–strain term Φ_{ij} is the central modelling challenge and is typically decomposed into slow (return-to-isotropy), rapid (mean-strain dependent) and wall-reflection contributions. The choice of Φ_{ij} determines much of the RSM behaviour and calibration.

Elliptic-blending RSM (EB-RSM) further improves near-wall representation by introducing a non-local elliptic relaxation (or blending) function f that captures wall-induced redistribution without purely local damping functions. A compact form of the elliptic-blending operator is:

$$L^2\nabla^2 f - f \;=\; S_f, \tag{22}$$

where L is a turbulence length scale and S_f a source constructed from local invariants; f is then used to blend between outer-layer and near-wall pressure–strain models. EB-RSM improves prediction of anisotropy and non-local wall effects at the cost of solving additional equations (one elliptic equation plus the full RSM system) Manceau and Hanjalić [2002].

As a final remark, while SST provides a robust baseline and has been used for supplementary comparisons, **the RSM–EB model constitutes the reference RANS framework for this study**. Its ability to reproduce anisotropy and non-local near-wall redistribution makes it more consistent with the objectives of the SNGR validation, which aims to analyse how turbulence structure and stress anisotropy influence reconstructed velocity fluctuations. The k–ω **SST** model, however, is also employed in complementary runs for sensitivity analysis and baseline comparison, given its proven reliability for separated flows such as the backward-facing step (BFS). Using both models enables a clear assessment of how eddy-viscosity versus stress-transport closures affect the reconstruction accuracy of the SNGR method.

2.2 Turbulence Closure models

The Boussinessq hypothesis comes from the concept based on expressing the turbulent shear stress and heat fluxes in terms of mean flow variables, modeling them as proportional to the mean flow gradients (similar to viscous stresses in laminar flows). The Reynolds stress tensor is defined as:

$$\tau_{ij} = 2\mu_t \left(S_{ij} - \frac{1}{3}\frac{\partial u_k}{\partial x_k}\delta_{ij} \right) - \frac{2}{3}\rho k \delta_{ij} + \underbrace{\frac{2}{3}\mu_t \frac{\partial u_k}{\partial x_k}\delta_{ij}}_{\text{Dilatation term}} \quad \textit{Marvin} \text{ [1977]}. \tag{23}$$

where S_{ij} is the mean strain rate tensor, μ_t is the eddy viscosity and k is the turbulent kinetic energy. This framework applies to both incompressible and compressible flows, with the latter requiring additional terms for dilatational effects (for instance, pressure-dilatation correlations). Expressing the Navier-Stokes equations considering mean flow variables can be applied to both incompressible and compressible flow, the latter being an extension of techniques used for the former Marvin [1977].

2.2.1 Closure Approaches

Turbulence closure models can be classified based on their Reynolds stress tensor. In first-order closures, second-order correlations, such as the Reynolds stresses, are modeled using first-order correlations (i.e., mean flow variables). Higher-order closures, in contrast, introduce additional transport equations to express third- and higher-order correlations in terms of second-order or higher correlations. Pope [2000].

An example of first-order closures are eddy-viscosity models, whose turbulent stresses are defined as proportional to mean strain rated through the Bousinessq hypothesis, as seen in Equation (24).

$$-\rho\overline{u_i' u_j'} = 2\mu_t S_{ij} - \tfrac{2}{3}\rho k\, \delta_{ij}, \tag{24}$$

where μ_t is the turbulent viscosity, $S_{ij} = \frac{1}{2}\left(\partial\overline{u}_i/\partial x_j + \partial\overline{u}_j/\partial x_i\right)$ is the mean strain-rate tensor, and k is the turbulent kinetic energy.

Other examples are the mean velocity, and simple algebraic models like Prandtl's mixing length theory that define μ_t as expressed in Equation (25).

$$\mu_t = \rho\, \ell^2 \left| \frac{\partial \overline{u}}{\partial y} \right|, \tag{25}$$

with ℓ being the mixing length.

More advanced two-equation models, such as k–ϵ and k–ω, introduce additional transport equations (PDE´s) for turbulent scales, balancing accuracy with computational efficiency. They are computationally efficient, but lack accuracy when evaluating flows strong separation and streamline curvature, because they assume turbulence as isotropic.

These limitations are solved by directly solving transport equations for Reynolds stresses, as done by higher-order closures. For incompressible flow, the Reynolds stress transport equation reads:

$$\frac{\partial \overline{u_i' u_j'}}{\partial t} + \overline{u}_k \frac{\partial \overline{u_i' u_j'}}{\partial x_k} = P_{ij} + \Phi_{ij} - \varepsilon_{ij} + D_{ij}, \tag{26}$$

where P_{ij} is production, Φ_{ij} is pressure–strain redistribution, ε_{ij} is dissipation, and D_{ij} is turbulent diffusion Launder et al. [1975].

For instance, Bradshaw's formulation accounts for directional effects by relating shear stress to turbulent kinetic energy ($\overline{u'v'} \propto k$) Bradshaw [1973]. To properly analyze complex physical phenomena, it is required to solve 5 to 7

additional transport equations compared to standard two-equation models, severely increasing the computational cost of the models Launder et al. [1975].

Having explained the main differences between first and higher-order approaches, it is possible to blend them using hybrid strategies, combining the efficiency of eddy-viscosity models with the accuracy of Reynolds stress transport models. One of the most common hybrid implementations is the k-ω Shear Stress Transport (SST) model Menter [1994], which blends the k-ω (near the wall) and k-ϵ (in the freestream region) formulations through an interpolation function F_1:

$$\phi = F_1\phi_{k\omega} + (1 - F_1)\phi_{k\epsilon}, \tag{27}$$

where ϕ denotes the generic model variable. The SST model achieves improved accuracy for separated flows without compromising the Boussinesq assumption Menter et al. [2003], and will be used to evaluate the performance of the SNGR method when evaluating a backward-facing step geometry. Other hybrid strategies include the ζ-f model Hanjalic et al. [2004] , which incorporates elliptic relaxation to add wall effect modeling, and the Explicit Algebraic Reynolds Stress Model (EARSM) Wallin and Johansson [2000], which combines anisotropic effects within a two-equation framework, keeping the cost close to two-equation models.

2.3 Anisotropy

Turbulence anisotropy is important for predicting the right velocity and pressure fluctuations in turbulent flows that drive flow-induced vibration (FIV). This huge influence on the aerodynamic mixing noise is shown by Khavaran [1999], being a crucial concept as pressure fluctuations caused by turbulent structures cause significant forces on surfaces.

In this work anisotropy is enforced by modifying the target Reynolds-stress tensors prior to mapping them onto the synthetic snapshots. Concretely, the SNGR method computes group-averaged target R_{ij} fields from the RANS data exported from *StarCCM+* (grouping cells by wall-normal coordinate), and then applies a near-wall Wilcox-style correction to those group targets. This adjusts the synthetically generated turbulent fluctuations to match the expected stress distributions observed in real turbulent flows. The corrected tensor array is then used to map isotropic synthetic fields into anisotropic fields via a local matrix square root (Cholesky) so that

$$\mathbf{u}^{\text{aniso}}(\mathbf{x}, t) = \mathbf{L}(\mathbf{x})\,\mathbf{u}^{\text{iso}}(\mathbf{x}, t), \qquad \text{with} \qquad \mathbf{R}(\mathbf{x}) = \mathbf{L}(\mathbf{x})\,\mathbf{L}(\mathbf{x})^{\top}, \tag{28}$$

which guarantees that the mapped snapshots recover the corrected Reynolds stresses, $\overline{u_i' u_j'} = R_{ij}$, when applied to zero-mean, unit-variance isotropic input fields. Furthermore, diagnostics are provided when running the SNGR code, using them to validate the near-wall correction and the stress recovery.

2.3.1 Wilcox correction

Wilcox - whose k-ω turbulence model forms the basis for correcting the turbulent eddies in the Stochastic Noise Generation and Radiation (SNGR) method, Wilcox [1998] - developed a two-equation model to approximate the governing equations of the flow for boundary layer flows. The model consists of:

- A turbulent kinetic energy (k) transport equation:

$$\frac{Dk}{Dt} = P_k - \beta^* k\omega + \frac{\partial}{\partial x_j}\left[(\nu + \sigma^*\nu_t)\frac{\partial k}{\partial x_j}\right] \tag{29}$$

- A specific dissipation rate (ω) equation derived from Saffman's pseudo-vorticity transport theory Saffman [1970]:

$$\frac{D\omega}{Dt} = \alpha\frac{\omega}{k}P_k - \beta\omega^2 + \frac{\partial}{\partial x_j}\left[(\nu + \sigma\nu_t)\frac{\partial \omega}{\partial x_j}\right] \tag{30}$$

In collaboration with Tracy, Wilcox maintained the canonical turbulence intensity ratios:

$$\langle u'^2\rangle : \langle v'^2\rangle : \langle w'^2\rangle = 4 : 3 : 2 \tag{31}$$

This ratio, originally established through boundary layer measurements by Bradshaw [1967] and later incorporated into turbulence modeling by Wilcox [1998], is preserved in the SNGR method to adjust the components of the turbulent kinetic energy extracted from a RANS simulation, improving the estimation and representation of the Reynolds stresses.

3 Methodology of the SNGR method

3.1 Method overview

In practice SNGR proceeds in four conceptual stages:

1. **Mode sampling:** choose a discrete set of wave-numbers $\{\kappa_n\}_{n=1}^{N}$ that cover the spectral band of interest and sample random directional/phase parameters for each mode. This is discussed in Section 3.4.
2. **Isotropic Fourier synthesis:** generate an initial, divergence-free, isotropic velocity field as a finite Fourier sum where the modal amplitudes are drawn from the target one-dimensional energy spectrum $E(\kappa)$. This aspect is elaborated in Section 3.3.
3. **Convection & temporal correlation:** this stage imposes mean-flow convection and temporal correlation to generate realistic time dynamics. A control parameter is used to balance the relative influence of convection and time correlation, ensuring that the generated fluctuations evolve realistically in time and are physically transported by the underlying flow. This is explained in detail in Section 3.6.
4. **Anisotropic scaling (tensor mapping):** transform the isotropic synthetic field to match the target Reynolds stress tensor R_{ij} using a linear mapping (e.g. Cholesky factorisation). A comprehensive analysis is provided in Section 3.5.

3.2 Role of Key Turbulence Parameters in SNGR

The primary input from the RANS simulation for scaling the synthetic turbulence is the Turbulent Kinetic Energy (TKE), which sets the overall energy budget of the fluctuations. However, TKE alone is insufficient to fully define a physically realistic turbulent field. The SNGR methodology therefore relies on a hierarchy of statistical inputs to distribute this energy correctly in space, direction, and time:

- **The Reynolds Stress Tensor (R_{ij})** determines how the total turbulent energy is partitioned among the different velocity components ($\langle u'u' \rangle$, $\langle v'v' \rangle$, etc.) and establishes the correlation between them ($\langle u'v' \rangle$). This is essential for capturing the turbulence anisotropy, as discussed in Section 3.5.
- **Integral Length Scales and the Energy Spectrum ($E(\kappa)$)** define the characteristic size of the energy-containing eddies and dictate how the TKE is distributed across different wavenumbers (or frequencies). This ensures the synthetic field has the correct spectral content, as detailed in Section 3.4.
- **Convection Velocity and Temporal Correlation** govern the dynamic evolution of the synthetic field. The mean convection velocity transports the eddies through the domain, while temporal correlation parameters ensure that these eddies persist and decorrelate over physical timescales, creating a realistic time history (Section 3.6).

By combining these parameters, the SNGR method aims to generate a synthetic field that is not only energetically correct (matching TKE) but also structurally and dynamically representative of real turbulence. An incorrect representation of any of these components can lead to a synthetic field with the right total energy but the wrong directional or spectral content.

3.3 Fourier synthesis and divergence-free constraint

The synthetic velocity field is built as a finite Fourier sum of N modes. A compact, divergence-free formulation for the i-th component of the velocity is described in Equation (32).

$$\boldsymbol{u}_i^{\mathrm{iso}}(\mathbf{x},t) \;=\; 2\sum_{n=1}^{N} \hat{u}_n\, \boldsymbol{\sigma}_i^{(n)} \cos\!\big(\boldsymbol{\kappa}_n \cdot \mathbf{x} + \psi_n\big), \tag{32}$$

where $\hat{u}_n$ is the scalar amplitude for mode n, κ_n the wave-vector (with magnitude $\kappa_n = |\kappa_n|$), ψ_n an independent random phase uniformly distributed on $[-\pi, \pi]$, and $\boldsymbol{\sigma}^{(n)}$ the unit direction vector of the velocity for that mode. The divergence-free condition is enforced by construction requiring the expression shown in Equation (33).

$$\boldsymbol{\kappa}_n \cdot \boldsymbol{\sigma}^{(n)} = 0 \qquad \text{for all } n, \tag{33}$$

which can be achieved by choosing an auxiliary vector $\zeta^{(n)}$ and setting the unit direction vector as expressed in Equation (34).

$$\boldsymbol{\sigma}^{(n)} = \frac{\zeta^{(n)} \times \boldsymbol{\kappa}_n}{\left|\zeta^{(n)} \times \boldsymbol{\kappa}_n\right|}. \tag{34}$$

3.4 Modal amplitudes from the energy spectrum

The modal amplitude $\hat{u}_n$ is chosen so that the synthesized field has an integral energy consistent with the one-dimensional energy spectrum $E(\kappa)$. For a discretisation in wavenumber space with shell width $\Delta\kappa_n$, a consistent choice is described in Equation (35).

$$\hat{u}_n \;\propto\; \sqrt{E(\kappa_n)\,\Delta\kappa_n}, \tag{35}$$

since the total variance is recovered by $\sum_n \hat{u}_n^2 \sim \int E(\kappa)\,d\kappa$ Kolmogorov [1941], Pope [2000].

The modal amplitude used in the synthetic field must be normalised so that the modal energy reproduces the target turbulent kinetic energy supplied by the RANS solution. For a discretisation in wavenumber space with shells centred at k_n and width Δk_n the standard form expressed in Equation (36) is adopted.

$$\hat{a}_n \;=\; C\,\sqrt{E(k_n)\,\Delta k_n}, \tag{36}$$

where $E(k)$ is the one-dimensional energy spectrum and C is a scalar normalisation constant chosen such that the modal energy sum matches the target TKE. Using the convention exhibited in Equation (37)

$$\int_0^\infty E(k)\,\mathrm{d}k \;=\; k_{\text{target}}, \tag{37}$$

the constant C is computed as shown in Equation (38).

$$C \;=\; \sqrt{\frac{k_{\text{target}}}{\sum_n E(k_n)\,\Delta k_n}}. \tag{38}$$

In practice, the proportionality constant is set to match the desired turbulent kinetic energy $k = \frac{1}{2}\overline{u_i' u_i'}$ supplied by the RANS solution. If a different spectral convention is used (for example a spectrum defined per velocity component or a one-half factor in the definition of k), an appropriate multiplicative factor must be included in Equation (38).

A phase-advection velocity $\mathbf{U}_c$ is *not* prescribed for the modes. Consequently, modal phases are not advanced using $\omega_n = \boldsymbol{\kappa}_n \cdot \mathbf{U}_c$ and the synthetic fluctuations are not explicitly convected by a chosen convection velocity.

3.5 Anisotropic scaling to match Reynolds stresses

The isotropic field produced above is mapped to the required anisotropy by a linear transformation that enforces the target Reynolds stress tensor R_{ij}. If $\mathbf{u}^{\text{iso}}$ denotes the vector of modal velocity components, the anisotropic field is obtained as

$$\mathbf{u}^{\text{aniso}}(\mathbf{x},t) = \mathbf{L}\,\mathbf{u}^{\text{iso}}(\mathbf{x},t), \qquad \text{with} \quad R = \mathbf{L}\mathbf{L}^{\mathrm{T}}, \tag{39}$$

where $\mathbf{L}$ is the lower-triangular Cholesky factor of R (or any other square root of R). This operation guarantees $\overline{u_i^{\text{aniso}} u_j^{\text{aniso}}} = R_{ij}$ when applied to zero-mean, isotropic unit-variance input fields. In some implementations a correction (e.g. Wilcox near-wall correction) is applied to the RANS-derived R_{ij} prior to decomposition to improve near-wall behaviour Weitz et al. [2019]. Anisotropy plays a very important role in the proper estimation of Reynolds stresses, and it is further explained in Section 2.3.

3.6 Convection and temporal correlation

The canonical SNGR per-mode phase advection and *Ornstein-Uhlenbeck* (OU) amplitude evolution are not implemented in the code used for the results. Instead, modal geometry and random phases are resampled, and finite temporal correlation is introduced by applying an Eulerian advection operator to previously computed fluctuation fields. The equations used in the original SNGR method to update phases between time steps are presented in Equation (40) and Equation (41).

$$\phi_n(\mathbf{x},t) \;=\; \boldsymbol{\kappa}_n \cdot \mathbf{x} - \omega_n(\mathbf{x})\,t \tag{40}$$

$$\omega_n(\mathbf{x}) \;=\; \boldsymbol{\kappa}_n \cdot \overline{\mathbf{U}}(\mathbf{x}) \tag{41}$$

Here, $\phi_n(\mathbf{x},t)$ denotes the total phase of mode n, which depends on both space and time. Its spatial gradient gives the wavevector $\boldsymbol{\kappa}_n$ and its temporal derivative yields the local angular frequency $-\partial_t \phi_n = \omega_n(\mathbf{x})$. The term $\omega_n(\mathbf{x})$ represents a *spatially varying convective frequency field* determined by the local mean velocity $\overline{\mathbf{U}}(\mathbf{x})$. In regions where the mean flow is approximately uniform, $\omega_n(\mathbf{x})$ reduces to a constant scalar value $\omega_n = \boldsymbol{\kappa}_n \cdot \overline{\mathbf{U}}_{\text{ref}}$, where $\overline{\mathbf{U}}_{\text{ref}}$ is a representative convection velocity.

In the implemented SNGR version, this spatially varying frequency field is not computed. Instead, each snapshot is generated independently by resampling modal directions and random phase offsets ψ_n. Consequently, the discrete per-mode phase advection step typically expressed as

$$\psi_n^{m+1} = \psi_n^m - \omega_n \, \Delta t, \tag{42}$$

is not performed. In Equation (42), ω_n is treated as a single scalar convective frequency per mode rather than a spatially varying quantity.

This design choice is motivated by the characteristics of the available CFD input data (discrete instantaneous fields rather than time-resolved sequences) and by the intended validation objectives of this study, which focus on reproducing correct spatial and spectral statistics rather than generating a phase-accurate time evolution of the fluctuations. The simplified approach also avoids the additional storage and computational overhead associated with maintaining per-mode temporal states.

Although using $\omega_n(\mathbf{x})$ from Equation (41) would be recommended when the mean velocity varies significantly across the domain, the present implementation neglects this dependence. For flows with nearly uniform mean convection, this simplification is acceptable, since either the convective phase form $\phi_n(\mathbf{x}, t)$ or the equivalent shifted-coordinate representation $(\mathbf{x} - \overline{\mathbf{U}}t)$ yield algebraically identical behavior.

The *Ornstein–Uhlenbeck stochastic differential equation* (OU SDE) and its exact discrete update are standard and are reproduced here for completeness. The temporal dynamics of the modal amplitude are modelled by the OU stochastic differential equation presented in Equation (43).

$$\frac{da_n}{dt} = -\frac{1}{T_n} \, a_n + \sqrt{\frac{2}{T_n}} \, \xi_n(t), \tag{43}$$

where $\xi_n(t)$ is unit white noise and T_n the decorrelation time of mode n. Equation (43) yields an exponential autocorrelation, $\mathbb{E}[a_n(t)\, a_n(t+\tau)] \propto e^{-|\tau|/T_n}$, and stationary unit variance when driven by unit white noise Schoder [2023]. For a time-discrete implementation (time step Δt) the exact OU update is used to preserve the correct stationary variance and to avoid stiffness issues. If $a_n^m \equiv a_n(m\Delta t)$ then

$$a_n^{m+1} \;=\; a_n^m \, e^{-\Delta t/T_n} \;+\; \sqrt{1 - e^{-2\Delta t/T_n}} \, \eta_n^m, \tag{44}$$

where $\eta_n^m \sim \mathcal{N}(0,1)$ are independent standard normal samples. Equation (44) gives the exact increment for the OU process over a time step Δt. However, this OU SDE and its discrete update were not implemented in the version of the SNGR method developed in this work. Instead, modal amplitudes are generated once per snapshot through the energy spectrum and are not time-integrated via an OU process.

In practice, the contribution of a single mode n to the i-th component of the velocity, $u_i^{mode,n}$, is assembled (for real-valued fields) as a sum of cosines with modal amplitudes set from the energy spectrum and with random phases and mode directions. In the present implementation the factor a_n^t is not integrated in time via an OU update because, as previously explained, modal amplitudes are obtained from the spectrum per snapshot, as exhibited in Equation (45)

$$u_i^{\mathrm{mode},n}(\mathbf{x}, t) \;=\; 2 \, (\hat{u}_n \, a_n(t)) \, \sigma_i^{(n)} \cos\big(\phi_n(\mathbf{x}, t) + \psi_n\big), \tag{45}$$

where $\hat{u}_n$ is the modal amplitude set from the spectrum, ψ_n is a random phase, and $\sigma_i^{(n)}$ is the i-th component of the unit direction vector for that mode.

Furthermore, the relative importance of advection and intrinsic decorrelation for each mode may be defined as shown in Equation (46)

$$\mathcal{C}_n \;=\; \frac{\tau_{\mathrm{adv},n}}{T_n} \;=\; \frac{1}{\kappa_n \, U_{\mathrm{conv}} \, T_n}, \tag{46}$$

where $\tau_{\mathrm{adv},n} \approx 1/(\kappa_n U_{\mathrm{conv}})$ and U_{conv} is a characteristic convection speed. Nonetheless, in teh current implementation shown in this report, this C_n parameter is not explicitly computed to alter modal evolution, since modal-phase advection and OU amplitude updates are not performed.

3.7 Mode time scales and the tuning factor f_τ

Per-mode decorrelation times are related to RANS-derived integral scales and inertial-range scaling. A practical choice is

$$T_n \;=\; f_\tau \, T_{\mathrm{int}} \left(\frac{\kappa_n}{\kappa_{\mathrm{ref}}} \right)^{-2/3}, \tag{47}$$

where $T_{\mathrm{int}} = \frac{k}{\varepsilon}$ is the integral time scale obtained from RANS, $\kappa_{\mathrm{ref}} = \frac{2\pi}{L_{\mathrm{int}}}$ is the reference wavenumber, and f_τ is a dimensionless tuning factor used to adjust temporal coherence. The exponent $-2/3$ follows Kolmogorov scaling in the inertial range Kolmogorov [1941], Pope [2000].

3.8 Amplitude normalization and TKE matching

Modal amplitudes are initially set by Eq. (35). After assembling the isotropic field u_i^{iso}, it is recommended to compute the generated turbulent kinetic energy

$$k_{\mathrm{gen}} = \frac{1}{2}\overline{u_i^{\mathrm{iso}} u_i^{\mathrm{iso}}},$$

and apply a global (or local) scaling factor

$$\gamma = \sqrt{\frac{k_{\mathrm{target}}}{k_{\mathrm{gen}}}}$$

so that $\gamma\, u_i^{\mathrm{iso}}$ has the target TKE k_{target} from the RANS field. This compensates for sampling and discretisation bias in the finite-mode synthesis.

3.9 Diagnostics and validation

Key diagnostics to report in the thesis are:

- Comparison of generated one-dimensional spectra with input $E(\kappa)$ and inertial-range $-5/3$ scaling where applicable Kolmogorov [1941], Pope [2000].
- TKE match (target vs generated) and residual mean component (verifying zero-mean fluctuations).
- Reynolds-stress recovery after anisotropic mapping: $\overline{u_i^{\mathrm{aniso}} u_j^{\mathrm{aniso}}} \stackrel{?}{=} R_{ij}$.

These checks are used both for parameter tuning (N, f_τ) and for reporting SNGR uncertainties versus DNS baselines.

3.10 Computational Implementation and Performance

The SNGR pipeline is implemented in MATLAB and designed to leverage high-performance computing (HPC) resources where possible. It is important to clarify how the model utilizes the available hardware, as its performance characteristics differ between the two main SNGR approaches.

The MATLAB implementation leverages the Parallel Computing Toolbox, specifically using `parfor` loops to distribute the computation of the **ensemble-averaged statistics**. Since each of the M statistical realizations is an independent calculation, this part of the analysis scales effectively with the number of available cores.

However, the **time-marching (single-loop) simulation** is inherently sequential, as each timestep depends on the result of the previous one. This loop therefore runs on a single core. While MATLAB's built-in multi-threading may accelerate individual matrix operations (such as those involving linear algebra) within this loop, the overall time-marching process does not benefit from the massively parallel architecture in the same way the ensemble calculations do. This sequential dependency is the primary bottleneck for the time-marching simulations.

3.11 Methodology for CGNS/HDF5 data ingestion and preprocessing

This section describes the conceptual processing steps applied when importing CFD data from a CGNS/HDF5 file into MATLAB for subsequent analysis, focusing on the data objects and the transformations performed on them.

The CFD General Notation System (CGNS) is a standardized data format and set of conventions for the storage and exchange of computational fluid dynamics (CFD) data. It was originally developed by NASA in collaboration with industry partners to provide a vendor-neutral, self-describing file format that facilitates interoperability between different CFD solvers, post-processing tools, and long-term archival of simulation results CGNS Steering Committee [2022], Acuña et al. [1999].

CGNS is typically implemented on top of the Hierarchical Data Format (HDF5), which provides a portable and efficient binary container capable of storing large structured or unstructured mesh datasets together with associated flow solutions. Within a CGNS file, the mesh geometry, boundary conditions, solution fields, and auxiliary metadata are stored hierarchically following the CGNS Standard Interface Data Structures (SIDS). This ensures that any compliant tool can navigate the file, extract information in a consistent manner, and map it into its own data structures without relying on case-specific export scripts.

In this work, MATLAB serves as the environment in which CGNS/HDF5 data is read, transformed and processed into the pipeline's internal structures. MATLAB is chosen because of its robust support for HDF5 I/O, efficient array operations, and ease of integration with custom diagnostics and turbulence-processing routines developed for the SNGR method. These aspects guarantee the proper handing of large CFD datasets. Furthermore, the MATLAB code is executed on a high-performance computing (HPC) cluster, where MATLAB's built-in parallelization tools (e.g. `parpool`, `parfor`) are employed to reduce execution time and scale the workflow to large meshes.

In practice, CGNS files enable the seamless transfer of CFD data from solvers to analysis environments such as MATLAB. By adhering to the standardized SIDS definitions, fields such as velocity, pressure, turbulence quantities, or Reynolds stresses can be located reliably and compared across simulations. For large-scale studies, this guarantees reproducibility and reduces the risk of errors arising from inconsistent file conventions.

The procedures described in the remainder of this subsection correspond to the actual MATLAB CGNS/HDF5 ingestion and preprocessing pipeline implemented for this work. All major components listed below have been implemented and verified on the datasets used in this study. Where a behaviour is conditional (for example the Boussinesq linear-eddy-viscosity assembly), the code performs the fallback automatically at runtime and records the event in the read-diagnostics log.

3.11.1 Geometry and mesh reconstruction

The first step is to reconstruct the computational geometry and its discrete representation:

1. **Vertex coordinates:** The file provides global vertex coordinates as three arrays (for the x, y, z components). These are read and assembled into an $N_v \times 3$ array of vertex coordinates.

2. **Cell centroids and volumes:** Per–cell centroid coordinates and cell volumes are read into an $N_c \times 3$ centroid array and an N_c-length volume vector. All subsequent per-cell fields are expected to align with this centroid ordering.

3. **Face and cell topology:** Face connectivity is reconstructed by reading face vertex indices and assembling a face list. The cell–face incidence relation (which cells touch each face) is computed from the face connectivity. From that, face centroids are computed by averaging the coordinates of the face vertices.

4. **Orientation and ordering:** The imported face and cell indexing is normalized to conform to the pipeline's internal conventions (consistent ordering, orientation). Any geometry offsets required to convert between coordinate systems (for example a vertical offset applied to match experimental reference) are applied explicitly and documented as a separate geometric transformation.

3.11.2 Primary flow fields and derived quantities

Once the mesh has been processed in MATLAB, the MATLAB-based data-processing pipeline imports per-cell flow fields (velocity components, pressure, turbulence quantities, etc.), aligns them with the internal centroid ordering, and computes the principal derived quantities used throughout the analysis:

- **Velocities and pressure:** Per–cell velocity components and pressure (or pressure coefficient) are loaded into arrays aligned to the centroid ordering.

- **Viscosity and density:** Dynamic viscosity and density are read. The kinematic viscosity is computed elementwise as defined in Equation (48).

$$\nu = \frac{\mu}{\rho} \tag{48}$$

  These values are stored for use in nondimensionalization and viscous computations.

- **Turbulence quantities:** Turbulent kinetic energy (TKE) and turbulent (eddy) viscosity are read when available and stored per cell. These quantities are used in model fallbacks and in diagnostics (for example to compute friction velocity).

3.11.3 Reynolds–stress reconstruction: hierarchy of strategies

Comparisons and anisotropy corrections require a full Reynolds–stress tensor field per cell. The ingestion routine assembles this tensor using a robust, multi-tiered strategy (use the richest available data and gracefully fall back when fields are absent):

1. **Direct tensor read:** If the file contains a packed nine-component stress dataset (one 3×3 tensor per cell), it is reshaped into an $N_c \times 3 \times 3$ array and used directly. This is the case when considering a turbulence model such as the RSM-EB model.

2. **Six independent components:** If the file provides separate datasets for the six unique symmetric stress components (e.g. $R_{11}, R_{22}, R_{33}, R_{12}, R_{13}, R_{23}$), these are combined into the symmetric 3×3 tensor at each cell.

3. **Model closure fallback (Boussinesq approximation):** If explicit Reynolds stresses are not directly available, an approximate tensor is constructed from turbulent kinetic energy k values, turbulent viscosity ν_t, and the rate-of-strain components S_{ij} that were generated using field functions in *StarCCM+*, using the linear eddy-viscosity (Boussinesq) closure described in Equation (49) Boussinesq [1897], Tennekes and Lumley [1972a], Schmitt [2007].

$$R_{ij} \approx \frac{2}{3} k \, \delta_{ij} - 2 \, \nu_t \, S_{ij} \tag{49}$$

where $S_{ij} = \frac{1}{2} (\partial_j U_i + \partial_i U_j)$ and δ_{ij} is the Kronecker delta (i.e. $\delta_{ij} = 1$ if $i = j$, and $\delta_{ij} = 0$ otherwise). This strategy is explicitly a fallback and should be considered an approximation suitable for diagnostics rather than a replacement for explicit Reynolds stresses in high-accuracy comparisons.

3.11.4 Wall shear stress handling

Wall shear stress data can be stored with various naming or component conventions; the ingestion routine therefore:

- Attempts to locate wall shear stress components exported by the solver, matching multiple naming variants.
- If components are found, they are read and associated with face geometry (or mapped to cell quantities if needed).
- If no wall shear stress data are present, a zero field is provided and a warning is issued. A missing wall shear stress dataset disables friction-velocity estimates derived from direct wall stress and prompts alternative computations if needed (e.g., viscous gradient approximations).

3.11.5 Data consistency checks and diagnostics

To catch common export or read issues early, the routine performs a set of automated sanity checks and writes concise diagnostics:

1. **Shape consistency:** Verify that each per-cell field has N_c entries. If a mismatch is detected, the pipeline flags the field and either aborts (strict mode) or issues a clear warning (diagnostic mode).
2. **Physical checks:** Compare turbulent kinetic energy to the diagonal components of the stress tensor. Two hypotheses are tested depending on the convention used in the file:

$$\text{if stresses are variances:} \quad k \stackrel{?}{=} \tfrac{1}{2} (R_{11} + R_{22} + R_{33}),$$
$$\text{if raw RMS values were exported instead of variances:} \quad k \stackrel{?}{=} \tfrac{1}{2} \left((u')^2 + (v')^2 + (w')^2 \right).$$

3. **Range/histogram outputs:** Print summary statistics (min, max, mean) for key fields (e.g., R_{11}, R_{22}, R_{33}, TKE) and location indices of extreme values to make spot checks easier.
4. **Naming and provenance log:** Record which dataset names were actually used for each logical field (this log is invaluable when the same pipeline runs on files produced by different solver configurations).

3.11.6 Interpolation and topology robustness considerations

The pipeline depends on per-cell quantities being ordered consistently with the centroid array. Practical robustness measures include:

- **Topology validation:** Verify that face centroids computed from vertex coordinates are consistent with any face-based datasets; report if a per-face dataset appears to be per-cell or mismatched.
- **Memory and scaling controls:** For very large meshes, avoid unnecessary copying of arrays and provide options to read only required subsets (or to reduce sampling resolution for diagnostics).
- **Explicit coordinate transforms:** Any geometry offsets or coordinate system conversions are made explicit, parameterized and recorded in metadata rather than applied implicitly.

3.11.7 Common failure modes and recommended fixes

Importing a new CGNS/HDF5 dataset from StarCCM+ may result in some common issues:

- **Missing fields:** Some solver exports omit Reynolds stresses or strain-rate fields, which are necessary in some cases to run the full pipeline without crashing. Reconfiguring the export to include the needed quantities or accept the Boussinesq fallback only for diagnostic purposes solves this problem.
- **Dataset naming variability:** Different solver versions and export scripts use different dataset names. Thus, the generated diagnostics listing must be checked, and the export or ingestion mapping must be adjusted, keeping a short mapping table in project documentation.
- **Shape mismatch (per-face vs per-cell):** Datasets may be written per face rather than per cell. Therefore, it is highly recommended to detect mismatches early and either map face data to cell centers or request a per-cell export.
- **Coordinate offsets:** Implicit coordinate shifts can lead to misalignment with other data (e.g., experimental reference), and may happen as a result of considering a different (0,0) coordinate in MATLAB compared to the geometry from the CFD solver. This can be avoided by making any shifts explicit with parameters and log them.
- **Sign conventions for strain rates:** When using the closure fallback, check the sign convention of strain-rate components with a small test case.

3.11.8 Practical recommendations for reproducible ingestion

To make the ingestion stage reliable across datasets and future users, the following best practices are recommended:

1. Provide a small "read diagnostics" report file with each run listing found datasets, shapes, ranges and any fallbacks applied.
2. Offer a strict validation mode that aborts on any size or unit mismatch, suitable for automated testing; and a permissive diagnostics mode suitable for exploratory work.
3. Document the expected dataset names and units that should be enabled on the CFD export side so that future exports are immediately compatible.
4. Use explicit metadata attributes (where available) to capture units and mesh provenance, and convert units to the pipeline's internal standard at read time.

4 Description of validation cases

As has been previously explained, it is important to compare the current model with reliable turbulent data. As the highest accuracy is achieved by solving all turbulent length scales, DNS studies are used to identify sources of error and uncertainties using the SGNR method and eventually improve it.

4.1 Turbulent Channel Flow (TCF)

Turbulent Channel Flow (TCF) is an academic case that cannot be directly extended to industrial applications. It represents an idealized turbulent flow between two parallel plates driven by a mean pressure gradient parallel to the walls. The flow is statistically homogeneous in both the streamwise (x) and spanwise (z) directions. This statistical homogeneity implies that, when averaged over time in these directions, there is no net transport of momentum in the x and x plane, resulting in a zero covariance of streamwise and spanwise fluctuations. The Direct Numerical Simulation (DNS) by Kim et al. Kim et al. [1987] is a foundational high-fidelity benchmark widely used for validating computational models like SNGR and RANS-based reconstructions. The simulation modeled fully developed turbulent channel flow at a friction Reynolds number of 180, which was based on the wall shear velocity .

While the overall turbulence statistics from the simulation align well with experimental data, detailed comparisons reveal consistent discrepancies, particularly in the near-wall region . Specifically, the DNS predicts Reynolds stress components—both normal and shear—that are consistently lower than those from historical experimental measurements. Conversely, the simulation shows higher levels of small-scale vorticity fluctuations near the wall compared to the experimental values. These discrepancies, which are commonly observed for low Reynolds number channel comparisons, likely result from a combination of experimental uncertainty in the oil-channel measurements, minor mismatches in Reynolds number and normalization, and differences in near-wall grid resolution and data processing. Similar deviations between DNS and experiments have been reported by Eckelmann [1974], Moin and Mahesh [1998a], and Schlatter and Örlü [2010], emphasizing that measurement noise and spatial resolution strongly affect the near-wall turbulence intensities. Consequently, while Kim et al. remains a primary DNS benchmark for spectral and scaling validation, conclusions based on absolute near-wall amplitudes should be qualified with quantitative error measures (e.g. percent error in specified y^+ bands) and and supported by sensitivity analysis.

Furthermore, most of the experimental datasets used for comparison with the values obtained from the DNS analysis at these low Reynolds numbers was sourced from the oil channel at the Max-Planck-Institut für Strömungsforschung in Göttingen, West Germany Max Planck Institute for Dynamics and Self-Organization [2025].

4.1.1 Direct Numerical Simulation of Turbulent Channel Flow up to $Re_\tau \approx 5200$

A comparison to a DNS study of turbulent channel flow has been performed for friction Reynolds numbers up to $Re_\tau \approx 5200$ Lee and Moser [2015]. In this report, five DNS cases of fully developed flow were analyzed at different friction Reynolds numbers: $Re_\tau \approx 180, 550, 1000, 2000$ and 5200. This paper was selected to evaluate the accuracy of the SNGR method due to its comprehensive introduction, well-defined methodology, extensive set of turbulence statistics and spectra, and the public availability of the corresponding DNS data through the Texas Advanced Computing Center (TACC) Turbulence Database maintained by the University of Texas at Austin Lee and Moser [2015].

In the present work, the SNGR analysis focuses on the lowest Reynolds-number case ($Re_{tau} = 180$ from the dataset of Lee and Moser [2015], which corresponds to the canonical configuration first analyzed by Kim et al. [1987]. The bulk Reynolds number is estimated as seen in Equation (50).

$$Re_b = \frac{U_b \cdot 2\delta}{\nu} \tag{50}$$

where:

- U_b = average streamwise (bulk) velocity, set to 1.000 m/s.
- δ = channel half-width, defined as 1.000 m.
- ν = kinematic viscosity, specified as 5×10^{-5} m^2/s.

An evaluation was also performed assuming $Re_\tau = 1000.512$ and $Re_b = 40000$. The number of grid cells in the streamwise (N_x), wall-normal (N_y), and spanwise (N_z) directions are 2304, 512, and 2048, respectively. The domain size (L) in the directions x, y, and z is set relative to π, so that:

$$L_x = 8\pi, \quad L_y = \pi, \quad L_z = 3\pi.$$

For the case considering $Re_\tau = 1994.756$, the grid resolution is set to $N_x = 4096$, $N_y = 768$, and $N_z = 3072$. These values are significantly large and would increase the computational time of the SNGR model to more than

a day using the high-performance computing system (with up to 192 cores). As detailed in Section 3.10, the time-marching simulation is an inherently sequential process and does not scale with the number of available cores, making it the primary computational bottleneck. Therefore, this high-resolution case is not computationally feasible for this study.

Similarly, the study considering $Re_\tau = 5185.897$ used a number of grid cells of $N_x = 10240$, $N_y = 1536$, and $N_z = 7680$.

It is important to note that the DNS of Lee and Moser [2015] resolves all turbulent scales without employing any turbulence model. In contrast, the RANS simulations used in the present study rely on the Cebeci–Smith eddy-viscosity formulation Cebeci and Smith [1974], which combines Prandtl's mixing-length theory Prandtl [1925] in the inner layer, as shown in (51).

$$\ell = \kappa y \left(1 - e^{-y^+/A^+}\right) \tag{51}$$

where ℓ is the local mixing length, κ is the von Kármán constant, and A^+ is an empirical damping constant. Clauser's eddy-viscosity formulation is then applied in the outer layer. This hybrid model provides a reasonable approximation of near-wall behavior but is limited compared to two-equation transport models (such as the k–ϵ model) described in Section 2.1.3.

4.1.2 StarCCM+

As the main goal is quantifying the error and uncertainties of the SNGR method based on the RANS equations, simulations are performed using the StarCCM + software, operating using a High Performance Computing (HPC) cluster, reducing computational time. To ensure statistical homogeneity in the x- and z-directions, the geometry is created in the plane formed by the streamwise direction (x) and the wall-normal direction (y). Once the 2D computational domain is created on this x-y plane, it is extruded along the spanwise direction (z). The physical models are set to match the DNS study: the flow is modeled as air, the simulation is steady (as RANS solves for time-averaged quantities), and the domain is three-dimensional. A 3D domain is essential not to resolve instantaneous fluctuations, but to allow for the computation of the three-dimensional turbulence statistics (e.g., the full Reynolds stress tensor) that serve as input for the SNGR method. Furthermore, the flow is considered to be segregated, solving momentum and pressure equations from the Navier-Stokes equations separately in an iterative manner - it works well for incompressible and moderately turbulent flows.

Additionally, the equation of state is set to constant density, as the flow analyzed in the DNS study is incompressible. Finally, the viscous regime is defined as turbulent, and no transition model is chosen as the flow is fully turbulent at $Re_\tau = 180$. Finally, on the top and bottom walls (perpendicular to the wall normal direction) the no-slip condition is applied, as they represent the solid channel boundaries where viscous effects are dominant, defining the other boundaries with the slip condition.

The flow studied in the DNS analysis is driven by a uniform pressure gradient, which varies in time to ensure that the mass flux through the channel remains constant, meaning that there is no explicit inlet velocity, but instead, the flow is fully developed due to the applied pressure gradient. Thus, in StarCCM+ a constant pressure gradient force is applied in the streamwise direction by creating a Fully developed interface between the Inlet and Outlet, describing the periodicity as Translational along the streamwise axis, and setting the Pressure Jump to 2.49×10^{-3}. Similarly, in the spanwise direction (z-direction), a Fully developed interface is set, with periodic topology along Z axis, and no pressure jump value specified.

Mesh generation

Meshing is performed with the Trimmed Cell Mesher and Prism Layer Mesher. The base size is set to 0.0075 m, considering a target surface size is 100% of base, and prism layers are specified with:

- Number of prism layers: 15
- Stretching factor: 1.2
- Total prism thickness: 180% of base size. This is 0.0135 m

The *maximum core/prism transition ratio* is disabled, *growth rate* is set to slow in volume and disabled at surfaces, and *maximum cell size* is limited to 100% of the base value. Besides, *post-mesh optimization* is enabled (boundary vertex optimization and cell topology optimization).

This configuration yields a high-resolution mesh around the step and near-wall regions, ensuring that the recirculation and reattachment zones are adequately resolved. The mesh used to get the most accurate results of this report is shown in Figure 4.

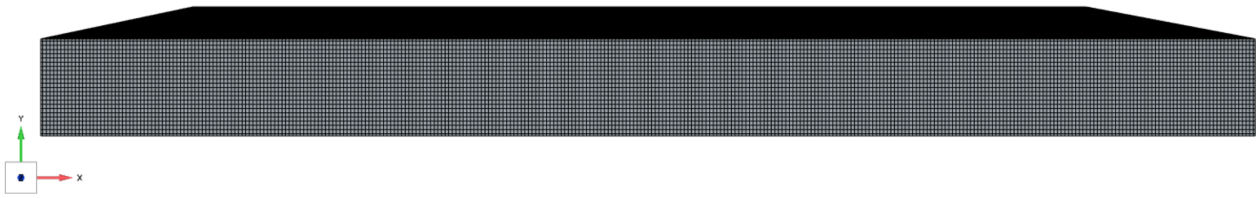

(a) Streamwise and wall-normal view of the refined mesh considered for the TCF.

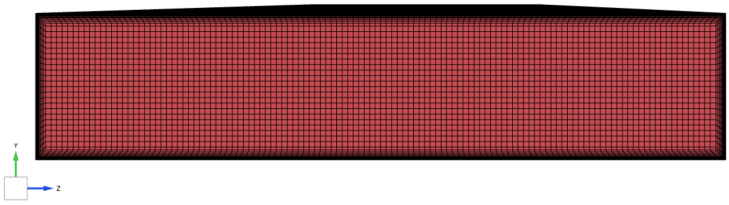

(b) Wall-normal and spanwise view of the refined mesh.

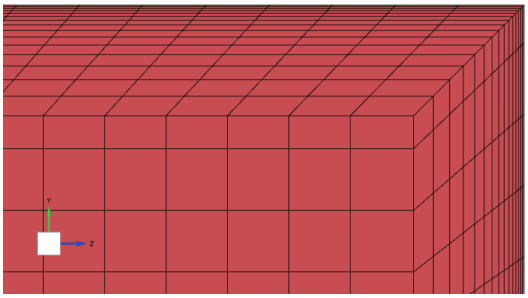

(c) Close look of the mesh along the wall-normal and spanwise directions of the refined mesh.

Figure 4: Mesh views of the TCF simulation

4.2 Backward facing step (BFS)

Recirculating eddies produced by flow separation strongly influence shear stresses and heat transfer, which makes the backward-facing step a useful canonical test for turbulent separated flows Armaly et al. [1983], Chiang and Sheu [1999]. The previous analysis considering turbulent channel flow offered a limited perspective regarding this aspect. As a subject of great importance in fluid mechanics - including fluid-structure interaction - it is investigated by simulating a flow over a backward-facing step shown schematically in Figure 5.

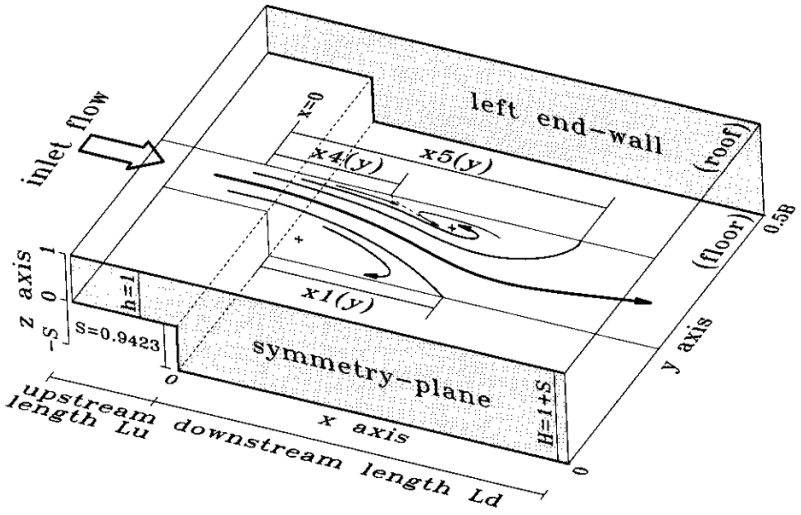

Figure 5: Schematic of laminar flow over a backward-facing step Chiang and Sheu [1999].

The recirculation zone is described in Figure 5 shows the inlet direction and symmetry conventions used in DNS studies. Besides, separation and reattachment lengths are defined as:

$$x_1 = \text{Primary reattachment length} \tag{52}$$
$$x_4 = \text{Separation length} \tag{53}$$
$$x_5 = \text{Secondary reattachment length} \tag{54}$$

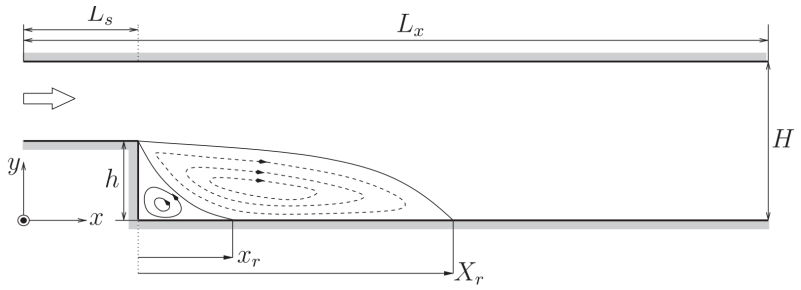

Figure 6: 2D Schematic of flow over a generic backward-facing step Barri et al. [2010].

The simplicity of this geometry makes it an ideal configuration for examining and comparing the behavior of vortical flows. In actual applications, the recirculation zone may be produced by a strong adverse pressure gradient, although the topology of the flow is similar to the flow separation caused by a geometric variation. Barkley et al. [2002] The availability of extensive experimental and simulation data has made this a standard test case for numerical validation. As shown in Figure 6, the separation length x_4 characterizes the upstream recirculation, while x_1 and x_5 mark critical reattachment points downstream. These lengths are defined as:

$$X_r = \text{Primary reattachment length} \tag{55}$$
$$x_r = \text{Separation length} \tag{56}$$

All reattachment positions reported later are given in terms of x/h where h is the step height, so comparisons with DNS and experimental references are straightforward.

4.2.1 Direct Numerical Simulation of Turbulent Flow over a Backward-Facing Step 1997

The DNS study by Le et al. [1997] is chosen because of the detailed methodology, clear explanation of the boundary conditions, multiple graphs of velocity and pressure fluctuations, turbulence intensities, Reynolds stresses, and extended energy plots. Furthermore, it is highly cited, and the authors have a long career in the field, being very successful. Previously, Moin and Le had worked on a similar analysis Le and Moin [1992].

4.2.2 StarCCM+

The computational domain consists of a streamwise length $L_x = 30h$, including an inlet section $L_s = 10h$ prior to sudden expansion, a vertical height $H = 6h$, and a spanwise width $L_z = 4h$, where h is the step height, which is set at 5 cm.

The physical models used in the *StarCCM+* simulation are constant density (the DNS study assumes incompressible flow), turbulent viscous regime, and a coupled pressure-velocity solution strategy. The coupled solver was selected for the BFS because separated flows exhibit a strong mutual dependence between velocity and pressure: steep pressure gradients, reverse-flow regions and intense shear produce rapid, nonlinear feedback between the momentum equations and the continuity constraint, which coupled solvers handle more robustly than segregated algorithms. This is especially important for accurately capturing turbulence and boundary layer dynamics). Furthermore, the turbulence closure is the Reynolds-averaged k-ω SST, the simulation is steady-state, and gas is chosen as the fluid to be evaluated in the computational domain.

The selection of the equation solver differs from the turbulent channel-flow case, which is fully developed and periodic and therefore exhibits weaker local pressure–velocity coupling. For this reason the BFS simulations were performed using the coupled pressure–velocity solver (simultaneous solution of momentum and continuity), whereas the turbulent channel-flow (TCF) validation used a segregated solver, as explained in Section 4.1.2..

In addition, a reference pressure of 101325.0 Pa is set as a standard convention (corresponding to regular atmospheric conditions). Since the DNS study models an incompressible flow, only pressure gradients affect the flow dynamics, making the absolute reference pressure arbitrary. The fluid density is set to 1.225 kg/m^3. This value is not specified in the reference paper; however, for an incompressible simulation defined by a Reynolds number based on kinematic viscosity ($Re_h = U_0 h/\nu$), the specific choice of density is not dynamically critical, provided the resulting kinematic viscosity matches that of the DNS study.

The inlet velocity profile used for the simulation is shown in Figure 7, adapted from Spalart (1996) Spalart [1986]. This profile represents the turbulent boundary layer upstream of the backward-facing step and serves as the inflow condition for the present study.

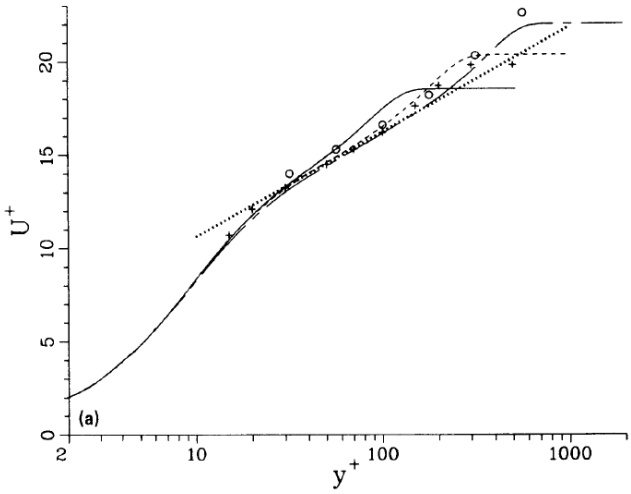

Figure 7: Inlet velocity profile used for the backward-facing step simulation, based on Spalart Spalart [1986], representing the turbulent boundary layer profile upstream of the step.

With respect to the boundary conditions, the upper boundary of the computational domain is set as a no-stress wall, which is a direct match to the boundary conditions specified in the DNS study Le et al. [1997]. To implement this, the upper boundary is defined as a slip wall in StarCCM+, which enforces the velocity conditions shown in Equation (57).

$$v = 0, \quad \frac{\partial u}{\partial y} = \frac{\partial w}{\partial y} = 0. \tag{57}$$

The flow is assumed to be statistically homogeneous in the spanwise direction, where periodic boundary conditions are imposed. This is done by creating an interface between the two parallel boundaries extending on the x-y plane,

defining this type of interface as fully-developed, considering a periodic topology, as well as setting the periodicity as translational (because of the statistic homogeneity in the spanwise direction of the flow, repeating the structures without rotation), and specifying the axis direction as [0.0, 0.0, 1.0]. Besides, the pressure jump is set to 0 Pa.

With respect to the inlet, the mean velocity profile is extracted from the direct numerical simulation of a turbulent boundary layer by Spalart Spalart [1986]. This choice is a direct match with the DNS study by Le et al. [1997], which used Spalart's data for a momentum-thickness Reynolds number of $Re_\theta = 670$. This specific profile was selected by Le et al. because it results in a boundary layer thickness at the step of $\delta_{99} = 1.2h$, matching the conditions of their experiments. As shown in Figure 7, for this Reynolds number, the profile closely follows the logarithmic law, as described in Equation (58).

$$U^+ = \frac{1}{\kappa}\log(y^+) + C \tag{58}$$

with constants $\kappa = 0.41$ and $C = 5$, starting from $y^+ \approx 30$. Further details on the profile's implementation can be found in Appendix C.

Furthermore, the outlet is defined as a pressure outlet, and no-slip boundary conditions are applied at all solid walls (both the bottom surface and the step face). These conditions are also consistent with the setup described by Le et al. [1997].

To assess the mean pressure field, the pressure coefficient (C_p) along the bottom wall was extracted from the CFD simulation. The C_p is a dimensionless quantity that describes the relative pressures throughout a flow field and is defined as:

$$C_p = \frac{p - p_\infty}{\frac{1}{2}\rho_\infty U_\infty^2} \tag{59}$$

where p is the local static pressure, and p_∞, ρ_∞, and U_∞ are the reference pressure, density, and velocity in the freestream, respectively.

Figure 8 shows the pressure coefficient (C_p) along the bottom wall, downstream of the flow reattachment. The profile exhibits three distinct regions:

- **Pressure Recovery ($10 < x/h \lesssim 18$):** The C_p rises sharply as the reattached flow decelerates, converting kinetic energy into static pressure.
- **Plateau Region ($18 \lesssim x/h \lesssim 26$):** The pressure reaches a broad peak, indicating the flow is approaching a redeveloped, equilibrium state.
- **Outlet Influence ($x/h \gtrsim 26$):** The final pressure drop is a non-physical numerical artifact. It is caused by the influence of the pressure outlet boundary condition, which enforces a fixed pressure at the domain exit and creates an artificial gradient that propagates a short distance upstream.

The small, final uptick right at $x/h = 30$ is also part of this numerical artifact and is not a physical phenomenon. It is likely a localized effect from the interaction of the last computational cell with the boundary.

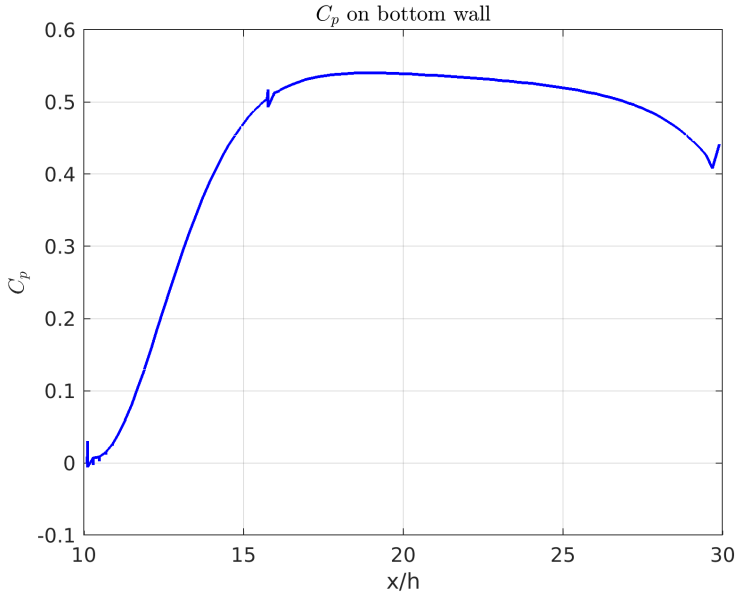

Figure 8: Pressure coefficient (C_p) distribution along the bottom wall of the backward-facing step, as predicted by the CFD simulation.

Mesh generation

Meshing is performed with the Trimmed Cell Mesher and Prism Layer Mesher. The base size is set to 0.008 m. The primary objective of the meshing strategy is to accurately resolve the turbulent boundary layer, particularly the viscous sublayer near the wall. To achieve this, a set of 15 prism layers with a stretching factor of 1.2 is applied to all wall boundaries. This configuration is specifically designed to ensure that the non-dimensional wall distance of the first grid cell, the y^+ value, remains low. As shown in Table 3, This configuration results in a y^+ value for the first grid cell that varies along the wall, ranging from a **minimum of approximately 1.2** in the high-shear region near the step corner to a **maximum of approximately 8.5** in regions of lower shear stress. The higher y^+ value at the inlet (~ 4.5) is expected due to the lower wall shear stress in the upstream, fully-developed boundary layer.

While the minimum value approaches the ideal target, the presence of regions where $y^+ > 5$ indicates that the viscous sublayer is **not fully resolved** across all wall boundaries. This limitation is a likely source of inaccuracies in the RANS prediction of wall shear stress and near-wall turbulence production, which in turn could contribute to discrepancies in the SNGR results.

Additional settings were used to ensure high mesh quality. The *maximum core/prism transition ratio* is disabled, *growth rate* is set to slow in volume and disabled at surfaces, and *maximum cell size* is limited to 100% of the base value. Besides, *post-mesh optimization* (boundary vertex optimization and cell topology optimization) is enabled.

Table 3 compares the grid resolution of the current RANS simulation with that of the reference DNS by Le et al. [1997]. While the RANS mesh is understandably coarser in the streamwise (Δx^+) and spanwise (Δz^+) directions, the critical near-wall resolution (Δy^+_{min}) is of a comparable order of magnitude. This confirms that the mesh is well-suited for capturing the dominant wall-bounded physics required for this validation study.

Table 3: Comparison of non-dimensional grid spacing in wall units between the current RANS mesh and the reference DNS.

Parameter	**Current RANS**	**Le et al. (1997) DNS Le et al. [1997]**
Δx^+	≈ 31.3	≈ 10
Δy^+_{inlet}	≈ 4.5	≈ 0.3
Δy^+_{wall} (min/max)	≈ 1.2 / 8.5	≈ 0.3
Δz^+	≈ 31.3	≈ 15

This configuration yields a high-resolution mesh around the step and near-wall regions, ensuring that the recirculation and reattachment zones are adequately resolved. An overview of the mesh used to get the most accurate results of this report is shown in Figure 9.

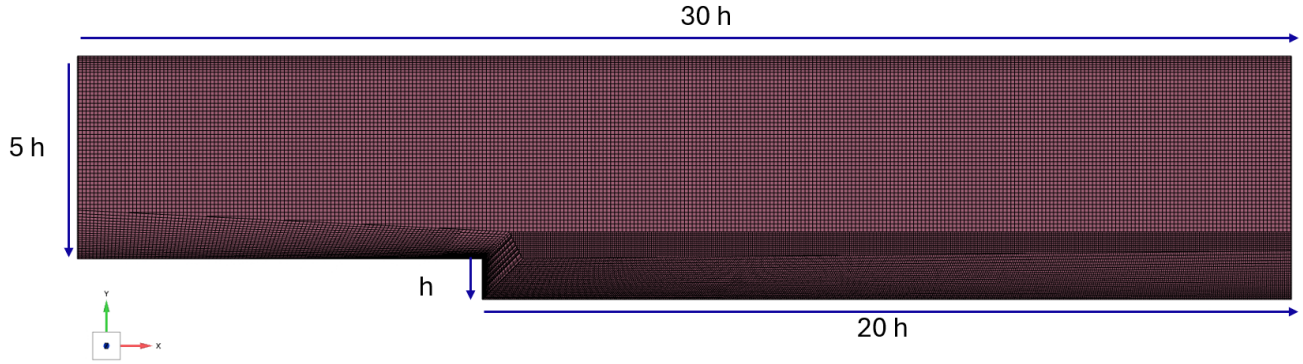

(a) Streamwise and wall-normal view of the refined mesh considered for the BFS.

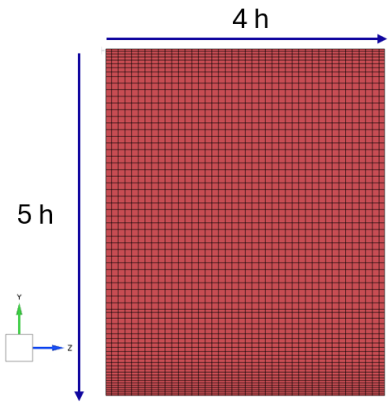

(b) Wall-normal and spanwise view of the refined mesh.

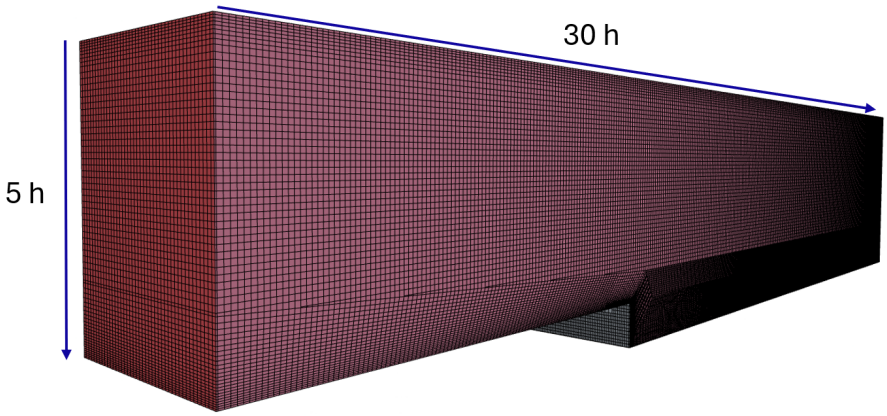

(c) Three-dimensional view of the refined mesh.

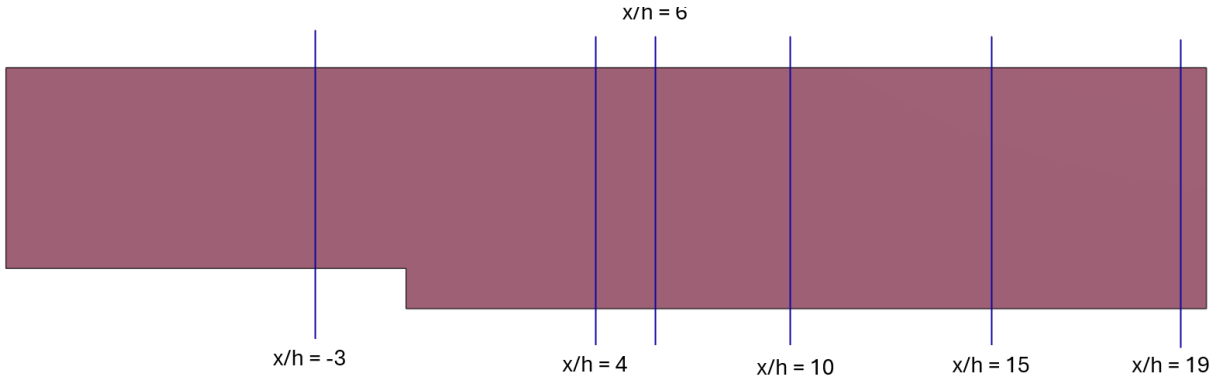

(d) Location of the six slices evaluated in the DNS study and compared to CFD and SNGR results.

Figure 9: Overview of the refined BFS mesh and slice locations: (a) streamwise/wall-normal, (b) wall-normal/spanwise, (c) 3D view, (d) slice locations used in comparisons.

5 Results

This chapter presents the quantitative and qualitative validation of the SNGR method. The analysis is performed on the Backward-Facing Step (BFS) validation case, comparing the SNGR-generated results against both the high-fidelity DNS reference data and the input RANS solution obtained from StarCCM+.

The validation is structured as follows: First, a detailed examination of the second-order turbulence statistics is conducted - in Section 5.1 - through comparisons of the Reynolds stress tensor components ($\langle u_i' u_j' \rangle$). This is followed by an analysis of the mean flow field, comparing streamwise velocity profiles in both physical units (U) - described in Section 5.2 - and non-dimensional wall units (U^+ vs. y^+) - shown in Section 5.3. Subsequently, the Turbulent Kinetic Energy (TKE) is evaluated to assess the overall energy content of the reconstructed fields under various mesh and inlet conditions, as described in Section 5.4. The frequency content of the synthetic turbulence is then analyzed via Power Spectral Density (PSD) plots to verify correct spectral scaling, including them in Section 5.5. Finally, the chapter concludes with a qualitative visualization of the instantaneous turbulent fluctuation fields, illustrating the dynamic behavior of the reconstructed eddies in Section 5.6. For all these comparisons, data corresponding to the RSM EB model are always shown first, followed by the plots using RANS k-omega SST as the turbulence model.

5.1 Reynolds Stresses Comparisons

The importance of the proper estimation of the Reynolds Stresses lies in its quantitative link between the turbulent fluctuations generated and the mean flow. As they control momentum transport and shear, the Reynolds Stresses are the primary drivers of the flow-induced vibrations that are the end goal of the SNGR quantification method. Thus, they are the most relevant second-order statistics for validating the performance of the SNGR method against the baseline data within this analysis of the BFS model.

5.1.1 Reynolds Stresses Comparison: CFD vs. DNS

Before evaluating the SNGR reconstruction, it is essential to first establish a baseline by comparing the RANS simulation results directly against the DNS data. This comparison quantifies the accuracy of the CFD model itself, which serves as the input for the SNGR method. Any discrepancies identified here are inherent to the RANS turbulence model (RSM-EB in this case) and will inevitably influence the final SNGR output. The following figures compare the Reynolds stress components from the corresponding CFD simulation against the DNS reference.

This analysis directly addresses **Research Question 2 (RQ2)**, which investigates how accurately the SNGR code can replicate DNS reference data for turbulent internal flows. By establishing the deviation between CFD and DNS, the uncertainty introduced by the RANS turbulence modeling is quantified, forming the baseline for evaluating the SNGR reconstruction in subsequent sections.

Results for the Elliptic Blending Reynolds Stress Model

Figure 10 shows the Reynolds stresses profiles along the streamwise direction having considered as turbulence model the RSM Elliptic Blending.

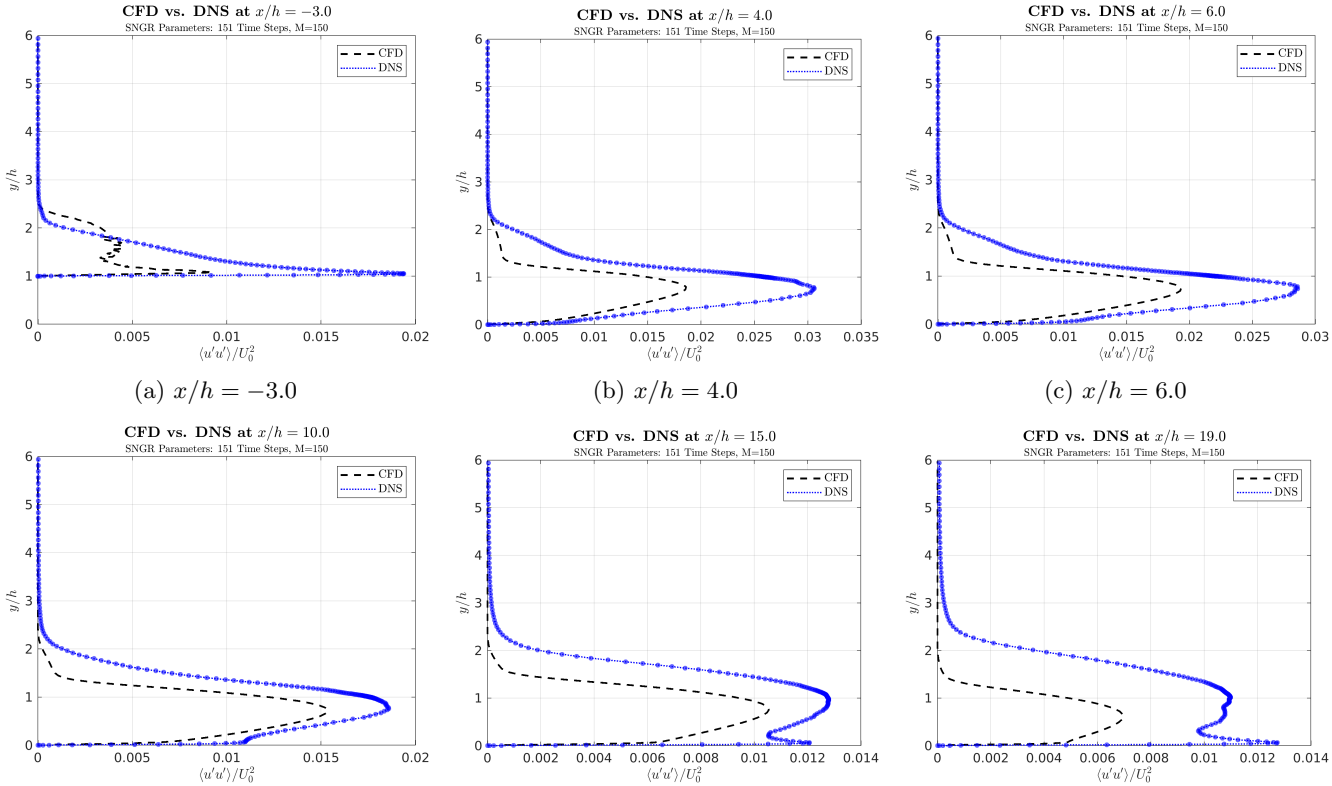

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 10: Streamwise Reynolds stress $\langle u'u' \rangle$ — CFD vs. DNS baseline comparison evaluated at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0), where $h = 0.05$ m.

It can be seen in Figure 10 that the prediction made by StarCCM+ of the Reynolds stresses in the streamwise direction matches quite close the overall profile described by the DNS study. Closest results between both methods are observed at $x/h = 4.0$ and $x/h = 6.0$ - corresponding to Figure 10b and Figure 10c - with a gap between $1 < y/h < 2$, where CFD underpredicts the $u'u'$ component, and a 15% overprediciton of the highest values of $< u'u' >$ according to CFD results. Overall, the location of the maximum stresses match the y/h estimated by DNS and CFD. The mentioned gap reaches its maximum difference right after the step (at $x/h = 4.0$) and decreases downstream. Biggest difference is near the wall, whose region is significantly overpredicted by StarCCM+ from, approximately, $x/h = 10.0$ (getting 100% higher values using CFD than DNS) and increasing this difference as the flow moves towards the outlet, reaching a 200% difference near the outlet, at $x/h = 19.0$.

Following this analysis, the wall-normal component of the Reynolds stresses is exhibited in Figure 11.

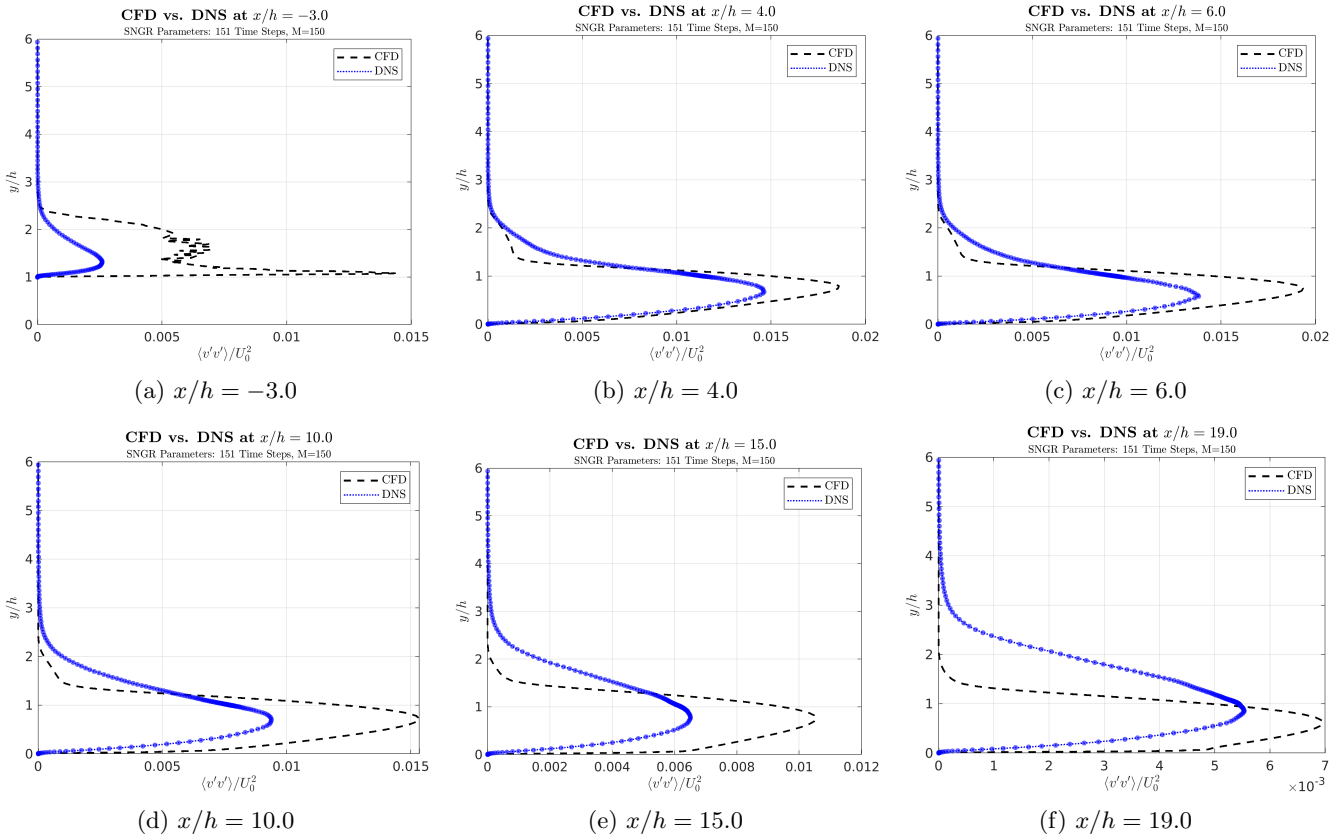

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 11: Wall-normal Reynolds stress $\langle v'v' \rangle$ — CFD vs. DNS baseline comparison evaluated at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0), where $h = 0.05$ m.

It is clearly shown in Figure 11 that the StarCCM+ model considerably overpredicts the maximum values of the wall-normal Reynolds stresses compared to the DNS reference. A key discrepancy is the absence of the sharp stress peak typically observed very close to the wall. Although the vertical axis is plotted in terms of y/h, the corresponding viscous-sublayer region ($y+ < 5$) is very thin and therefore not visually distinguishable in these plots. Instead, the CFD results, which are plotted from the raw cell-centered data imported into MATLAB, show a broad overprediction in the buffer region, starting just upstream of the step at $x/h = -3.0$.

This behavior is a direct consequence of the near-wall mesh resolution. As detailed in the mesh generation section, the first grid cell has a non-dimensional wall distance y^+ ranging from a minimum of 1.2 to a maximum of 8.5. With no grid points located within the viscous sublayer ($y^+ \leq 1$), the RANS simulation is unable to capture the steep velocity gradients at the wall, leading to an incorrect prediction of turbulence production. This results in the model failing to capture the near-wall stress peak while overcompensating further away from the wall, causing the pronounced overshoot seen in the buffer region. While the location of the outer peak is predicted reasonably well up to $x/h = 10.0$ (from where DNS predicts a maximum value at approximately $y/h = 0.8$, whereas CFD estimates it to be at the step's height, at $y/h = 1$), the model's inability to resolve the sublayer remains a primary source of error in the near-wall statistics. It is especially noticeable from $x/h = 10.0$.

The representation of the spanwise Reynolds stresses is observed in Figure 12.

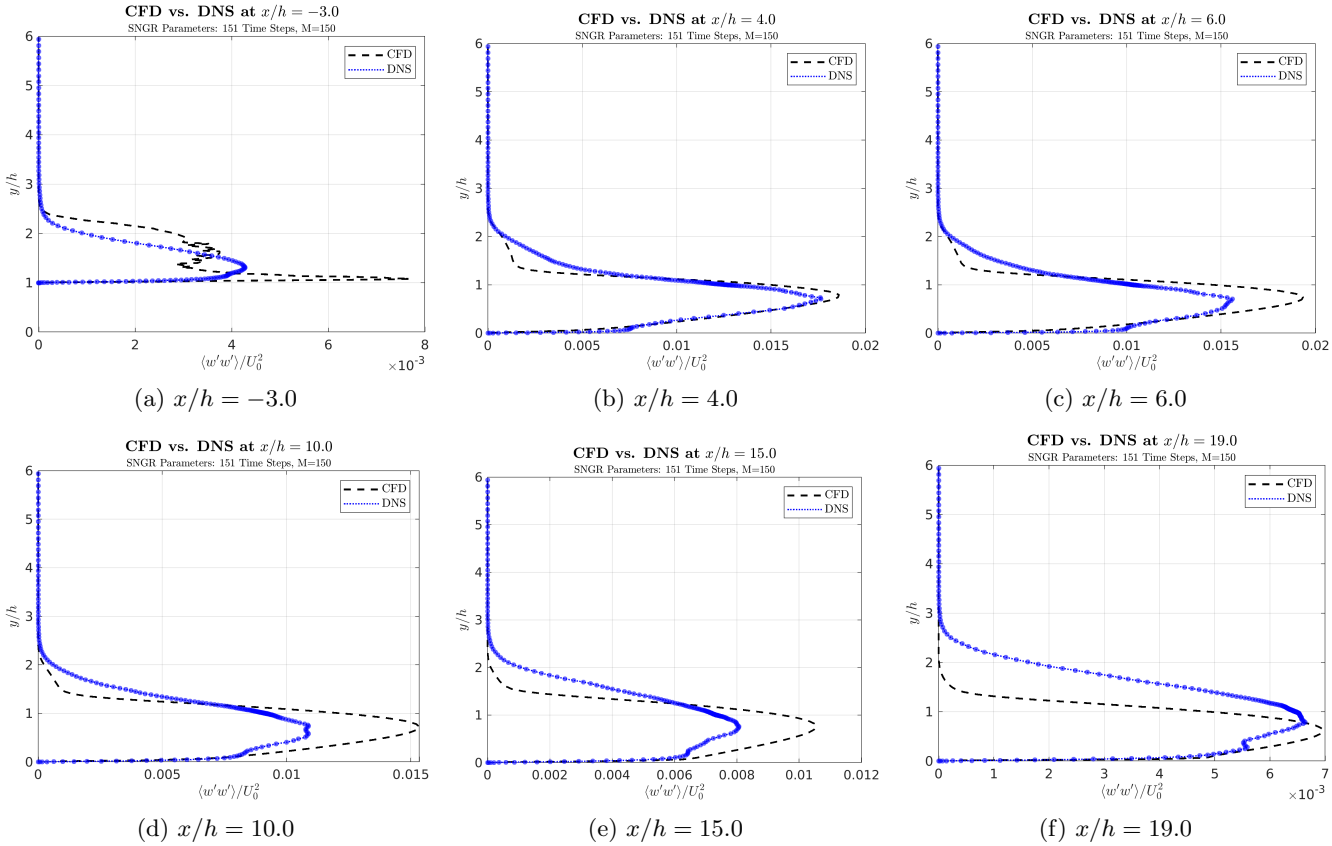

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 12: Spanwise Reynolds stress $\langle w'w' \rangle$ — CFD vs. DNS baseline comparison evaluated at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0), where $h = 0.05$ m.

The current representation of CFD values shows very similar profiles for the spanwise component - represented in Figure 12 - compared to the wall-normal component (previously described in Figure 11: the peak of the magnitude of these two Reynolds stresses components are almost the same at every slice, from upstream the step at $x/h = -3.0$, to near the outlet, at $x/h = 19.0$. On the other hand, DNS also shows profiles that look alike the wall-normal component, with slightly higher values overall, and less smooth profiles in the region below the step's corner, compared to values observed in Figure 11.

With respect to the cross-component $u'v'$, it is observed in Figure 13 that StarCCM+ overpredicts the peak of these stresses at slice $x/h = 4.0$, compared to DNS values. From there, the CFD predicts a lower maximum value than DNS (at least from $x/h = 6.0$), increasing this difference towards the outlet. Additionally, the overprediction made by StarCCM+ of the stresses near the bottom wall - at $y/h = 0$ - is also observed here (from Figure 13d at $x/h = 10.0$ to Figure 13f at $x/h = 19.0$), although with lower significance compared to normal Reynolds stresses components, as observed before from Figure 10 to Figure 12. Besides, the shape described by both methods is in general pretty similar, with the only major difference observed in Figure 13a, with a non-physical representation of the CFD values between $1 < y/h < 2$.

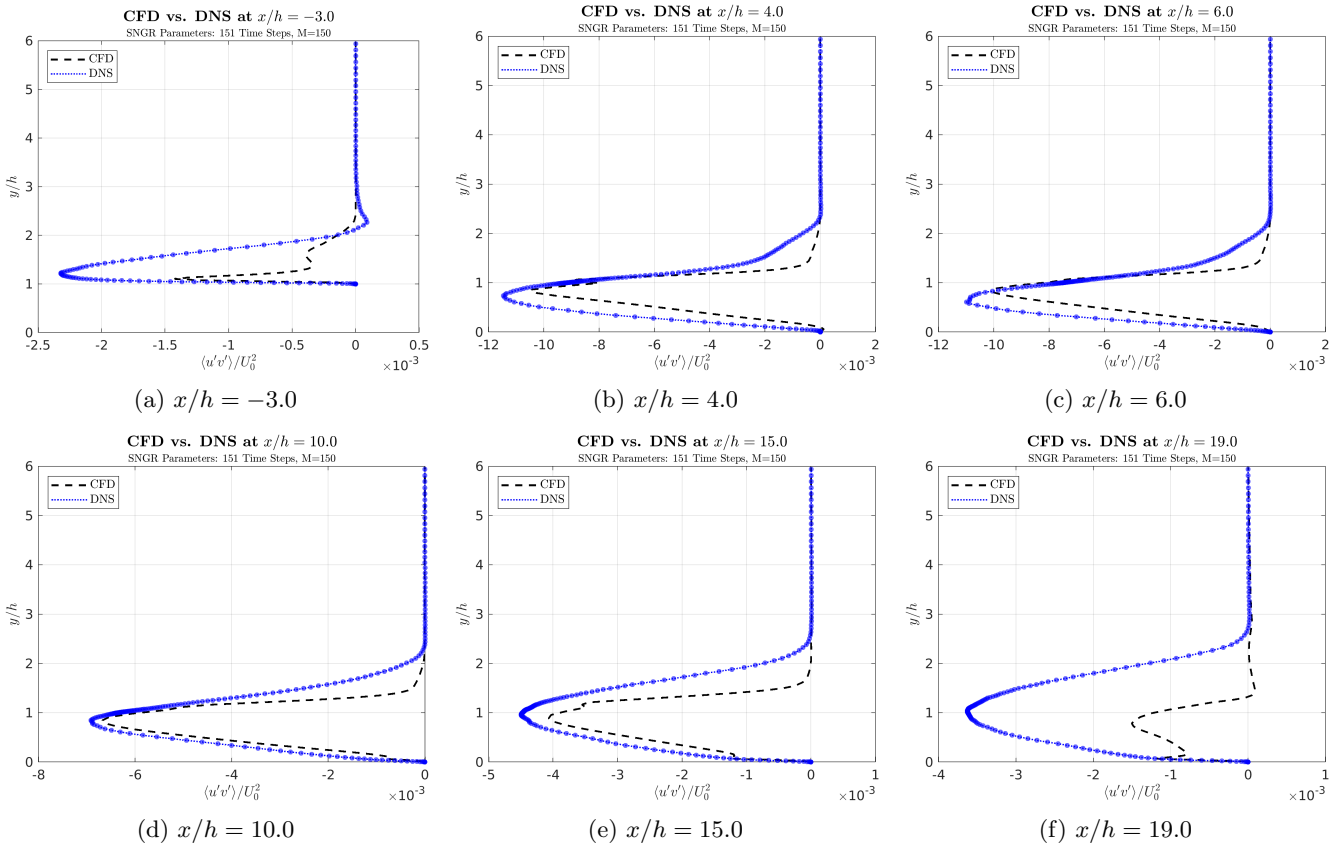

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 13: Reynolds shear stress $\langle u'v' \rangle$ — CFD vs. DNS baseline comparison evaluated at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0), where $h = 0.05$ m.

Results for RANS k- SST

After presenting the results obtained using the RSM Elliptic Blending model as the turbulence closure in $STAR-CCM+$, a similar comparison is now provided for the RANS k-ω SST model. Figure 14 illustrates the streamwise evolution of the Reynolds stress profiles computed with the k-ω SST approach.

The inlet condition predicted by CFD shows the same magnitude as when considering the RSM EB model, being this stage the only one where DNS is estimated to be smaller than the values provided by $Star-CCM+$. Besides, this difference is higher than in Figure 10, and it increases as the fluid flows towards the outlet, observing a major underprediction of the largest Reynolds stresses values along the streamwise direction, as seen at $y/h = 1$ from Figure 14b to Figure 14f.

Furthermore, CFD estimates larger mangitude for the stresses at the bottom wall - at $y/h = 0.0$ - getting up to 250% larger values at $x/h = 19.0$ compared to DNS. However, this difference is less pronounced than what was observed assuming RSM EB as the turbulence model, represented in Figure 10.

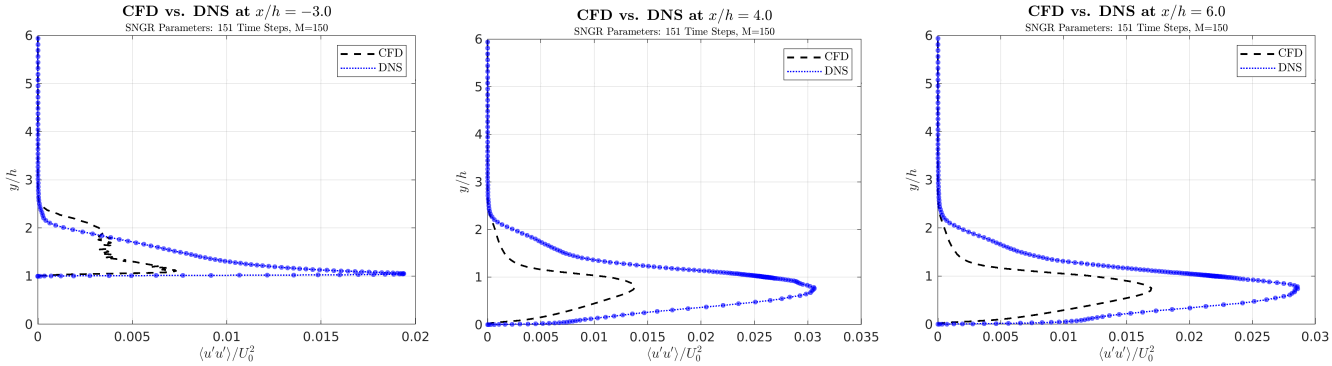

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

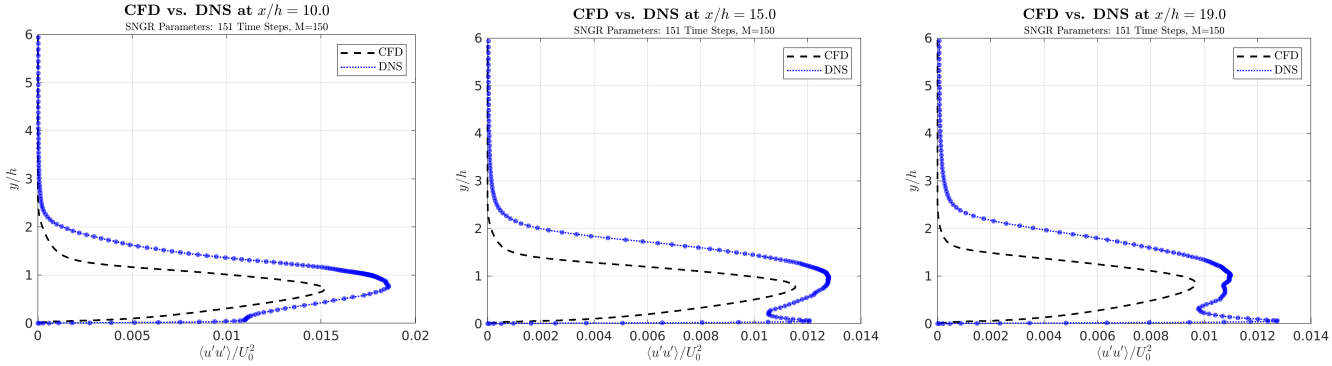

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 14: Streamwise Reynolds stress $\langle u'u' \rangle$ — CFD (k- SST) vs. DNS baseline comparison at six slices.

Following this analysis, the wall-normal component of the Reynolds stresses is exhibited in Figure 15.

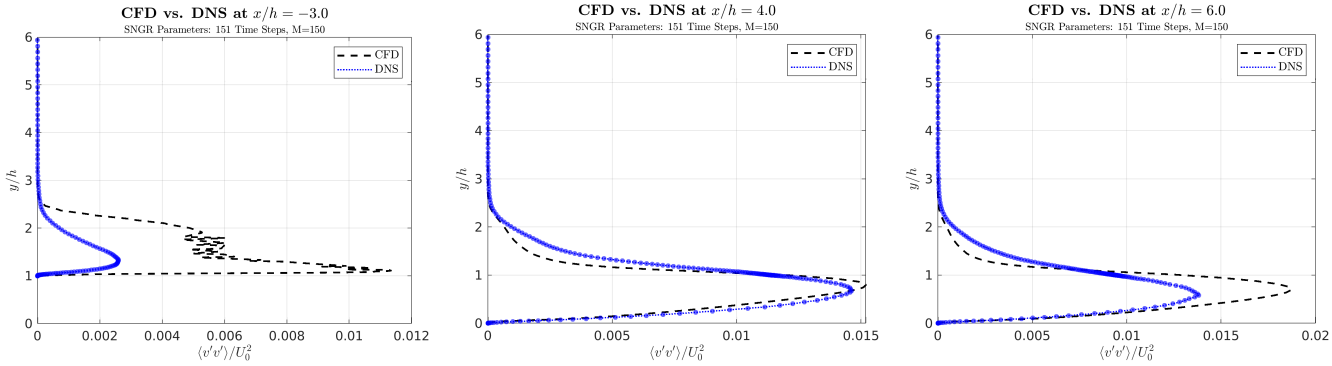

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

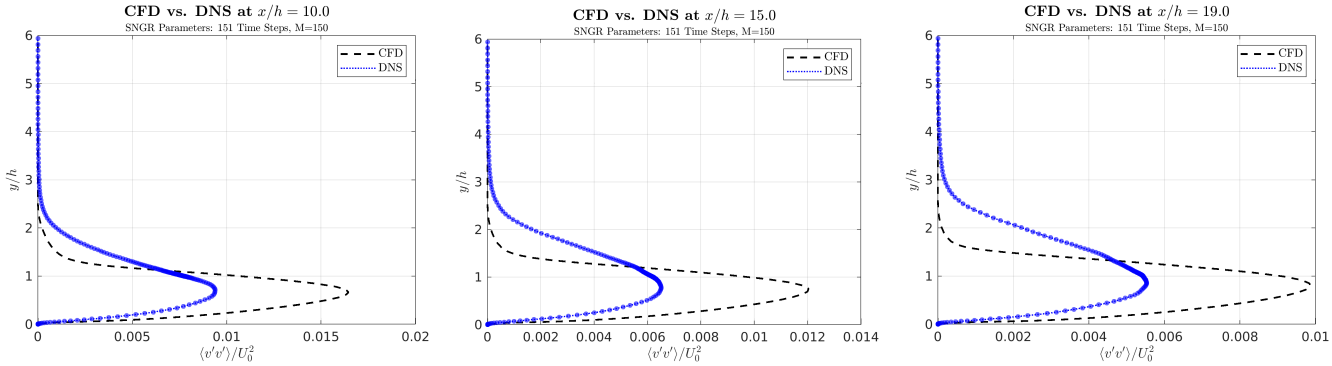

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 15: Wall-normal Reynolds stress $\langle v'v' \rangle$ — CFD (k- SST) vs. DNS baseline comparison at six slices.

A similar behaviour compared to the streamwise component is seen, predicting by assuming the RANS approximation significantly larger magnitudes of the Reynolds stresses compared to DNS. Additionally, these gaps

between both lines in each graph is generally bigger than what is represented assuming the RSM Elliptic Blending model, indicating that this method may provide as accurate values as the RSM EB chosen for this study.

Overall, the shape described by CFD is similar to the one indicated by DNS, obtaining the closest results in first slice evaluated downstream the step, at $x/h = 4.0$. From there, the difference between the largest magnitude predicted by DNS with respect to CFD increases downstream. Furthermore, the region above the corner's step shows a smoother profile considering $Star - CCM+$ than when performing a DNS analysis. This is clearly visible in Figure 15f.

The representation of the spanwise Reynolds stresses is observed in Figure 16. This turbulence model predicts 25% lower maximum Reynolds stresses value at $x/h = 3.0$ compared to RSM EB, as observed in Figure 16a. Near downstream, at $x/h = 4.0$ and $x/h = 6.0$, the SST profile shows a similar shape to DNS, predicting slightly lower values than RSM EB. At $x/h = 10.0$ and at $x/h = 15.0$, the SST estimates about a 5% larger maximum magnitude of the Reynolds stresses in the spanwise direction, increasing the difference with respect to DNS. Near the outlet, at $x/h = 19.0$, the lower part of the profiles overlaps, but the peak is overpredicted by about 35% by using $Star - CCM+$ compared to DNS.

To sum up, the RANS k-ω SST model provides an overall good estimation of the Reynolds stresses in the spanwise direction, close to those values computed by the RSM Elliptic blending model, even though the latter is more accurate than the former.

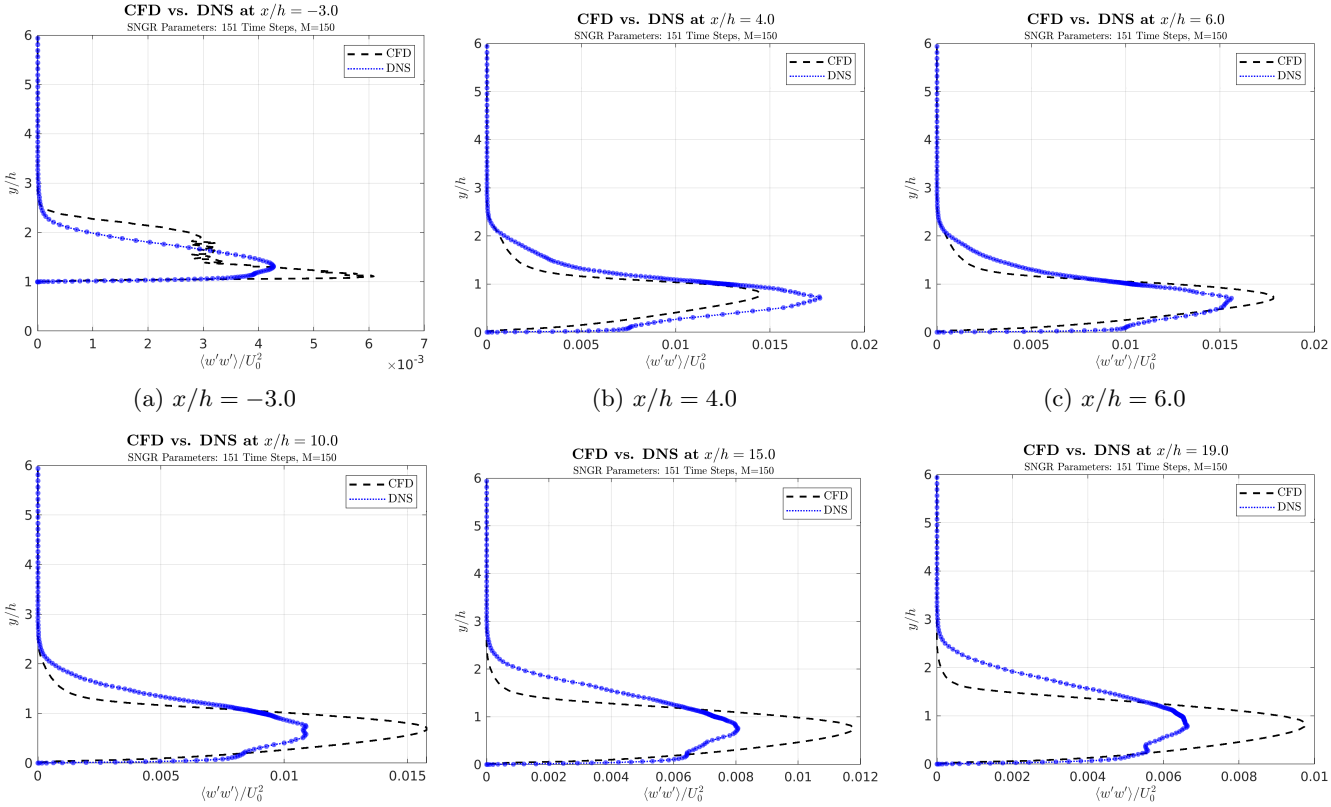

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 16: Spanwise Reynolds stress $\langle w'w' \rangle$ — CFD (k- SST) vs. DNS baseline comparison at six slices.

Having evaluated the three normal components of the Reynolds stresses, the main cross-component uv is represented in Figure 17a. The figures shown in the upper row show that RANS k-ω SST model underpredicts the largest magnitude of these Reynolds stresses up to $x/h = 10.0$, where both profiles - DNS and CFD - describe a very similar shape. From there, the approximation made by the SST turbulence model predicts a smoother transition between the values computed in the region $0 < y/h < 2$ compared to DNS, that shows a much sharper shape. This creates a gap in the region above the step, as observed for the normal Reynolds stresses components.

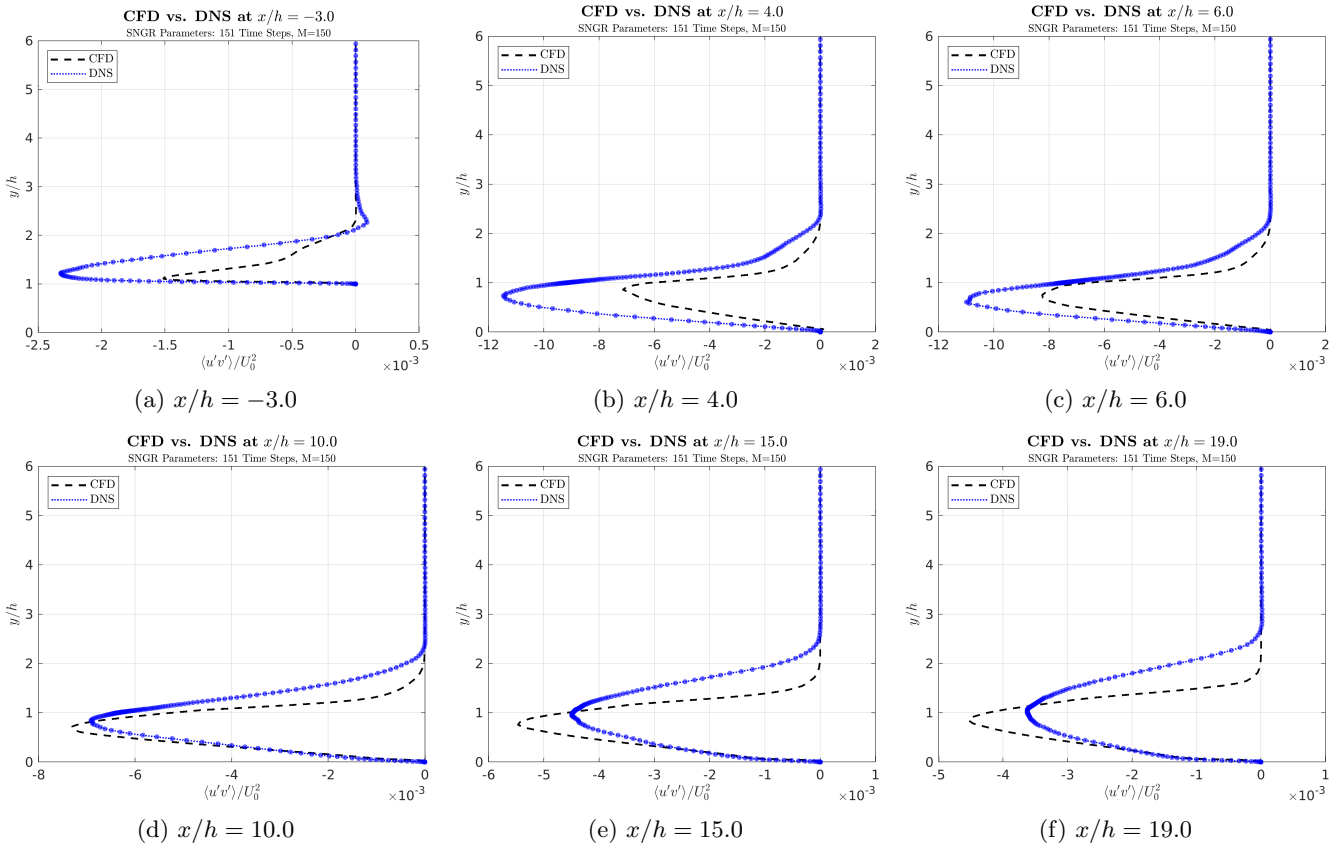

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 17: Reynolds shear stress $\langle u'v' \rangle$ — CFD (k- SST) vs. DNS baseline comparison at six slices.

5.1.2 Reynolds Stresses Comparison: CFD vs. SNGR Methods

Before comparing the SNGR results against the DNS benchmark, it is instructive to validate the SNGR reconstructions against their own input data from the CFD simulation. This section presents a three-way comparison between the RANS solution (CFD), the time-averaged statistics from the single time-loop simulation ($SNGR_{time}$), and the ensemble-averaged statistics ($SNGR_{ens}$). This analysis serves two critical purposes:

- **Input Fidelity Check:** It verifies that both SNGR approaches successfully reproduce the target Reynolds stress tensor provided by the parent RANS simulation.
- **Internal Consistency Check:** It assesses the statistical convergence of the time-marching simulation by comparing its time-averaged results to the statistically independent ensemble average.

These checks directly support the project research questions. Verifying that $SNGR_{time}$ and $SNGR_{ens}$ reproduce the RANS Reynolds stress tensor establishes the input fidelity that is a prerequisite for RQ 2. Quantifying differences between the time-averaged and ensemble-averaged SNGR statistics provides a measure of single-realization statistical uncertainty and therefore informs RQ 1. Finally, any systematic discrepancies between $SNGR_{time}$ and $SNGR_{ens}$ that point to temporal or convective deficiencies identify targets for improvement relevant to RQ 3.

Results for the Elliptic Blending Reynolds Stress Model

The following figures illustrate this comparison for all four Reynolds stress components across the six validation slices, with the RSM Elliptic Blending turbulence model considered.

Figure 18 shows the streamwise Reynolds stresses predicted by both approaches of the SNGR method, as well as estimated by $StarCCM+$.

Overall, both SNGR methods provide pretty similar results to CFD. The spikes observed in all six slices indicate that the simulation is not long enough to capture smooth curves. Besides, the representation shows that the SNGR single time-loop approach - described with the continuous red line - matches a bit better the CFD prediction.

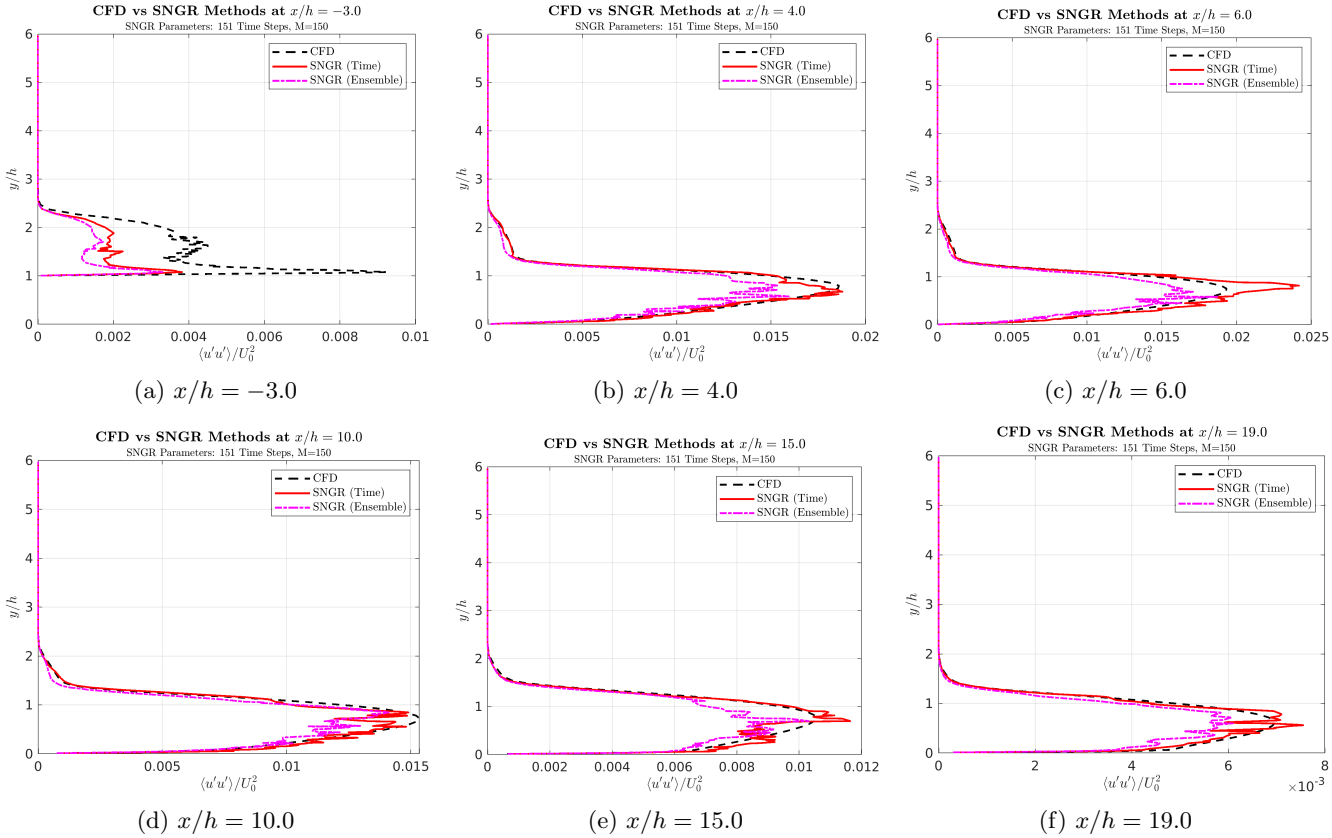

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 18: Streamwise Reynolds stress $\langle u'u' \rangle$ — Comparison of CFD, SNGR time-averaged, and SNGR ensemble-averaged results.

The wall-normal Reynolds stresses are represented in Figure 19. Results show similar behaviour to the previous normal component that has just been evaluated: single time-loop approach matches very closely the CFD values, followed by the ensemble average (with only a significant difference at $x/h = 10.0$). The three methods also share the same shape across the entire domain, and there is only an anomalous horizontal offset observed between the CFD and SNGR profiles at the $x/h = -3.0$ slice, as shown in Figure 19a. This offset was traced to a post-processing inconsistency in the exported CFD profile for that specific slice, rather than to a failure of the SNGR algorithm itself; the SNGR realizations reproduce the original CFD input as intended. All other evaluated slices display close agreement between CFD and SNGR, and the overall assessment of SNGR fidelity is therefore not changed by this single-slice artefact.

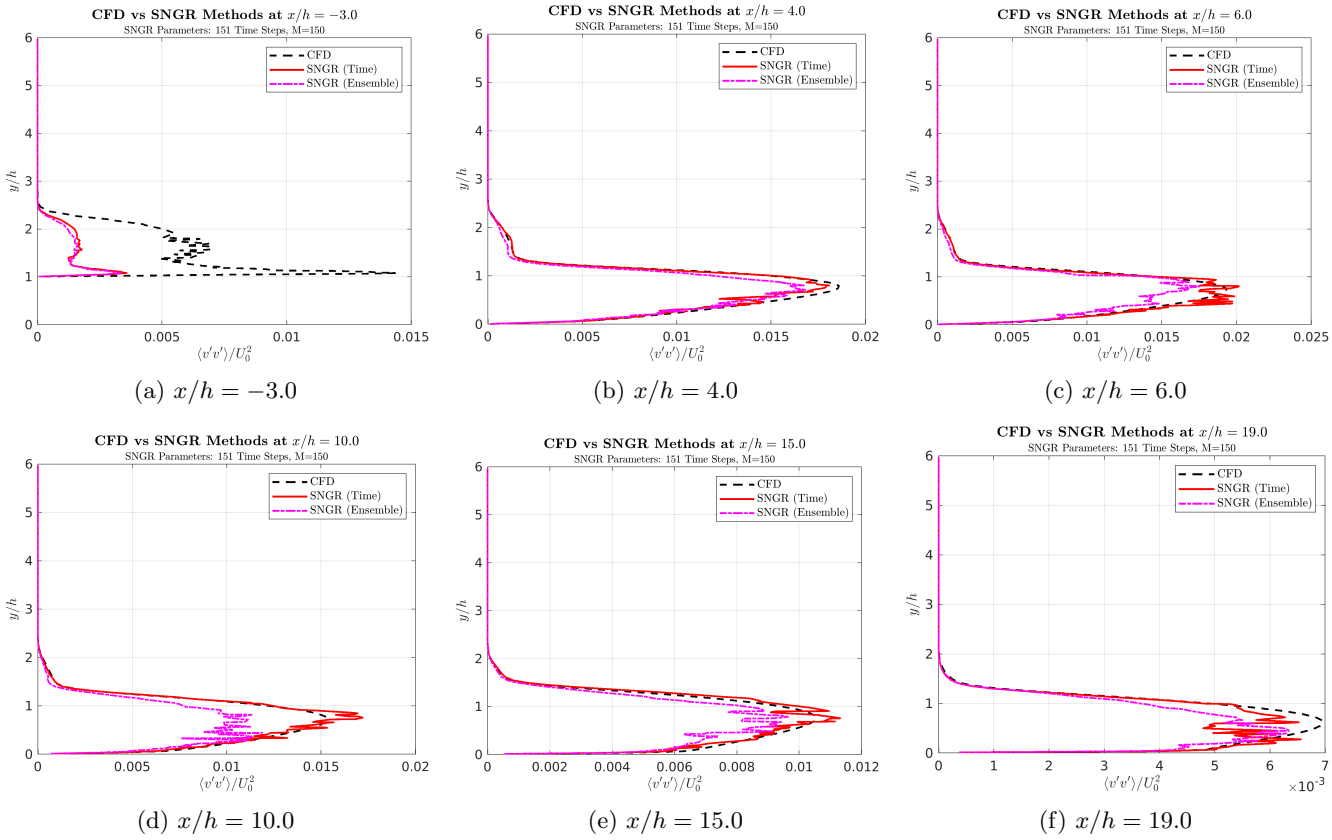

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 19: Wall-normal Reynolds stress $\langle v'v' \rangle$ — Comparison of CFD, SNGR time-averaged, and SNGR ensemble-averaged results.

The third normal component - describing the spanwise Reynolds stresses - is represented in Figure 20. These plots are characterized by a much sharper shape of the values estimated by the SNGR methods, clearly observed at $x/h = 4.0$. Besides, the ensemble average keeps predicting smaller magnitude of the stresses in comparison with the single time-loop approach.

These spikes are consistent with finite-sampling and intermittency effects rather than a fundamental algorithmic failure: turbulent fluctuations are temporally correlated and intermittent events produce large instantaneous contributions to second moments that survive short averages. Empirically we observe two complementary convergence behaviours from extended tests: (i) increasing the number of statistically independent realizations reduces sampling noise and yields smoother *ensemble* profiles, and (ii) increasing the simulated time T_{sim} for a single realization reduces temporal sampling noise and yields a smoother *time-averaged* profile.

For reference, the standard diagnostics that quantify these effects are:

$$T_{\text{int}} = \int_0^\infty \rho_{u'}(\tau)\, \mathrm{d}\tau \tag{60}$$

(the integral time scale computed from the autocorrelation $\rho_{u'}(\tau)$), and an estimate of the effective number of independent samples

$$N_{\text{eff}} \approx \frac{T_{\text{sim}}}{2\, T_{\text{int}}}. \tag{61}$$

The sampling error (standard error) of a Reynolds-stress estimate $R = \overline{u'v'}$ then scales approximately as

$$\mathrm{SE}(R) \approx \frac{\sigma_R}{\sqrt{N_{\text{eff}}}}, \tag{62}$$

where σ_R^2 is the variance of the instantaneous $u'v'$ samples. Equivalently, the variance of the ensemble mean decreases roughly as

$$\mathrm{Var}\big(\overline{R}\big) \sim \frac{\mathrm{Var}(R)}{N_{\text{eff}}}. \tag{63}$$

In the present manuscript no formal integral-time or uncertainty metrics are reported for the displayed profiles; the smoothing statements above are empirical observations from extended development runs. If desired, the diagnostics in equations (60)–(63) can be computed from the stored time series to produce quantitative confidence

intervals for the SNGR time-averages and to determine the T_{sim} (or ensemble size) required to achieve a target uncertainty.

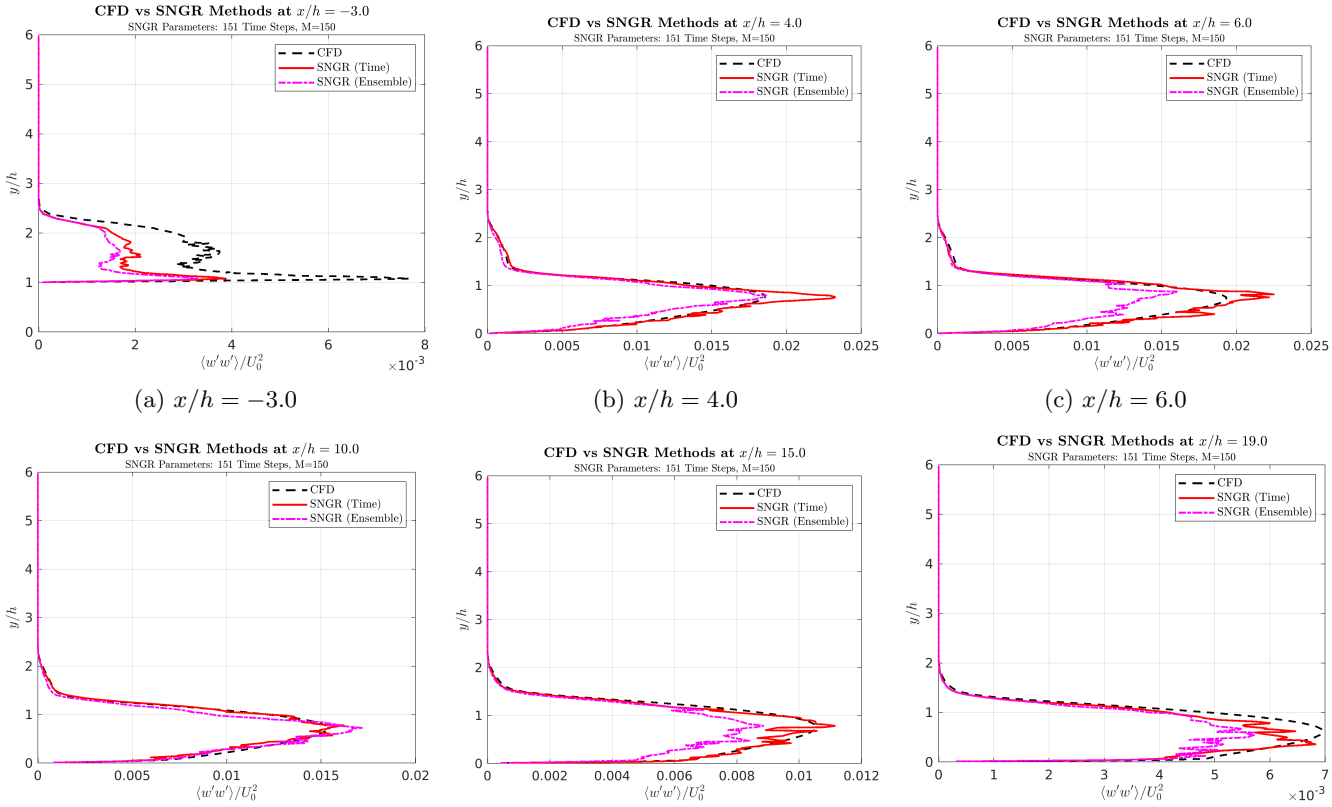

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 20: Spanwise Reynolds stress $\langle w'w' \rangle$ — Comparison of CFD, SNGR time-averaged, and SNGR ensemble-averaged results.

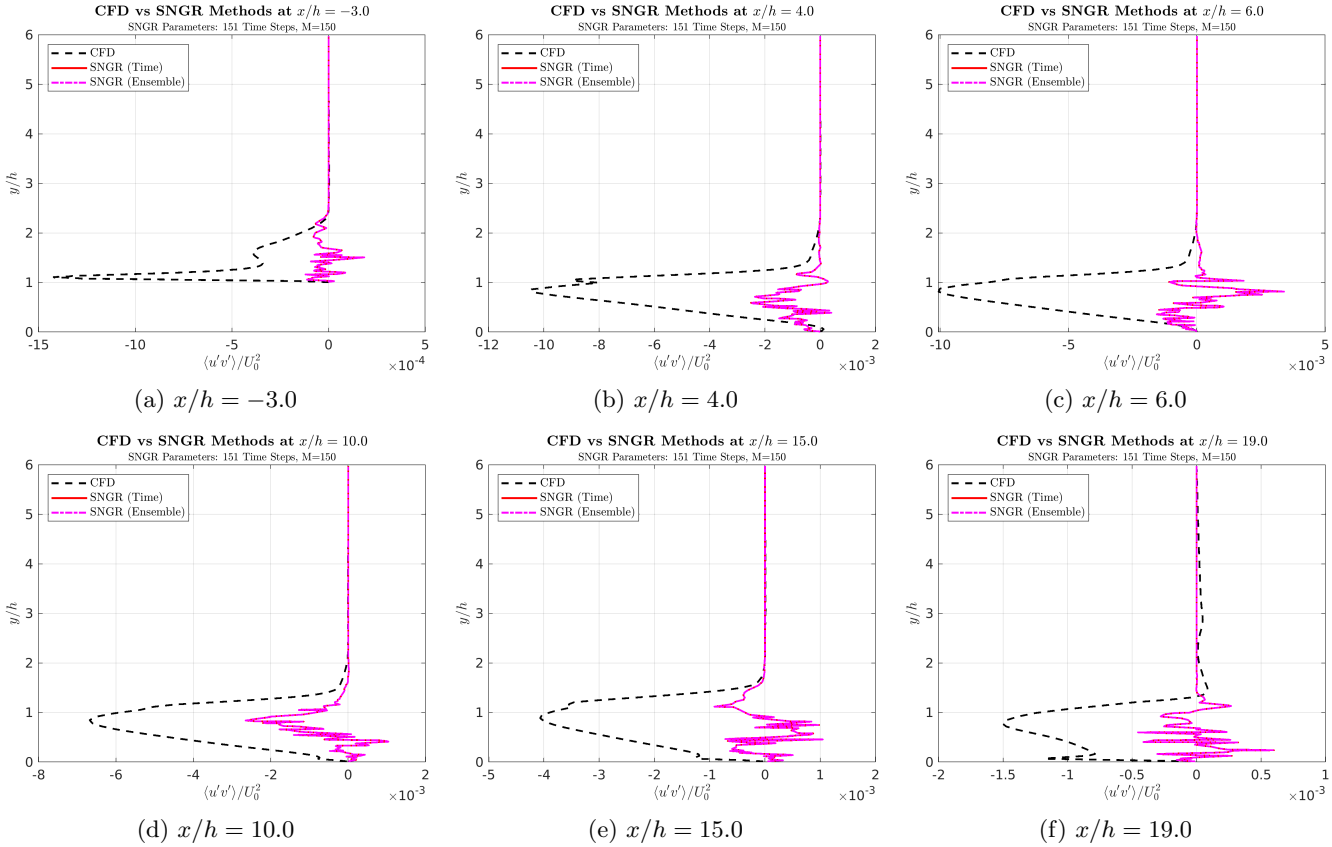

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 21: Reynolds shear stress $\langle u'v' \rangle$ — Comparison of CFD, SNGR time-averaged, and SNGR ensemble-averaged results.

As shown in Figures 18 to 20, both the SNGR time-averaged and ensemble-averaged methods demonstrate strong fidelity to the input CFD data for the normal stress components. The profiles for $SNGR_{time}$ and $SNGR_{ens}$ closely follow the shape and magnitude of the CFD reference profiles, confirming that the core reconstruction is working as intended for these variance-based quantities. Minor deviations between the two SNGR methods can be attributed to the finite duration of the time-marching simulation, which may not have fully achieved statistical convergence to the true mean represented by the large-sample ensemble average.

The perfect overlap observed between the single time-loop and ensemble-averaged results for the Reynolds shear stress component, $\langle u'v' \rangle$, stems from a common flaw in the underlying fluctuation generation method used by both approaches. As detailed in Appendix A, the initial isotropic velocity field generated by the function used to compute the turbulent velocity field contained a non-physical spatial correlation that artificially decorrelated the u' and v' components. Consequently, any snapshot or time-step generated using this function inherently produces a near-zero shear stress.

Since both the ensemble method (averaging many flawed snapshots) and the time-marching method (integrating a process continually fed by flawed fluctuations) rely on this same incorrect source, they both consistently and independently converge to the same erroneous result of $\langle u'v' \rangle \approx 0$. This is why their profiles match perfectly.

In contrast, the normal stress components ($\langle u'u' \rangle$, $\langle v'v' \rangle$, etc.) are variances and are primarily sensitive to the energy content of the fluctuations, which the SNGR method still produces. The minor differences seen in their plots are the expected statistical variations between a finite-time average with temporal correlation and a statistically independent ensemble average. This fundamental issue with the shear stress has since been identified, and a corrected implementation for generating a truly isotropic field is presented in Appendix A for future work.

Results for RANS k-ω SST

The following figures illustrate this comparison for all four Reynolds stress components across the six validation slices, with the RANS k- SST turbulence model considered. The length of the simulation is the same as that considered for the RSM Elliptic Blending model, as seen in the title of each individual figure.

The first normal component that is represented corresponds to the streamwise Reynolds stresses, shown in Figure 22.

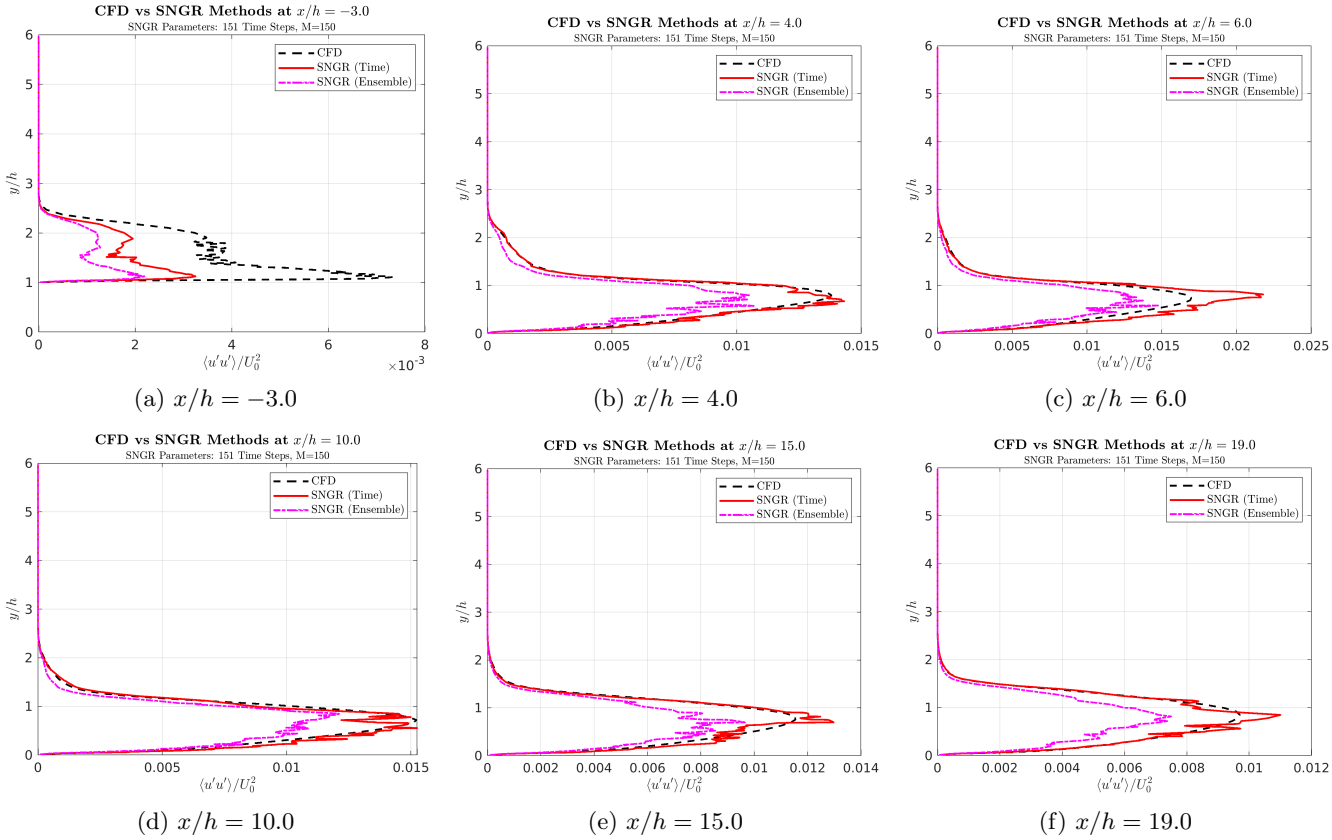

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 22: Streamwise Reynolds stress $\langle u'u' \rangle$ — Comparison of CFD (k- SST), SNGR time-averaged, and SNGR ensemble-averaged results.

The results for the wall-normal component are represented in Figure 23. The difference observed between both SNGR approaches is larger than for the RSM turbulence model, underpredicting the ensemble average method the magnitude of the streamwise Reynolds stresses for all six slices. This difference is significant in the entire downstream region of the step, from $x/h = 4.0$ to $x/h = 19.0$. Considering this large impact, it would be highly recommended to use the RSM EB model when evaluating the SNGR ensemble method, as the results are quite closer to CFD than what is seen in this specific case.

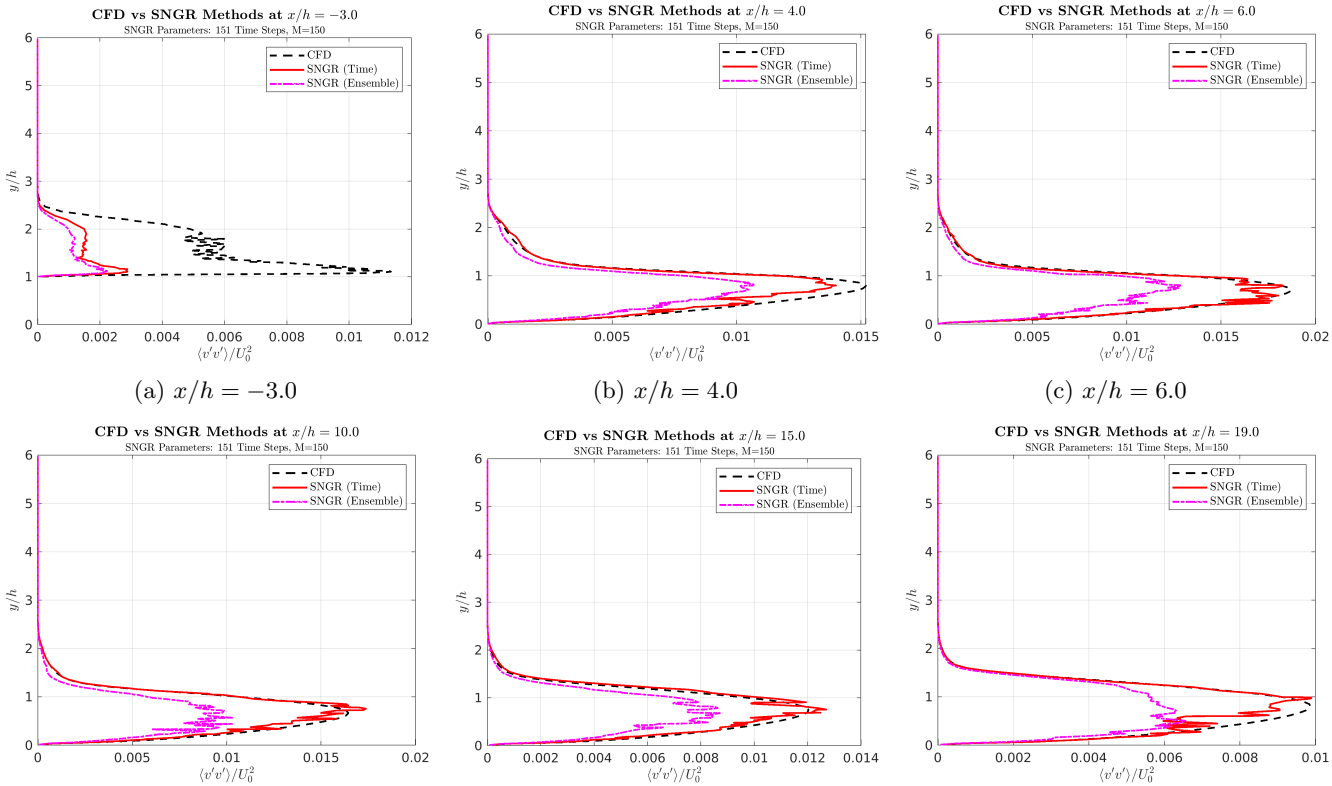

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 23: Wall-normal Reynolds stress $\langle v'v' \rangle$ — Comparison of CFD (k- SST), SNGR time-averaged, and SNGR ensemble-averaged results.

With respect to the Reynolds stresses computed in the spanwise direction - represented in Figure 24 - the slice with the most similar values between all three methods is at $x/h = 10.0$, the same behaviour as the one observed for the RSM Elliptic Blending model. Furthermore, the spikes and overall shape of the profiles are quite close to what was seen in Figure 20, even though the difference between the single time-loop approach and the ensemble average is usually a bit larger than for the RSM EB case.

The cross component uv is plotted in Figure 25 and, as it was observed for the other turbulence model, the values predicted by the SNGR methods are unphysical. Besides, values from both SNGR approaches are perfectly matching, which is a representation error. As explained in Appendix B, the methodology required to compute and represent these values is very different compared to the normal components. This explains why the streamwise and wall-normal components comparisons show profiles that follow the same shape, whereas their cross-component does not properly represent the SNGR values.

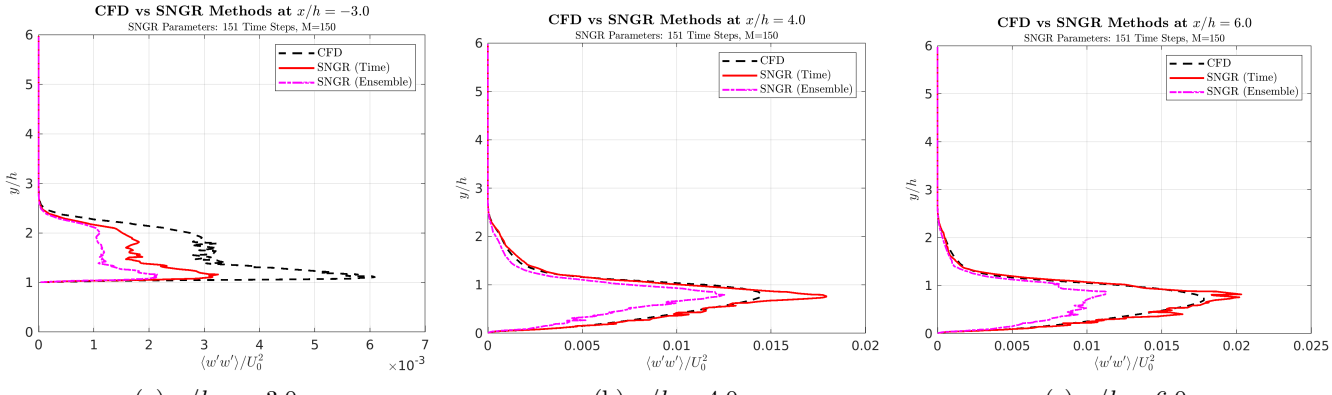

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

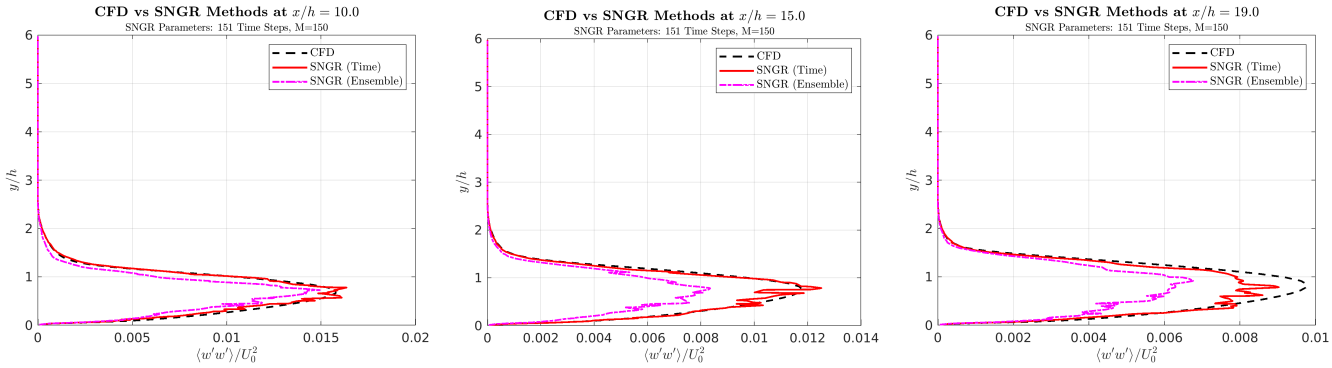

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 24: Spanwise Reynolds stress $\langle w'w' \rangle$ — Comparison of CFD (k- SST), SNGR time-averaged, and SNGR ensemble-averaged results.

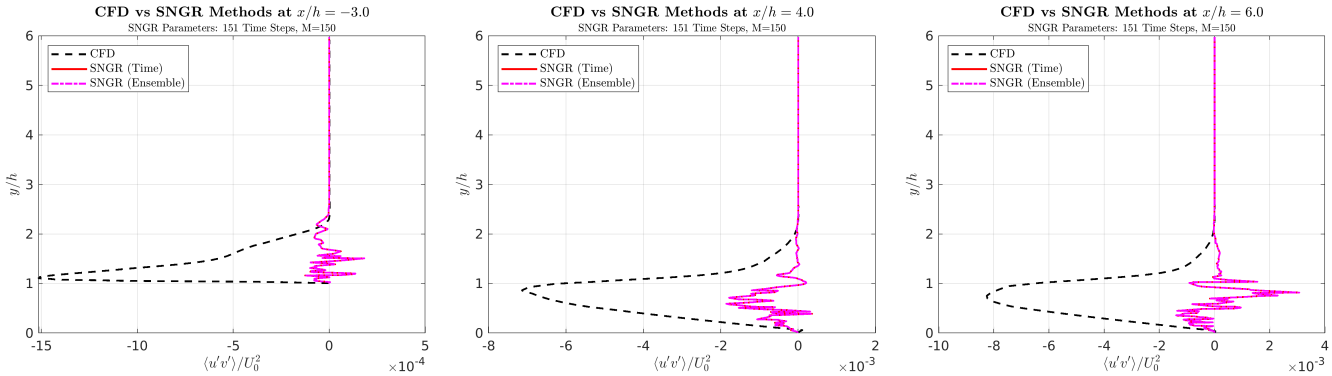

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

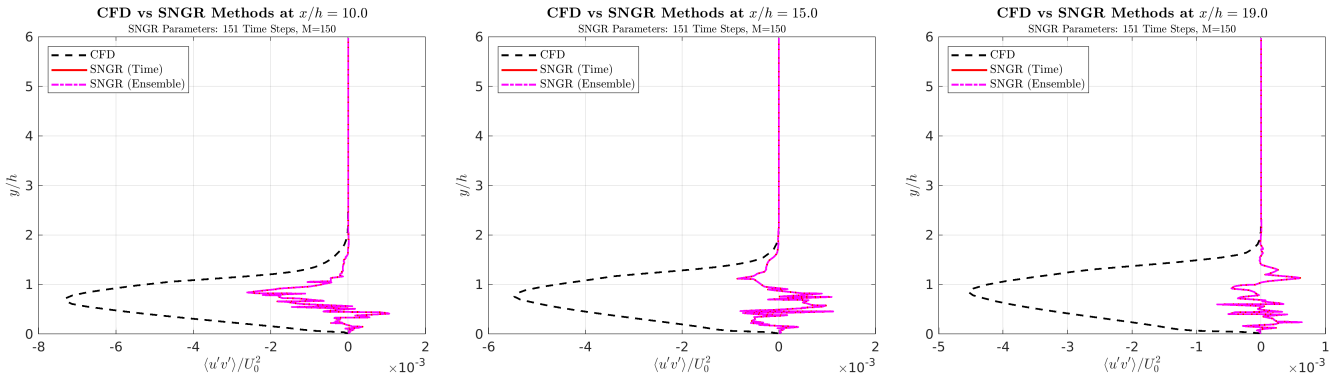

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 25: Reynolds shear stress $\langle u'v' \rangle$ — Comparison of CFD (k- SST), SNGR time-averaged, and SNGR ensemble-averaged results.

5.1.3 Reynolds Stresses Comparison: SNGR single time-loop vs. SNGR ensemble average vs. CFD vs. DNS

As previously explained in Section 1.4, the SNGR approach uses two distinct methods whose results vary as one of them is using uncorrelated data - the ensemble average, $SNGR_{ens}$ in the plots - whereas the single time-loop approach $SNGR_{time}$ computes data based on the previous iteration. Values obtained using these two techniques are plotted against CFD and DNS data, in 6-slices as shown in other Reynolds stresses , for every evaluated Reynolds stress component. As before, the first component to be studied is the streamwise direction $< u'u' >$, getting an overall lower prediction of Reynolds stresses considering the ensemble average instead of the single time-loop approach, as observed in Figure 26.

Results for the Elliptic Blending Reynolds Stress Model

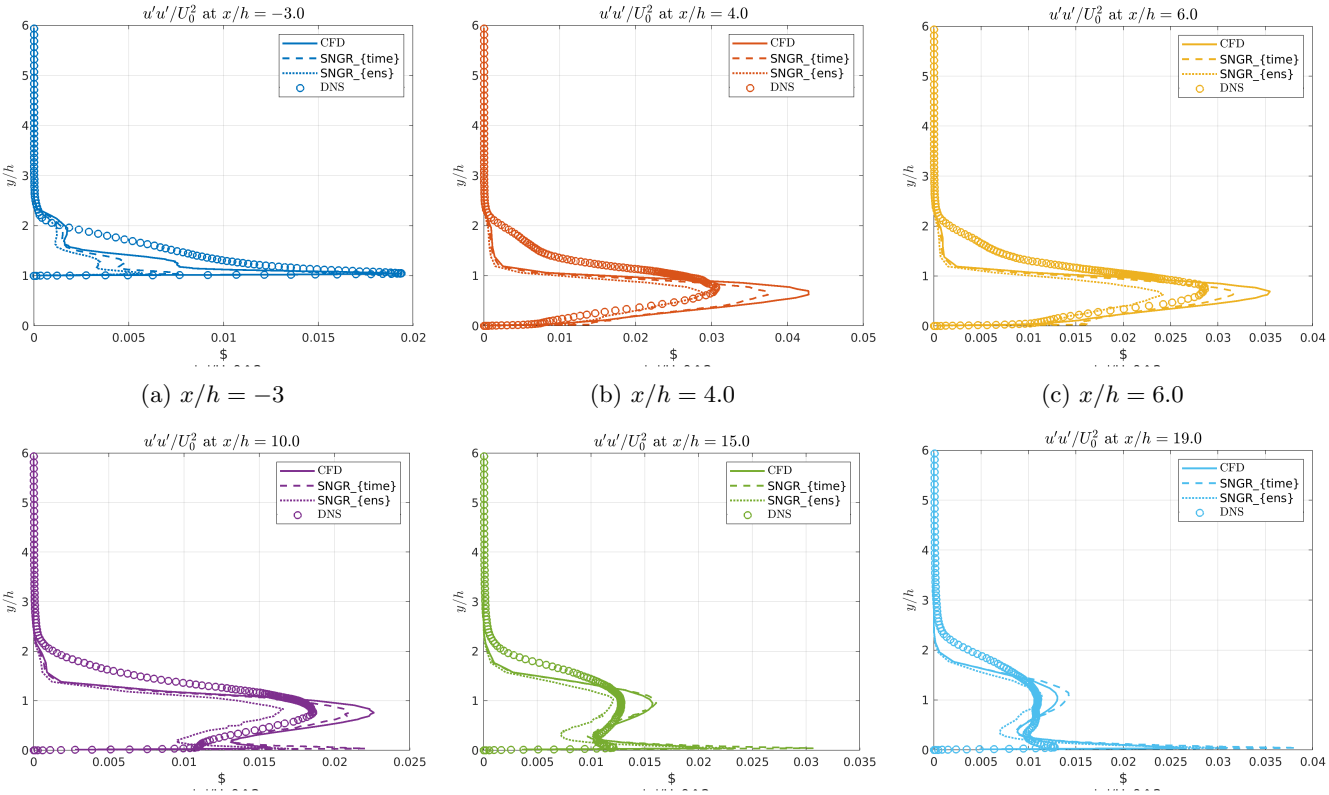

(a) $x/h = -3$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10$ (e) $x/h = 15$ (f) $x/h = 19$

Figure 26: Reynolds stress $\langle u'u' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices. Top row: $x/h = -3$, 4.0, 6.0. Bottom row: $x/h = 10$, 15, 19.

Upstream the step, it can be seen how the ensemble method provides a consistently lower estimation of the streamwise Reynolds stress compared to the single time-loop method, and both dotted lines have a significant lower magnitude as opposed to DNS and CFD values (whose maximum magnitude is matched at the bottom wall), represented using a string of circles and a thick line, respectively, in Figure 26a.

The first evaluation of the flow downstream the step shows that both SNGR methods follow a very similar shape as CFD, with a visible gap in the the region just above the step's height compared to DNS values, as exhibited in Figure 26b. As seen upstream the step, CFD overpredicts the streamwise Reynolds stresses compared to SNGR approaches, followed in magnitude by the time-loop estimation, and getting a good match between DNS and $SNGR_{ens}$ values.

A bit further downstream, a similar analysis can be extracted. SNGR matches the shape described by CFD, with a very similar distribution of maximum values compared to the previous slice, and with slightly smaller gap between DNS values and the other numerical methods, as reported in Figure 26c. In this case, DNS values have got just slightly lower (maximum magnitude of $0.0295m^2/s^2$ at $x/h = 6.0$ compared to $0.0305m^2/s^2$ at $x/h = 4.0$.

Following the flow further downstream, at $x/h = 10.0$, both CFD and SNGR approaches estimate values considerably larger than those provided by DNS at the bottom wall, as observed in Figure 26d. CFD, $SNGR_{time}$, and $SNGR_{ens}$ achieve a better representation of the outer part of the separated shear layer, close to the free-stream region. Moreover, largest difference between CFD and $SNGR_{time}$ keeps decreasing as the fluid flows, reaching an almost perfect match at $x/h = 15.0$, shown in Figure 26e. At this stage, DNS predicts a very similar value at the

bottom wall compared to its previous and following slice, represented in Figure 26d and Figure 26f, respectively. Nonetheless, CFD seriously increases the value at the bottom wall from the step to the outlet, reaching exceedingly large values that cause SNGR methods to overshoot the estimation of streamwise Reynolds stresses in this section. However, the outer separated shear layer computed by every numerical method seems to provide a closer look between them as the flow goes downstream.

The second normal cross-component evaluated corresponds to the wall-normal direction $< v'v' >$, representing all its slices combined in a single plot in Figure 29.

The plotted quantity in every panel is $\langle v'v' \rangle / U_0^2$, where v' is the wall-normal velocity fluctuation, and U_0 is the reference free-stream speed. If U_0 carries SI units ($\mathrm{m\,s^{-1}}$) then $\langle v'v' \rangle$ has units $\mathrm{m^2\,s^{-2}}$; dividing by U_0^2 produces a dimensionless profile. In this dataset $U_0 = 1$, so the plotted numbers are numerically equal to the dimensional variance in $\mathrm{m^2\,s^{-2}}$. Overall, SNGR single time-loop approach - represented using a long dotted line - computes the largest Reynolds stresses values for this wall-normal component, followed by a smaller overestimation made by the ensemble average method. Both SNGR methods describe a similar shape as the one provided by means of CFD data.

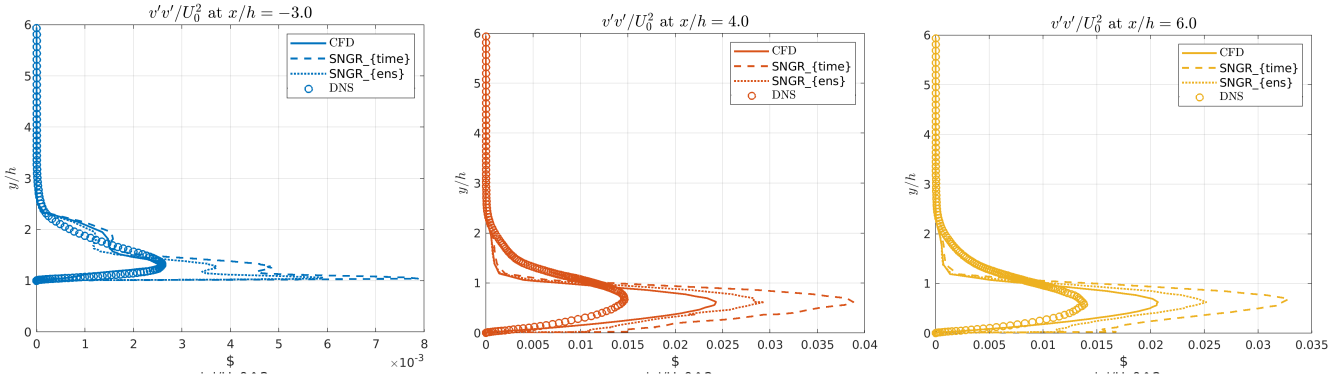

(a) $x/h = -3$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10$ (e) $x/h = 15$ (f) $x/h = 19$

Figure 27: Reynolds stress $\langle v'v' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

Upstream of the step (at $x/h = -3.0$), all methods predict very low Reynolds stresses in the wall-normal direction, as expected in the fully developed inlet region. The horizontal axis spans values of order 10^{-3}, indicating an attached, low-mixing flow. All data sources (CFD, a single SNGR time realization, the SNGR ensemble and the DNS markers) cluster close to the wall and show a monotonic decay with increasing y/h, reflecting the absence of a separated shear layer here. The profiles overlap closely in the outer shear layer, confirming mostly consistent initial conditions, even though at the bottom wall there is a remarkable difference between DNS and the other three methods, as seen in Figure 26a.

Shortly downstream of the step the magnitude of $\langle v'v' \rangle / U_0^2$ rises by more than an order of magnitude and a pronounced outer maximum appears (around $y/h \approx 0.7$), as observed in Figure 27b and Figure 27c. This outer hump is the signature of the separated shear layer and strong wall-normal mixing; the near-wall peak coexists with the shear-layer peak. DNS prescribes the reference peak location and amplitude as reaching a maximum value of $0.015 m^2/s^2$ at $y/h = 0.75$; the CFD field over predicts the outer peak amplitude by more than 150%, while the SNGR tends to produce a thinner distribution, capturing the mean shape. Discrepancies at this slice therefore reflect how well each method represents shear-layer production and component redistribution.

At $x/h = 10$, the outer hump is weaker and profiles trend toward a more relaxed state: wall-normal fluctuations decay as mixing and dissipation act on the separated shear layer, as it is shown in Figure 27d. $SNGR_{ens}$ produces

a somewhat fluctuating maximum peak, whereas the single time-loop approach produces a smoother shape. Any systematic over prediction in this region signals issues with global energy scaling or with the pressure–strain (redistribution) modelling used to convert scalar energy into component wise stresses.

Continuing downstream the flow approaches a reattached or re-equilibrated state: the near-wall peak becomes relatively more pronounced as the outer peak decrease its magnitude about 40% and the wall-normal fluctuation intensity reduces, as represented in Figure 27d

Far downstream the profiles are smaller in magnitude and smoother, with the outer hump largely gone and near-wall turbulence dominating the $v'v'$ budget according to the SNGR prediction made by both approaches, as observed in Figure 27f. Good agreement at this slice demonstrates that both SNGR (after any applied scaling) and the RANS baseline capture the downstream relaxation of wall-normal fluctuations; conversely, systematic biases here would suggest a global energy mismatch introduced earlier (for example in the SNGR input R_{ij} or in a scaling factor) rather than a localized modelling failure.

The third and final normal Reynolds stress component being analyzed is $\langle w'w' \rangle / U_0^2$, where w' is the spanwise velocity fluctuation and U_0 is the reference free-stream speed.

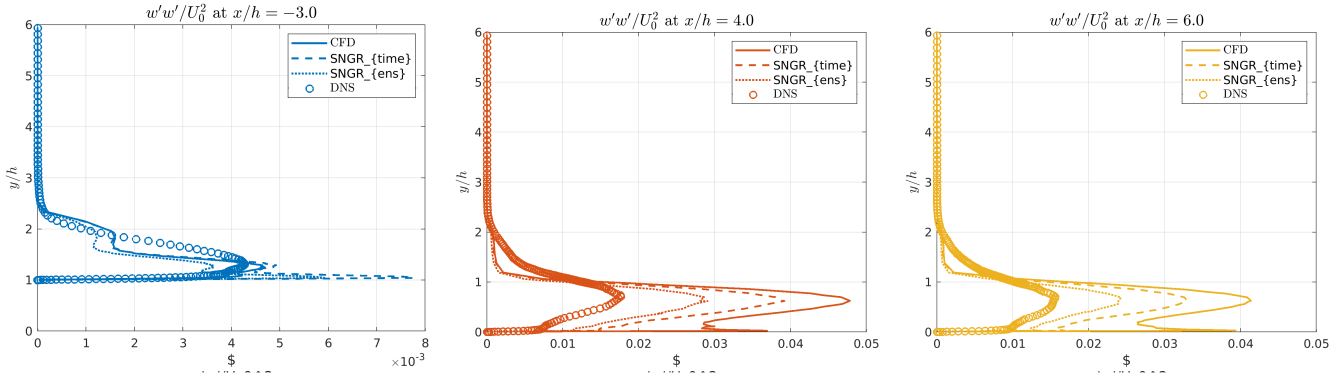

(a) $x/h = -3$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10$ (e) $x/h = 15$ (f) $x/h = 19$

Figure 28: Reynolds stress $\langle w'w' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

Upstream of the step the spanwise fluctuations are very small and tightly confined to the near-wall region: the profile is strongly peaked at low y/h and decays rapidly with height, as exhibited in Figure 28a. CFD, the SNGR single time-loop, the SNGR ensemble and the DNS markers are all clustered near zero across most of the channel, indicating that spanwise motions are weak in the attached upstream flow. CFD provides a quite close shape compared to DNS, and both SNGR methods have similar shapes as those two other approaches everywhere but at the bottom wall.

Immediately downstream of the step, at $x/h = 4.0$, the spanwise variance rises markedly and one observes an outer shoulder/peak associated with the separated shear layer at about $y/h = 1.0$. DNS prescribes the reference amplitude and shape of this shoulder, while the CFD solution underpredicts the outer structure due to turbulence-closure assumptions (RSM–EB redistributes stresses toward equilibrium anisotropy, and k–ω SST relies on the Boussinesq hypothesis, both of which damp the spanwise fluctuations in the shear layer). At the same time, CFD severely overpredicts the near-wall peak compared to DNS. Consequently, both SNGR approaches show an overshoot in this lower region, although the overall shape remains quite similar to DNS—particularly notable given the large near-wall overshoot present in the CFD input field.

Further downstream at $x/h = 6.0$ the outer shoulder persists but begins to broaden and shift as the shear layer spreads. The relative amplitude of the outer feature compared with the near-wall peak can change between methods: if SNGR_ens and DNS align in peak position and height that indicates the stochastic reconstruction plus anisotropy

correction captured the spanwise production correctly; if RANS remains smoother or systematically lower, that points to limitations of eddy-viscosity redistribution for this component.

Downstream of the reattachment zone, at $x/h = 10.0$, the outer hump weakens and the profiles begin to approach lower overall intensity. The $SNGR_{time}$ result exhibits larger scatter due to instantaneous variability, whereas $SNGR_{ens}$ provides a smoother mean profile, with CFD generally lying in between. The persistent amplitude offsets observed at this slice appear to be global rather than local, suggesting that they arise from the broader energy distribution inherited from the RANS input rather than from a local modelling deficiency at this specific streamwise position.

At $x/h = 15.0$ the spanwise fluctuations are diminished relative to the peak immediately downstream of the step; the profile tends to be dominated by the near-wall contribution as the separated shear layer dissipates and mixes into the bulk. Good agreement between SNGR_ens and DNS here indicates successful recovery of the redistributed component energy; systematic differences in shape or integrated area indicate remaining issues in the anisotropy or scaling corrections applied to the SNGR snapshots.

Finally, far downstream the profile is low in magnitude and smooth with the outer shoulder largely absent, as represented in Figure 28f; the near-wall region contains most of the residual spanwise variance. Agreement at this slice demonstrates that methods capture the downstream relaxation of spanwise turbulence.

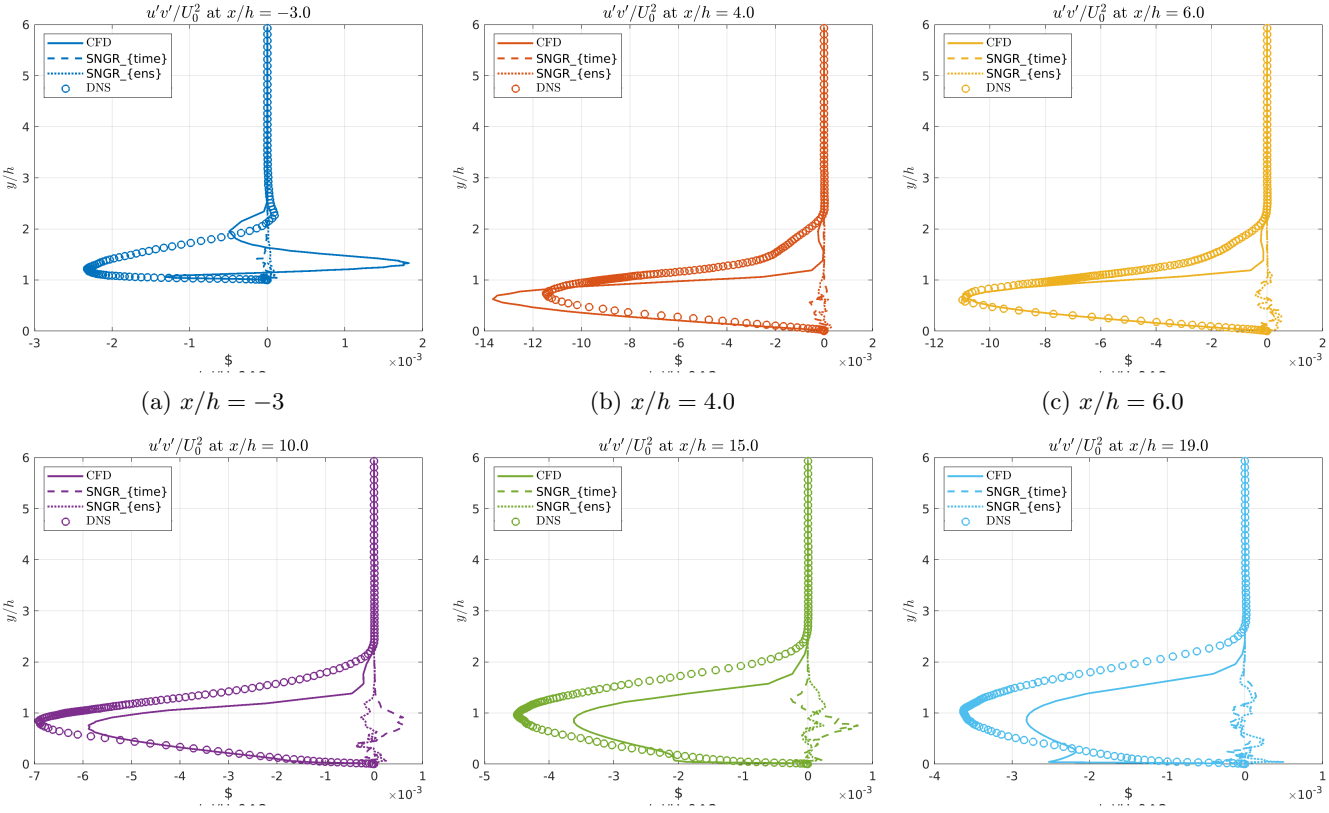

(a) $x/h = -3$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10$ (e) $x/h = 15$ (f) $x/h = 19$

Figure 29: Reynolds stress $\langle u'v' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

The cross-component $\langle u'v' \rangle$ is shown in Figure 29. As expected, CFD and DNS display consistent behaviour downstream of the step and opposite signs upstream (see Figure 29a), reflecting the change in shear direction across the separation region. In contrast, both SNGR approaches produce values that fluctuate around zero for all six slices, indicating that the synthesized fluctuations do not generate a physically meaningful correlation between u' and v'. This behaviour is especially notable given the otherwise good agreement obtained for the normal-stress components; despite accurately matching the streamwise and wall-normal variances, the SNGR reconstruction fails to reproduce the cross-component due to the lack of a sustained coupling between the two fluctuating velocity components.

Results for RANS k- SST

This section presents the comprehensive 4-way comparison for the simulation using the k- SST turbulence model, evaluating the performance of the full SNGR pipeline against the CFD input and DNS reference.

Figure 30 shows the full comparison of the stresses in the streamwise direction. It can be easily observed that SNGR and CFD underpredicts across the entire domain the mangnitude of the Reynolds stresses in this streamwise

direction, in comparison with the DNS analysis. However, the shape described by all four methods is overall pretty similar, getting the closest results at $x/h = 6.0$ and at $x/h = 10.0$. The largest difference occurs near the outlet, likely due to the massive pressure drop predicted by CFD, as described early in the report.

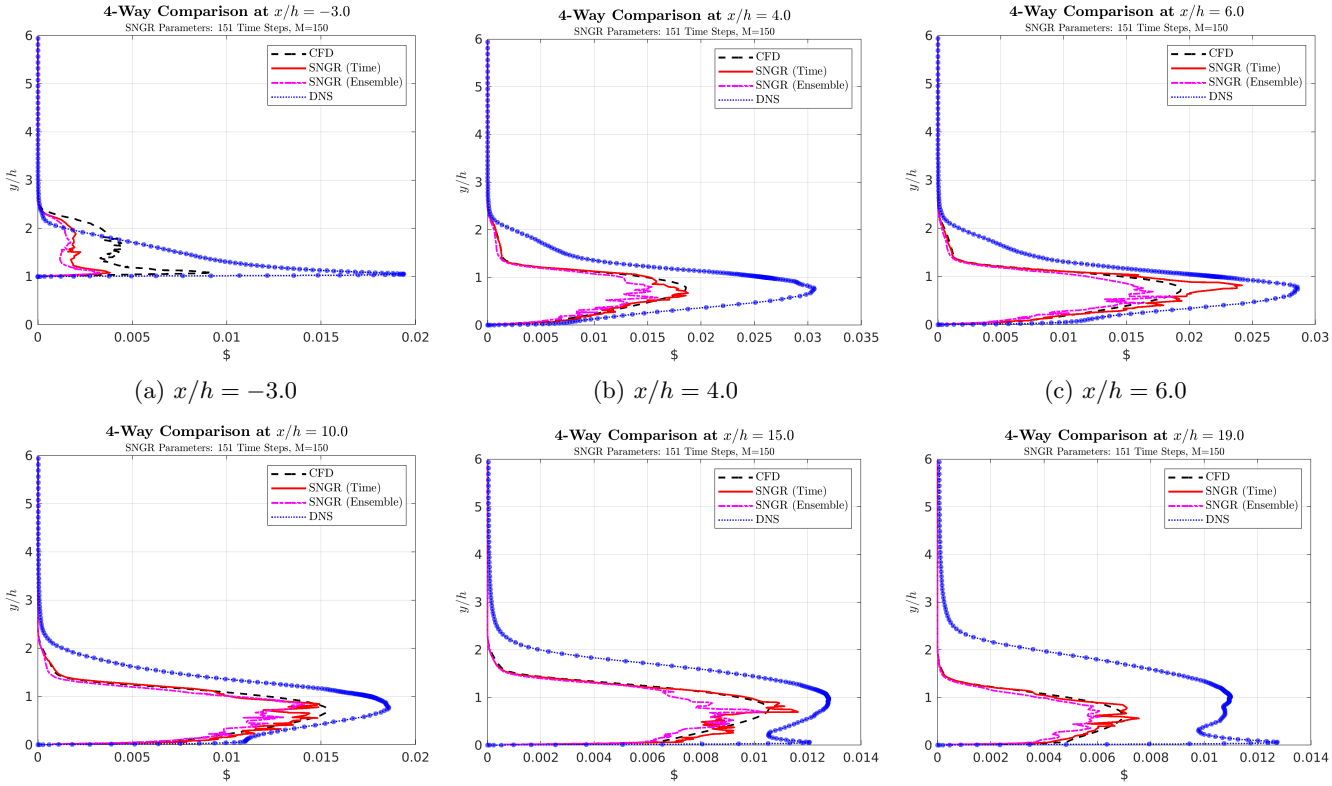

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 30: Reynolds stress $\langle u'u' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

The wall-normal Reynolds stresses are shown in Figure 30, describing a closer match between the four studied method compared to the streamwise component. Their difference is considerably smaller than what was observed for the equivalent slices and direction of the stresses, having considered the RSM Elliptic Blending as the turbulence model.

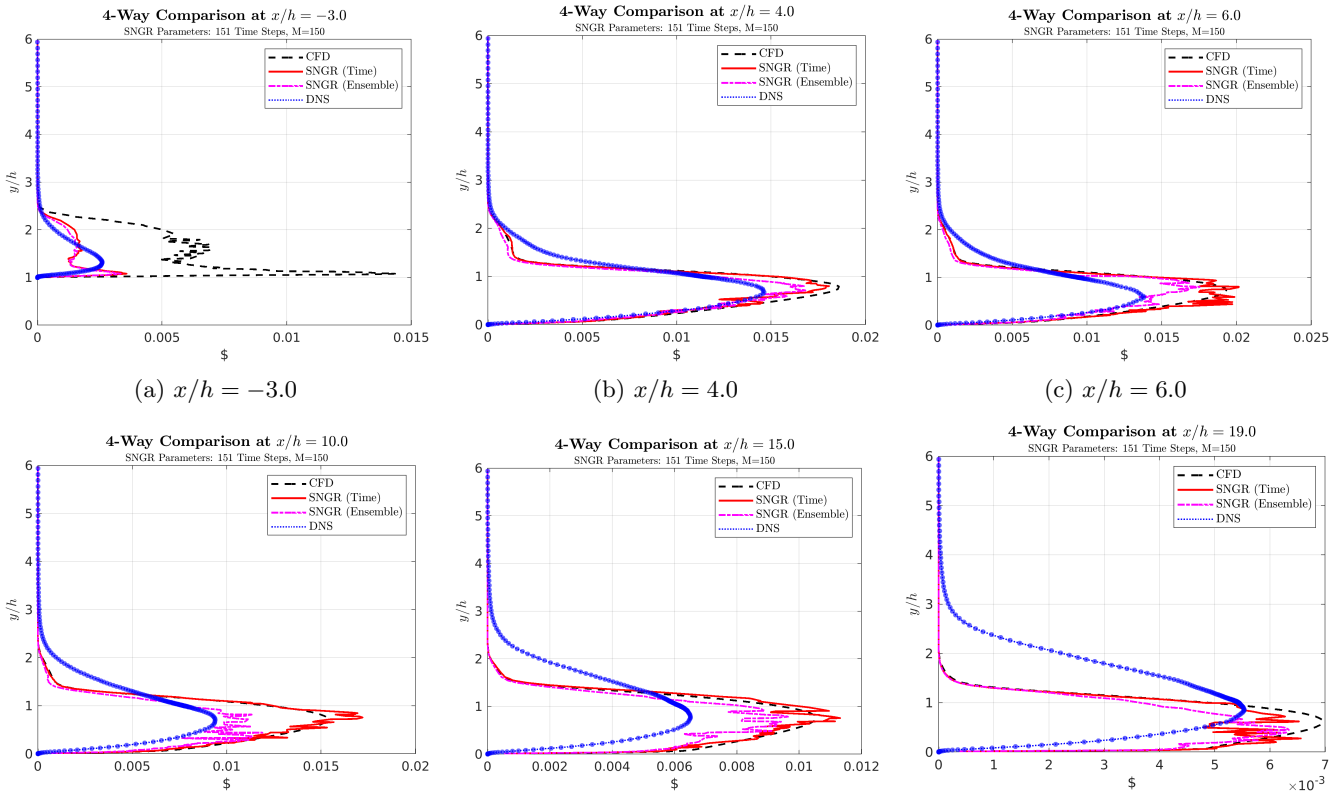

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 31: Reynolds stress $\langle v'v' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

With respect to the spanwise Reynolds stresses, the SST turbulence model plots the closest values of all components that have been evaluated, for both turbulence models. The inlet still describes a weird shape for CFD and SNGR values (as explained before), which seems not to affect the development of the flow downstream the step, as observed in Figure 32.

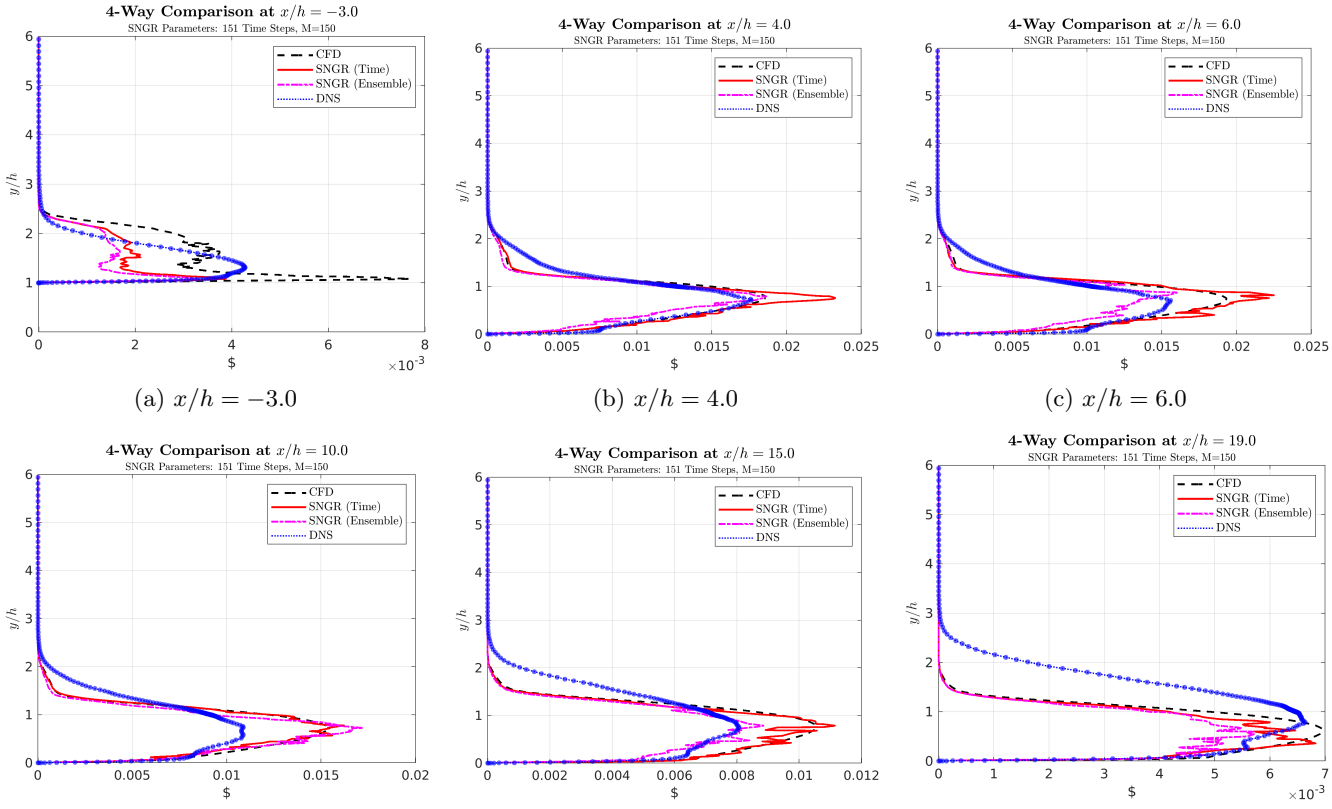

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 32: Reynolds stress $\langle w'w' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

The last component that is evaluated is - once more - the uv, which presents fairly close values between CFD and DNS (except near the inlet and the outlet) and massive difference when including the SNGR methods. In this case, the RSM Elliptic Blending model provided better results for the slice at $x/h = 19.0$ and also a quite good agreement for the other slices.

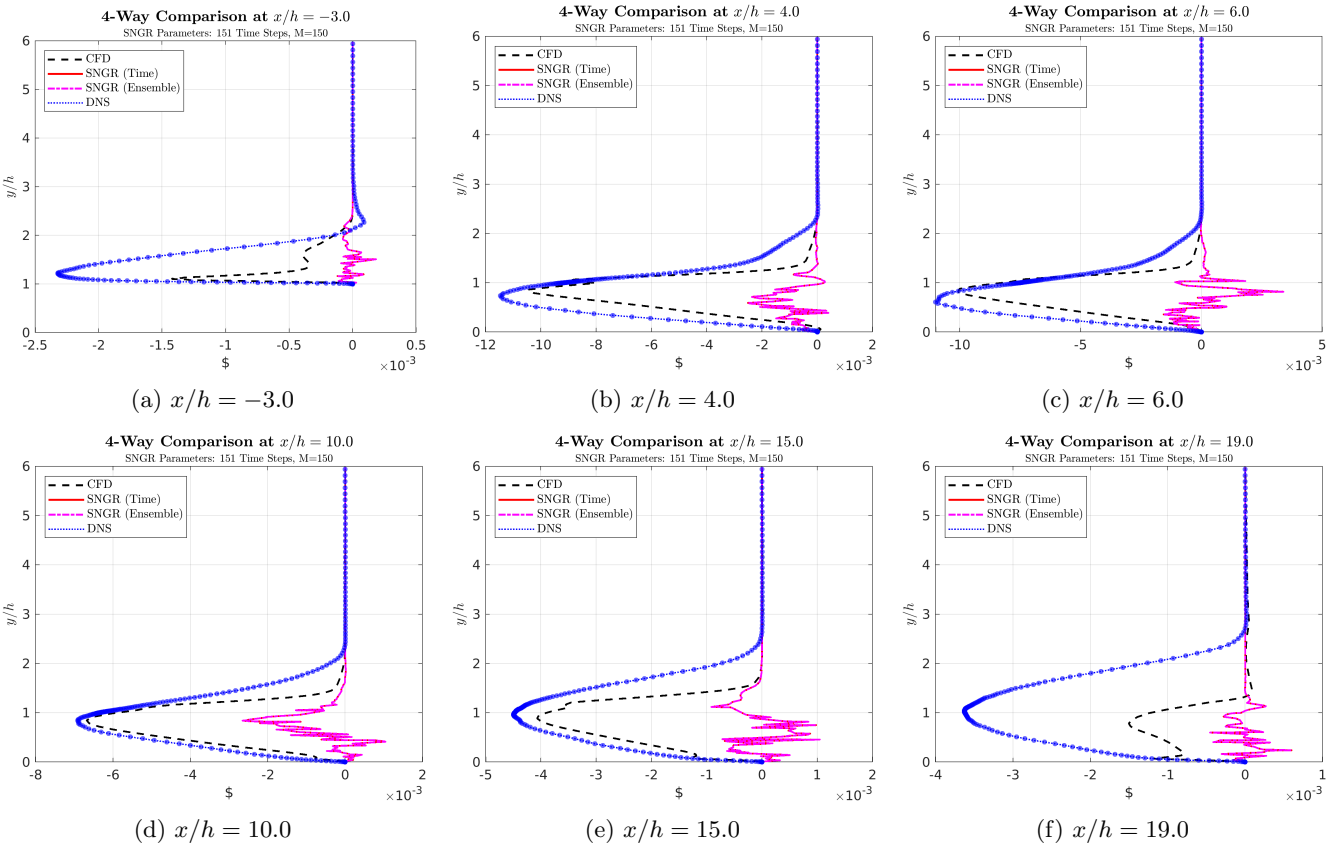

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 33: Reynolds stress $\langle u'v' \rangle$ (4-way comparison: CFD, SNGR_time, SNGR_ens, DNS) at six slices.

Analysis of the $\langle u'v' \rangle$ component discrepancy in SNGR. While the normal Reynolds stress components $\langle u'u' \rangle$, $\langle v'v' \rangle$, and $\langle w'w' \rangle$ generated by the SNGR method reproduce the expected magnitudes from CFD and DNS, the shear component $\langle u'v' \rangle$ shows significant deviation, remaining close to zero instead of the negative values characteristic of wall-bounded shear flows. Several factors can explain this behaviour:

(1) Coordinate or sign convention mismatch. The quantity $\langle u'v' \rangle$ represents the transport of streamwise momentum in the wall-normal direction and is expected to be negative in turbulent boundary layers, since high-momentum fluid moves downward ($-v'$) and low-momentum fluid moves upward ($+v'$). If the coordinate system used in SNGR differs from that in CFD or DNS (for instance, if the wall-normal axis is inverted), the sign of $\langle u'v' \rangle$ may flip or the correlation may cancel during averaging. Therefore, consistent definition of the wall-normal velocity component is crucial.

(2) Missing covariance in the anisotropy reconstruction. For an initially isotropic fluctuation field

$$\langle uu \rangle = \langle vv \rangle = \langle ww \rangle, \qquad \langle uv \rangle = \langle uw \rangle = \langle vw \rangle = 0,$$

a correct anisotropy step must apply a full linear transform (in this study, applying the Cholesky factor A such that $R = AA^T$) to generate the correlated field

$$\mathbf{U} = A\,\mathbf{u}. \tag{64}$$

This operation necessarily creates non-zero shear correlations when the off-diagonal terms of R are non-zero. A minimal illustrative example shows this explicitly. Let

$$U = u, \qquad V = v - 0.3\,u, \qquad W = w,$$

which corresponds to the linear transform

$$A = \begin{pmatrix} 1 & 0 & 0 \\ -0.3 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then, using only the isotropy properties of the seed field,

$$\langle UU \rangle = \langle u^2 \rangle, \tag{65}$$

$$\langle VV \rangle = \langle v^2 \rangle + 0.09\langle u^2 \rangle, \tag{66}$$

$$\langle UV \rangle = \langle u(v - 0.3u) \rangle = -0.3\,\langle u^2 \rangle \neq 0. \tag{67}$$

Equations (65)–(67) demonstrate that a proper anisotropic reconstruction *must* produce a non-zero Reynolds shear stress $\langle u'v' \rangle$ even if the seed field is isotropic.

Therefore, the persistent near-zero values of $\langle u'v' \rangle$ in the SNGR output imply that the current implementation does *not* apply a full covariance-preserving linear transform.

(3) Incomplete or simplified reconstruction of R_{ij}. When the Reynolds stress tensor is approximated through a Boussinesq model, as represented in Equation (68).

$$R_{ij} = \tfrac{2}{3} k \, \delta_{ij} - 2\nu_t S_{ij} \tag{68}$$

In the former expression, the off-diagonal terms depend directly on the local strain-rate tensor S_{ij}. If this tensor is noisy or locally weak, the resulting R_{12} may be underestimated, and the SNGR initialization loses the anisotropic shear component. Comparing the directly read CFD data with the R_{ij} used in SNGR helps identify this potential loss.

(4) Insufficient ensemble convergence. Although less likely, an inadequate temporal sampling may also lead to oscillations of $\langle u'v' \rangle$ around zero. The cross-correlation requires longer averaging times to converge than the variances of the normal components, since it depends on large-scale coherent motions. If the total simulated time does not cover several decorrelation periods of the shear layer, the computed mean remains unstable.

To sum everything up, the persistence of near-zero $\langle u'v' \rangle$ values in SNGR, despite accurate normal stresses, strongly suggests that inter-component coherence is not properly preserved in the reconstruction. This indicates that the fluctuation generation step lacks a covariance-based coupling (e.g., Cholesky or eigen decomposition of R_{ij}) or employs uncorrelated phase seeding between components. Correcting this aspect is expected to restore the negative $\langle u'v' \rangle$ characteristic of shear-driven turbulence and improve overall agreement with CFD and DNS data.

5.2 Streamwise Velocity Profiles

The representation of SNGR values is based on CFD data, directly depending on the RANS simulation. Plots are provided in the following subsections including individual representations considering only the SNGR method, combining these lines in a single plot, comparing them to DNS and CFD data in different graphs, as well as plotting all three methods together, either in 6-single slices or everything combined.

Unless otherwise stated, the streamwise velocity profiles presented in this section correspond to individual slices extracted at fixed spanwise locations and are **not spanwise-averaged**. The only exception is Figure 42b in Section 5.2.5, which explicitly presents spanwise-averaged results. This distinction is important when interpreting local versus integrated flow behavior: if the flow is statistically homogeneous in the spanwise (z) direction, the ensemble mean does not depend on z, so the spanwise average

$$\overline{u}(x,y) \equiv \frac{1}{L_z}\int_0^{L_z} u(x,y,z)\,\mathrm{d}z \tag{69}$$

theoretically converges to the same mean profile as a single representative z-plane. In practice, however, the finite domain length and finite sampling cause differences: spanwise averaging reduces sampling noise and yields a smoother, more converged estimate of mean and second-order statistics, whereas a single plane preserves local fluctuations and can be noisier. Therefore, when profiles are single-plane (showing a local behaviour) are noisier than when they are spanwise-averaged (integrated).

5.2.1 Individual slices - SNGR single time-loop approach

Considering the same six slices evaluated along the entire study - one upstream placed at $x/h = -3.0$ - the streamwise velocity profiles are plotted.

It is relevant to mention that the free-stream velocity is set to 1.0 m/s in *StarCCM+*, and as expected, it is only reached in the section upstream of the step.

Results for the Elliptic Blending Reynolds Stress Model

Figure 34 illustrates the evolution of the streamwise velocity profiles across the six selected slices.

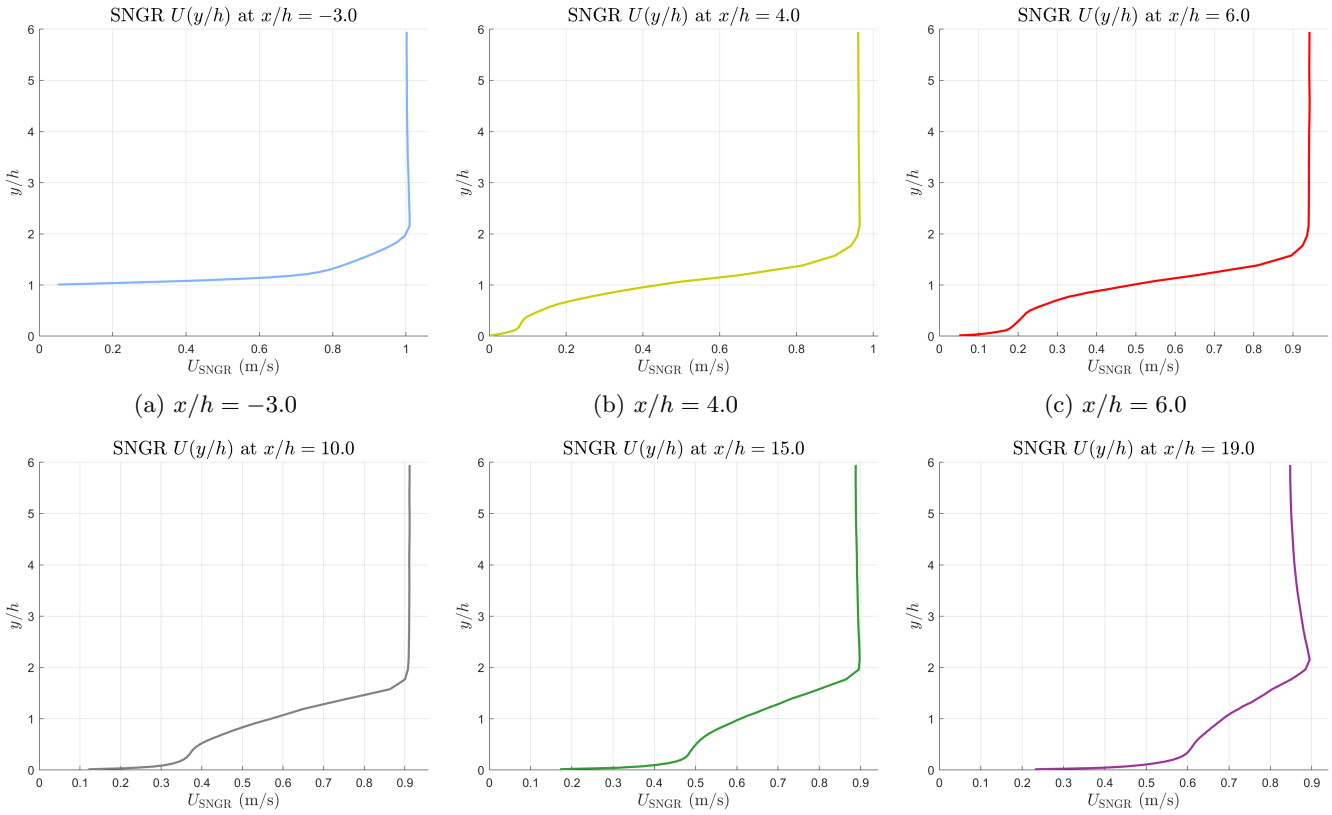

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 34: SNGR streamwise velocity profiles $U(y)$ at six downstream slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0).

Upstream of the step - observed in Figure 34a, the profile is nearly uniform, reaching the free-stream value of $1.0\,\mathrm{m\,s^{-1}}$, with only a thin boundary layer near the wall. This confirms that the inlet conditions are well-developed before the step.

Immediately downstream of the step ($x/h = 4.0$ and 6.0, represented in Figure 34b and Figure 34c, respectively), a strong velocity deficit appears close to the wall due to the separation bubble. The reversed flow region that is expected from this kind of flow is missing, suggesting there is no recirculation (negative flow speed) behind the step. Upwards of the separated layer the flow shows a modest overshoot above the free-stream speed, which is physically consistent with mass conservation as the fluid accelerates around the separated zone.

Further downstream, at $x/h = 10.0$ shown in Figure 34d, the recirculation zone begins to diminish, and the reversed flow near the wall becomes weaker. The velocity gradient close to the wall increases, indicating the onset of reattachment and the redevelopment of the boundary layer.

Finally, at $x/h = 15.0$ and 19.0, displayed in Figure 34e and Figure 34f, respectively, the profiles gradually recover toward a canonical turbulent boundary layer shape. However, even at $x/h = 19.0$, the wake effect persists, and the velocity near the wall remains lower than the free-stream value, showing that full recovery is not yet achieved within the studied domain.

Results for the RANS k-ω SST

Figure 35 illustrates the evolution of the streamwise velocity profiles across the six selected slices, based on the k-ω SST simulation.

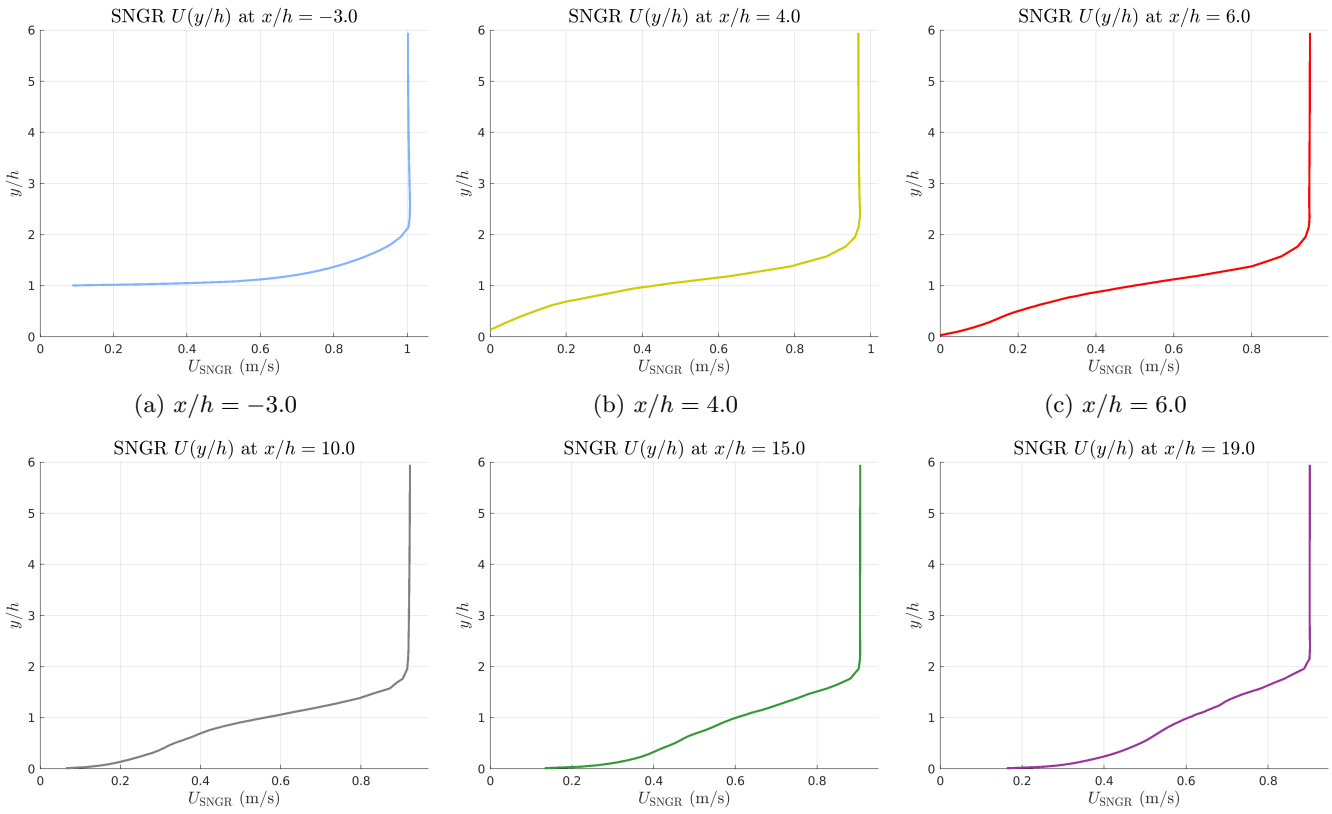

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 35: SNGR streamwise velocity profiles $U(y)$ at six downstream slices, based on k- SST input (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0).

The profiles described by the single time-loop approach are smoother when considering the RANS k–ω SST as the turbulence model than the RSM EB.

The profile is essentially uniform across most of the channel height and reaches the prescribed free-stream value $U \approx 1.0$ m/s, with only a very thin boundary layer close to the wall (visible as the small gradient near $y/h \approx 0$). This confirms that the inlet conditions and the upstream RANS boundary layer are well developed and that the SNGR synthesis preserves the expected upstream mean velocity.

At these slices the SNGR profiles show a pronounced velocity deficit in the near-wall region: U_{SNGR} is significantly lower close to the wall than in the outer flow. This behaviour is the signature of the separation bubble formed directly after the step. Compared with the RSM EB case, the k–ω SST–based SNGR profiles are noticeably

smoother: the deficit is spatially coherent and shows fewer small-scale fluctuations. instead a strong reduction of the streamwise velocity (but generally still non-negative) is observed.

Importantly, the k–ω SST based SNGR profiles do not exhibit a substantial reversed mean velocity in the near-wall region (i.e. $U \gtrsim 0$ for $y/h \lesssim 1$). This indicates that SNGR underpredicts the local recirculation strength observed in DNS studies of BFS flows and therefore does not fully capture the separation intensity. This may be due to limited near-wall resolution, RANS closure limitations, or incomplete covariance in the fluctuation reconstruction.

The outer region (away from the wall) recovers toward the free-stream velocity faster than the near-wall region, reflecting the separated shear layer that convects downstream.

At $x/h = 10.0$ the deficit associated with the recirculation diminishes and the near-wall velocity gradient increases. This indicates the beginning of reattachment and the redevelopment of the boundary layer. The k–ω SST–based SNGR profile shows a steady increase of U with y/h near the wall consistent with mass conservation as the separated region closes.

Profiles at $x/h = 15.0$ and $x/h = 19.0$ display progressive recovery toward a canonical turbulent boundary layer shape: the near-wall velocity rises and the overall profile becomes fuller. At $x/h = 19.0$ the flow has regained much of its momentum, but the velocity near the wall remains slightly lower than the free-stream value, indicating that complete recovery is not yet achieved inside the domain considered.

5.2.2 Streamwise Velocity Profiles: CFD vs. SNGR single time-loop approach

As SNGR data is based on CFD input, its comparison provides useful insights about how close the SNGR time-loop approach recreates data from CFD.

Results for the Elliptic Blending Reynolds Stress Model

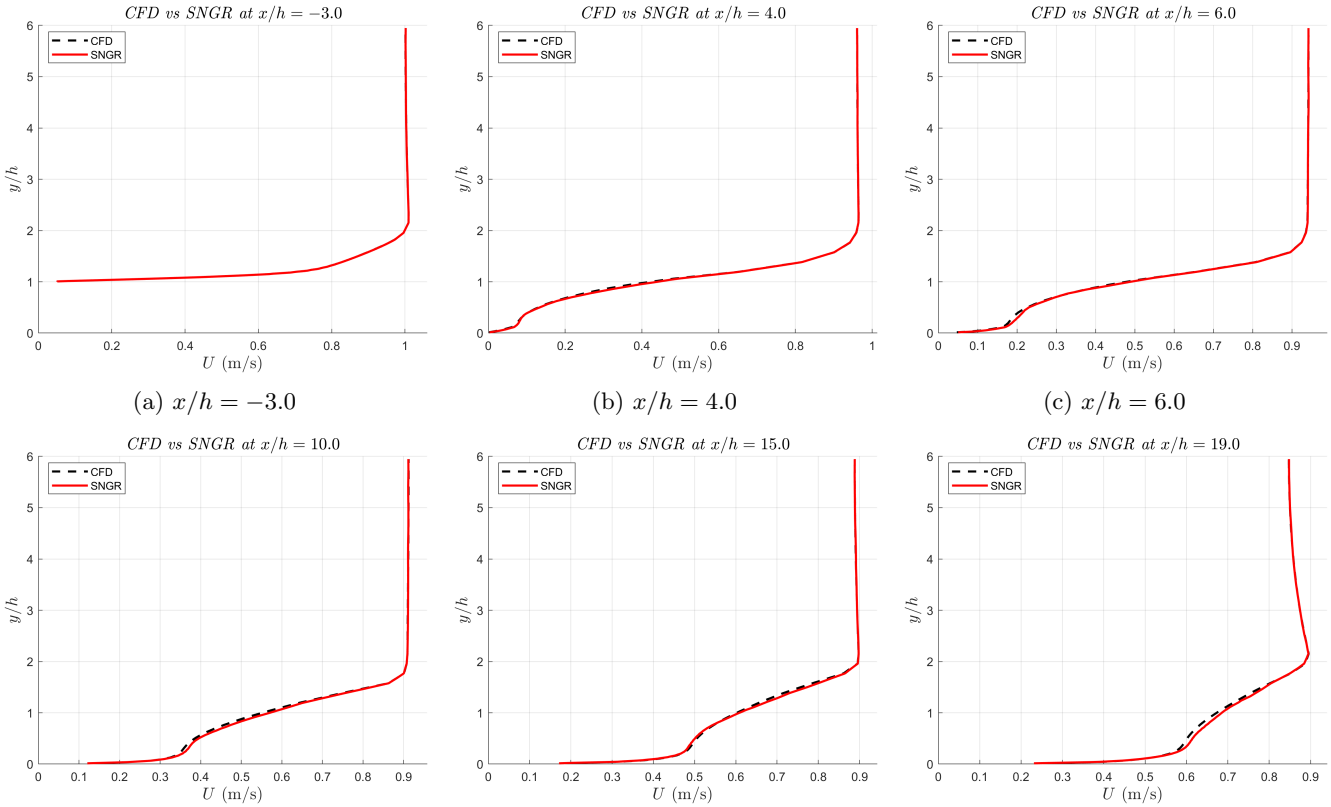

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 36: Comparison of streamwise velocity profiles $U(y)$ — CFD vs SNGR — at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0).

Overall, the agreement between the SNGR predictions and the CFD results is very good across all six slices. The SNGR fluctuations are constructed to be zero-mean ($\langle u' \rangle_= 0$) and, where enforced, approximately divergence-free, so they do not modify the prescribed RANS mean $\overline{U}$. Any apparent mean shift within a single time loop is a finite-sampling artefact that diminishes with longer averaging and vanishes in the ensemble (or long-time) mean.

Upstream of the step (showing in Figure 36a the slice corresponding to $x/h = -3.0$), both profiles coincide almost perfectly, confirming that the inlet conditions are consistent. Immediately downstream of the step (represented the values in Figure 36b and Figure 36c, indicating $x/h = 4.0$ and 6.0 slice, respectively), the SNGR profiles do not reproduce the velocity deficit and reversed flow region that is expected from DNS studies, as previously mentioned. SNGR shows a slightly less pronounced overshoot in the outer region compared to CFD, suggesting a minor underestimation of the acceleration caused by mass conservation.

It is important to clarify that the CFD curves shown here are not taken from the StarCCM+ post-processing interface; they are extracted from the CGNS/HDF5 file and imported into MATLAB. In this imported dataset the CFD streamwise velocity does not exhibit negative values at these slices, meaning that the recirculation bubble predicted by StarCCM+ is either very shallow or not captured by the specific slice and resolution extracted from the CGNS file. The SNGR profiles therefore correctly match the MATLAB-imported CFD data, even though both differ from the recirculation strength expected in DNS.

At $x/h = 10.0$ - presented in Figure 36d, the agreement remains strong, with both methods capturing the reduction of the recirculation zone and the steep velocity gradient near the wall. Further downstream (one meter after the step - Figure 36e, $x/h = 15.0$ and very close to the outlet - shown in Figure 36f, $x/h = 19.0$), the SNGR profiles continue to follow the CFD trend, although a small discrepancy persists in the near-wall region, where SNGR predicts slightly lower velocities than CFD. This indicates that the recovery of the boundary layer is somewhat delayed in the SNGR solution compared to the CFD reference.

In summary, the SNGR approach demonstrates excellent predictive capability for the overall velocity distribution and separation behavior compared to its input data, with only minor deviations in the outer region immediately

after the step and in the near-wall recovery far downstream.

It is important to note that the mean velocity and the overall separation behaviour are not predicted by SNGR itself: they are entirely prescribed by the RANS input field. SNGR only synthesises zero-mean turbulent fluctuations that are convected by this RANS mean flow, and therefore it cannot create or alter features such as the recirculation bubble or the outer-layer acceleration. The good agreement observed in the profiles thus reflects the consistency of the fluctuation reconstruction with the imposed RANS mean field rather than an independent prediction of the mean flow by SNGR.

Results for the RANS k-omega SST

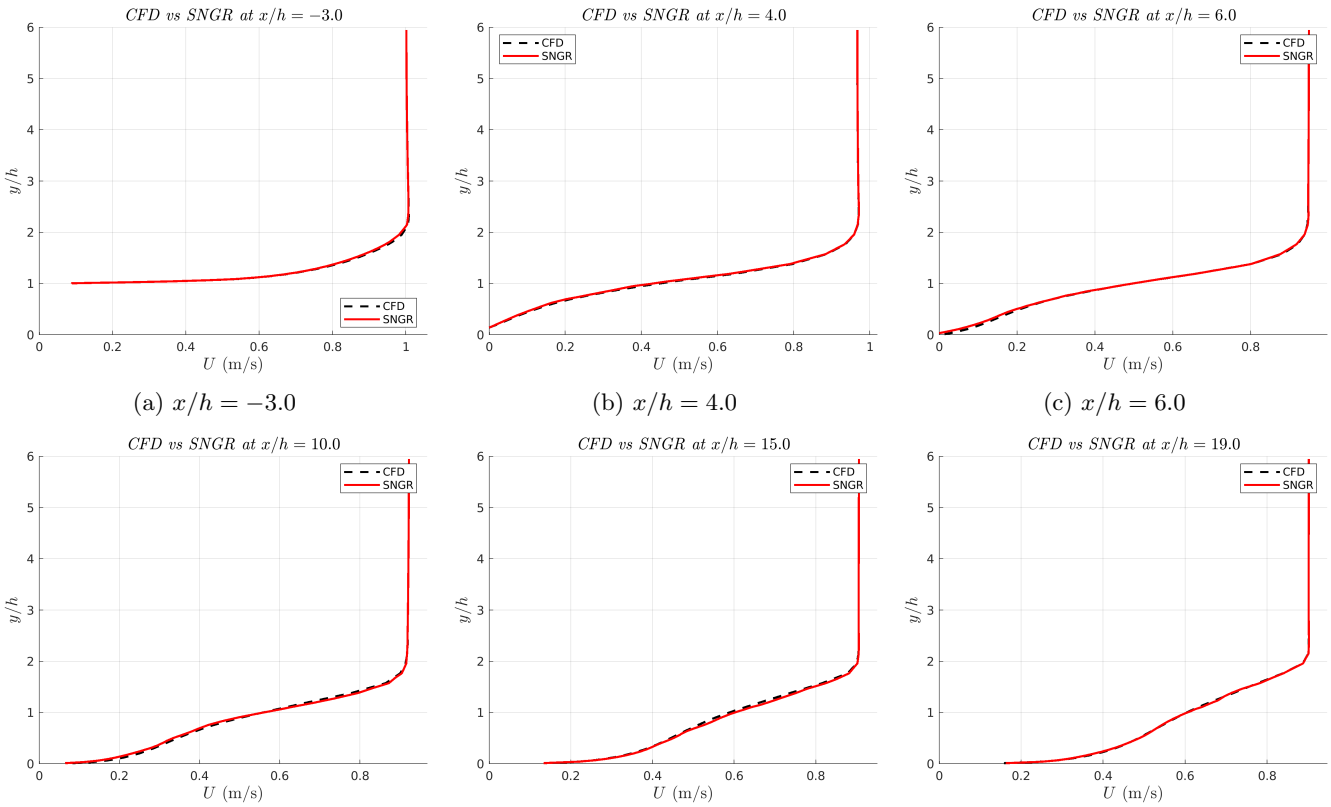

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 37: Comparison of streamwise velocity profiles $U(y)$ — CFD (k- SST) vs SNGR — at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0).

Overall, the agreement between the SNGR predictions and the CFD results based on the RANS k-ω SST model is again very close across all six slices. The mean velocity profiles coincide almost perfectly both upstream and downstream of the step, confirming that the SNGR approach successfully reproduces the flow separation, reattachment, and recovery trends predicted by CFD. Negligible differences appear at specific regions: either very clsoe to the bottom wall or at the step´s height. In summary, the SNGR solution driven by the k–ω SST input shows excellent consistency with the CFD reference, demonstrating that the time-loop approach accurately preserves the main flow features across all streamwise positions.

5.2.3 Streamwise Velocity Profiles: DNS vs. SNGR single time-loop approach

It is important to emphasize that SNGR does not compute or modify the mean flow: it only synthesizes zero-mean fluctuations on top of the RANS mean velocity field. Therefore, the purpose of examining the SNGR streamwise velocity profiles is not to assess how well SNGR "predicts" the mean flow, but rather to verify that the mean of the generated fluctuations indeed remains close to zero for the available sampling time. In other words, the following expression is checked

$$\langle u'(x,y)\rangle_{\text{time}} \approx 0, \tag{70}$$

so that the reconstructed SNGR mean velocity is statistically consistent with the imposed RANS mean field. Any departure from the RANS profile is interpreted as insufficient statistical convergence of the time-loop or residual numerical bias in the fluctuation synthesis—not as a physical prediction error. This distinction is essential when comparing the SNGR and DNS mean profiles.

Results for the Elliptic Blending Reynolds Stress Model

Figure 34 compares the streamwise velocity profiles $U(y)$ predicted by SNGR against DNS data at six streamwise locations. Overall, the agreement is strong, but some differences are observed depending on the region.

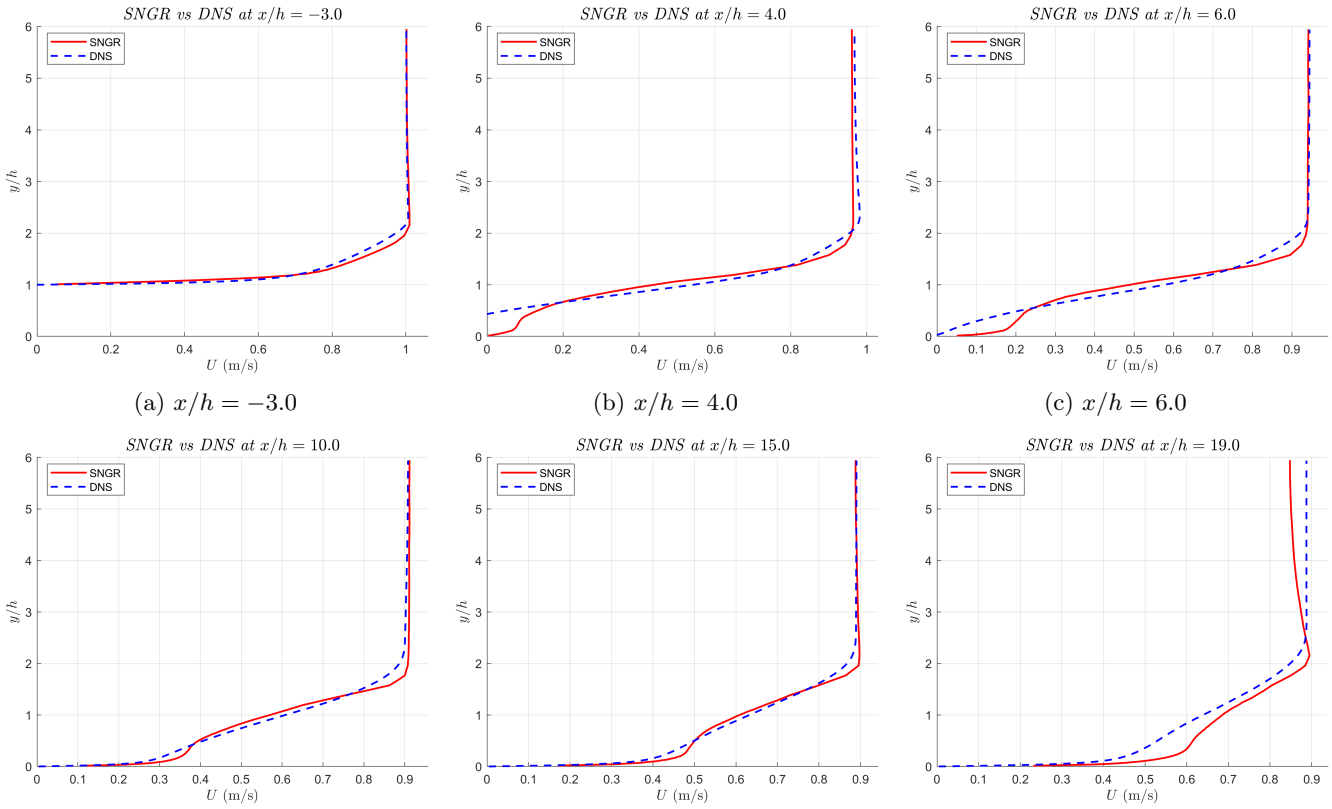

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 38: Comparison of streamwise velocity profiles $U(y)$ — SNGR vs DNS — at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0).

Upstream of the step at $x/h = -3.0$ - exhibited in Figure 40 - both profiles coincide almost perfectly, confirming that the inlet conditions are consistent and well-represented, with only a slight overestimation of the velocity values between $1 < y/h < 2$ predicted by the SNGR method compared to DNS.

Immediately downstream of the step (at $x/h = 4.0$), SNGR does not reproduce the main features of the separated flow. In the DNS profile, represented in Figure 38b, the streamwise velocity becomes negative over a substantial portion of the lower channel, with reversed flow extending approximately up to $y/h \approx 0.4$. This behaviour is characteristic of the recirculation bubble forming just after the backward-facing step. In contrast, the SNGR-generated profiles do not exhibit negative velocities, indicating that the recirculation predicted by DNS is essentially absent. This discrepancy arises because SNGR does not compute the mean flow: it simply adds zero-mean fluctuations to the RANS mean field. Since the RANS solution used here predicts only a very weak or shallow separation, the SNGR mean profile follows it closely and therefore cannot recover the stronger DNS backflow.

Additionally, the outer region acceleration (overshoot above the free-stream velocity) is less pronounced in SNGR, suggesting a weaker prediction of the shear layer growth. A bit further downstream, the gap between DNS and SNGR values increase at the region just over the step's height, as seen in Figure 38c between $1.3 < y/h < 2$.

At $x/h = 10.0$, indicated in Figure 38d, the profiles converge well, with SNGR accurately predicting the reduction of the recirculation zone and the steep velocity gradient near the wall. Minor differences remain in the wake region, where SNGR predicts slightly lower velocities than DNS, and a bit higher as the location moves upwards y/h.

Further downstream at $x/h = 15.0$, Figure 38e shows recovery toward a canonical turbulent boundary layer. Nevertheless, SNGR exhibits a small lag in boundary layer redevelopment, with slightly lower near-wall velocities than DNS, while the outer region matches well.

The largest disagreement between both methods is at $x/h = 19.0$, presented in Figure 38f. SNGR overpredicts the streamwise velocity profile from $0 < y/h < 2.5$ and underpredicts it compared to DNS in the section between $2.5 < y/h < 6$.

It is also important to note that the near-wall region in both SNGR and DNS profiles (particularly in Figures 38b–38f) exhibits some irregular behavior. This may be attributed to insufficient wall-normal resolution in the mesh, as the non-dimensional wall distance y^+ exceeds 1 in several regions, particularly near the step and downstream. Since the viscous sublayer is not fully resolved, this can lead to inaccuracies in capturing reversed flow and near-wall velocity gradients. The impact of y^+ on turbulence modeling and velocity profile fidelity is discussed in Section 4.2.2.

In summary, SNGR provides a consistent reconstruction of fluctuations around the imposed RANS mean field and thus reproduces the general outer-layer velocity distribution. However, SNGR cannot recover the DNS recirculation bubble or the sharp gradients associated with separation and reattachment, because these features originate from the mean flow, which in SNGR is prescribed entirely by the RANS input. The remaining discrepancies between SNGR and DNS are therefore expected and arise from the synthetic nature of the turbulence reconstruction and the limitations of the underlying RANS field.

Results for the RANS k-omega SST

Figure 39 compares the streamwise velocity profiles $U(y)$ predicted by SNGR against DNS data at six streamwise locations, using the k- SST model as the RANS input.

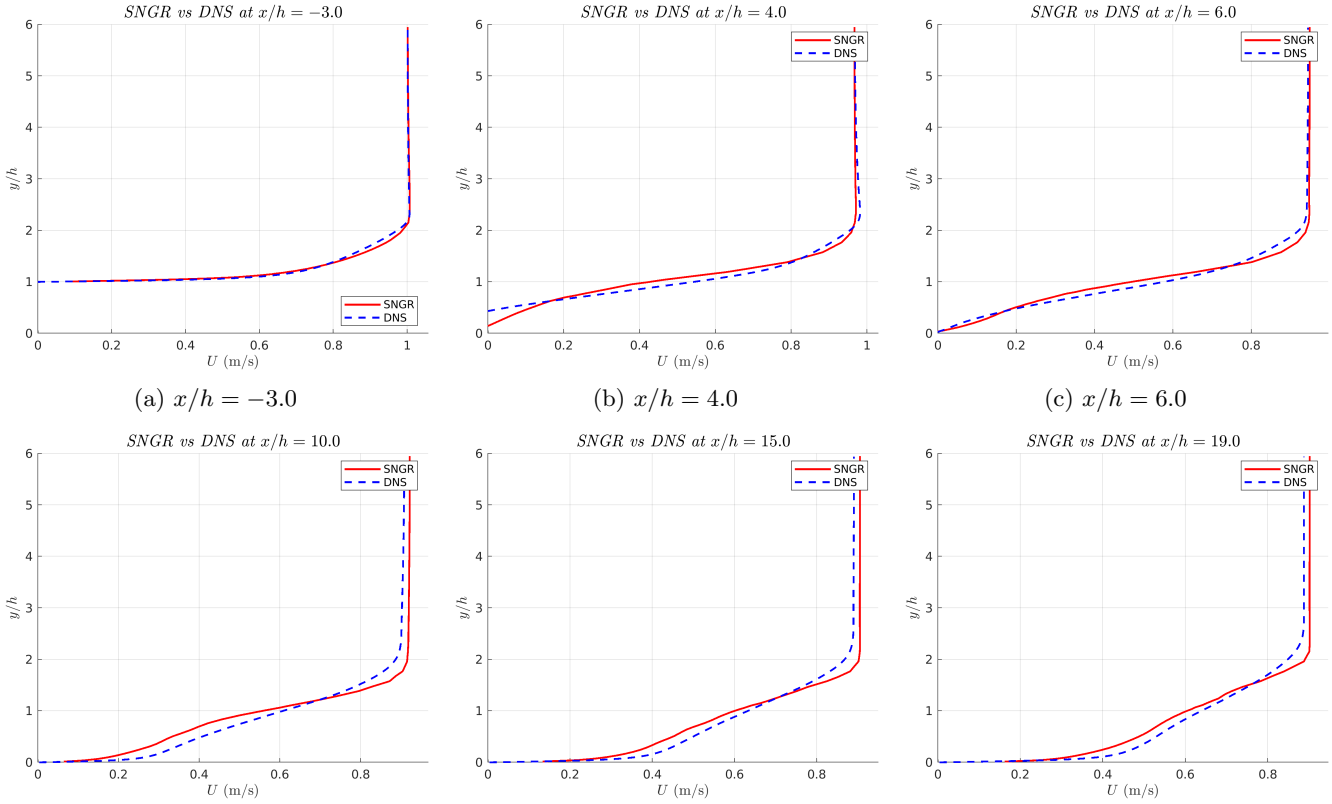

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 39: Comparison of streamwise velocity profiles $U(y)$ — SNGR (from k- SST) vs DNS — at six slices (top row: $x/h = -3.0$, 4.0, 6.0; bottom row: $x/h = 10.0$, 15.0, 19.0).

The SNGR single time-loop approach provides a better agreement with respect DNS up to $x/h = 6.0$ considering

the RANS k-ω SST model than the RSM EB, and also at the slice close to the outlet, at $x/h = 19.0$. However, when using as input data values from the RSM turbulence model simulation, the profiles converge closer to each other at slices $x/h = 10.0$ and $x/h = 15.0$, with SNGR accurately predicting the reduction of the recirculation zone and the steep velocity gradient near the wall.

5.2.4 Streamwise Velocity Profiles: DNS vs. CFD vs. SNGR single time-loop approach

Figure 40 presents the streamwise velocity profiles $U(y)$ at six streamwise locations, comparing three approaches: CFD (RSM EB), SNGR, and DNS.

While the previous section 5.2.3 quantified the absolute accuracy of SNGR against the DNS benchmark (addressing RQ2), this three-way comparison addresses RQ1 by isolating the sources of those discrepancies. By including the CFD (RANS) profiles, it is possible to distinguish between inherited model bias (deviations originating from the RANS input data, where CFD $\neq$ DNS) and reconstruction uncertainty (deviations introduced by the stochastic synthesis, where SNGR $\neq$ CFD). This distinction is critical for determining whether improvements should target the underlying turbulence model or the SNGR algorithm itself.

Results for the Elliptic Blending Reynolds Stress Model

Overall, all methods capture the main flow features, but some differences are observed in specific regions.

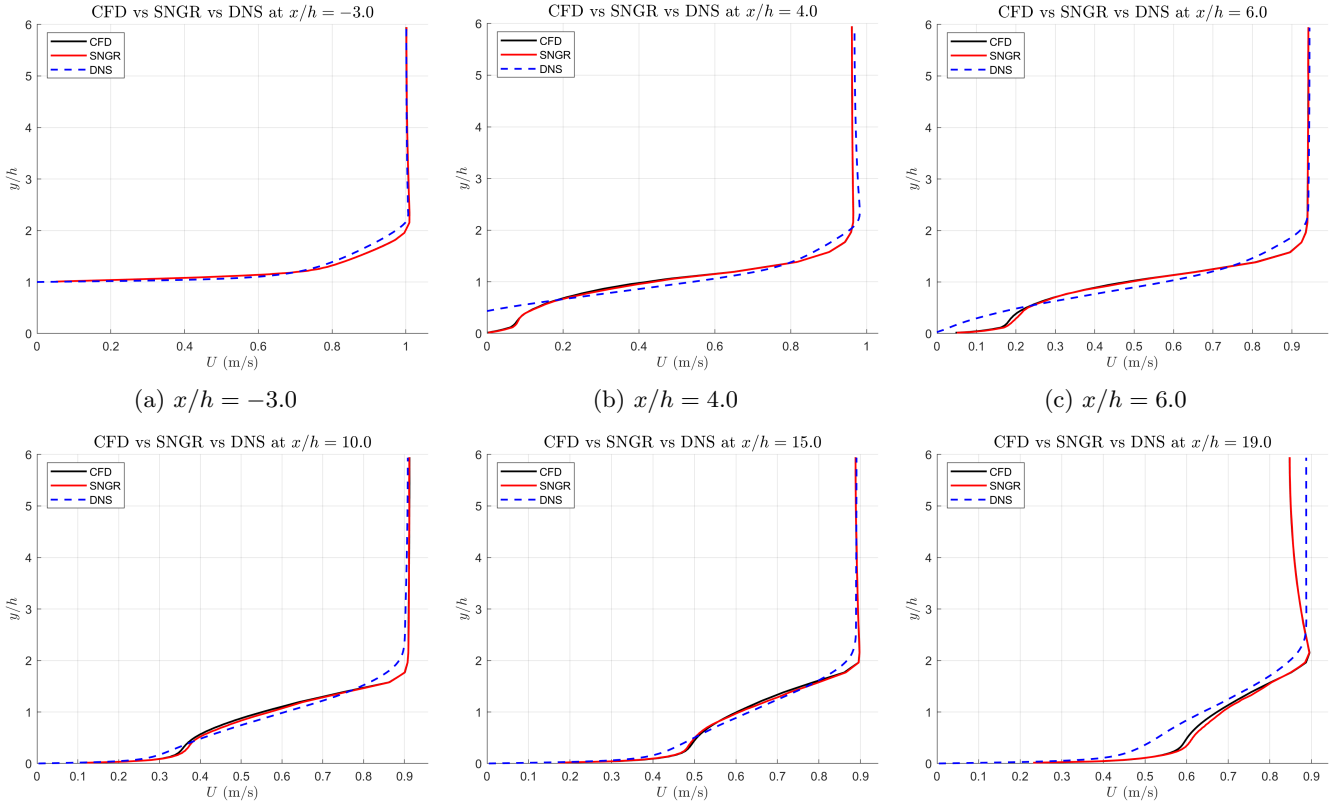

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 40: Comparison of velocity profiles at six streamwise locations. Results are shown for three methods: CFD (RSM EB), SNGR (time-marching), and DNS.

Upstream of the step ($x/h = -3.0$) - shown in Figure 40a - all three profiles coincide almost perfectly, confirming that the inlet conditions are consistent across methods.

Downstream the step, at $x/h = 4.0$ and 6.0, which are equivalent to $0.2m$ and $0.3m$ and are displayed in Figure 40b and Figure 40c, DNS shows the strongest reversed flow and the most pronounced outer-layer overshoot, which is typical of high-fidelity simulations. CFD closely follows DNS in these regions, while SNGR slightly underpredicts the reversed flow magnitude and the overshoot. This suggests that SNGR smooths the shear layer development compared to the other two methods.

At $x/h = 10.0$, halfway between the inlet and outlet of the geometry model, the three profiles converge significantly in the region close to the bottom wall at $y/h = 0$, and up to $y/h = 1.5$, as observed in Figure 40d. SNGR and CFD both capture the reduction of the recirculation zone and the steep velocity gradient near the wall, although SNGR still predicts slightly lower velocities in the near-wall region compared to DNS.

Further downstream at $x/h = 15.0$, all methods show recovery toward a canonical turbulent boundary layer, as exhibited in Figure 40e. CFD aligns very closely with SNGR, as seen in 5.2.2 exhibits a small lag in boundary layer redevelopment compared to DNS, with slightly lower near-wall velocities. However, the outer region matches pretty well, indicating that SNGR correctly predicts the overall momentum recovery trend.

Finally, as mentioned in previous subsections, the effect of the outlet causes a noticeable difference in the estimation of streamwise velocities by CFD - and thus, SNGR too - compared to DNS, as represented in Figure 40f.

To sum it up, DNS serves as the reference, showing the most detailed flow structure. CFD (RSM EB) provides a reasonable agreement with DNS across all regions, validating its accuracy for this case. SNGR demonstrates strong predictive capability for the overall velocity distribution and separation behavior, with only minor deviations in the near-wall region and immediately after the step. The reasoning behind these differences has been previously explained.

Results for the RANS k-omega SST

Figure 41 presents the streamwise velocity profiles $U(y)$ at six streamwise locations, comparing three approaches: CFD (k- SST), SNGR, and DNS.

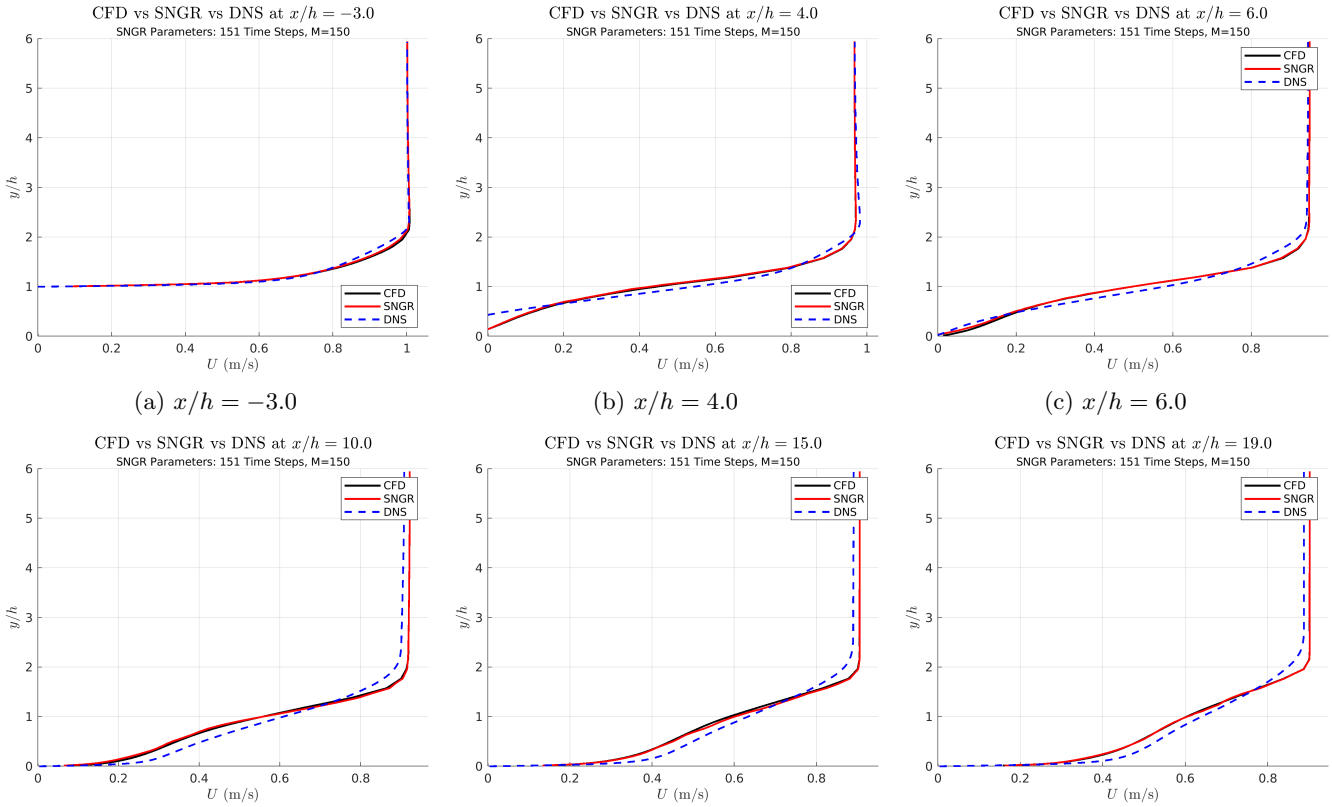

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 41: Comparison of velocity profiles at six streamwise locations. Results are shown for three methods: CFD (k- SST), SNGR (time-marching), and DNS.

As thoroughly explained before, this comparison describes an overall closer match between the three main method being considered in this study, with the biggest difference with respect the other turbulence model the much better definition of the velocity profile at $x/h = 19.0$.

5.2.5 Streamwise Velocity Profiles: Combined profiles

In this subsection, the results from the preceding individual slice comparisons are synthesized into combined representations to highlight the overall performance of the SNGR approach relative to CFD and DNS. The combined plots provide a holistic view of the velocity field evolution along the streamwise direction, allowing trends in flow recovery, separation, and overshoot to be more easily identified.

While the previous subsections quantified local discrepancies at specific streamwise locations, this section synthesizes the results into combined representations to assess the **global spatial consistency** and the **streamwise evolution** of the synthetic flow. This addresses the capability of SNGR to reproduce not just local magnitudes, but also the correct rates of flow recovery, separation bubble elongation, and boundary layer growth.

Results for the Elliptic Blending Reynolds Stress Model

Figure 42 presents two perspectives: (a) the combined SNGR velocity profiles across all slices, illustrating the general evolution of the flow predicted by the SNGR method, and (b) a spanwise-averaged comparison between CFD and SNGR, showing how well SNGR reproduces the CFD input in an integrated form.

Figure 43 extends this analysis to broader comparisons: (a) SNGR versus DNS across all slices to evaluate the stochastic reconstruction against high-fidelity reference data, and (b) a comprehensive comparison of CFD, SNGR, and DNS across all slices, highlighting the relative accuracy of each method in capturing the main flow features and the development of the turbulent boundary layer.

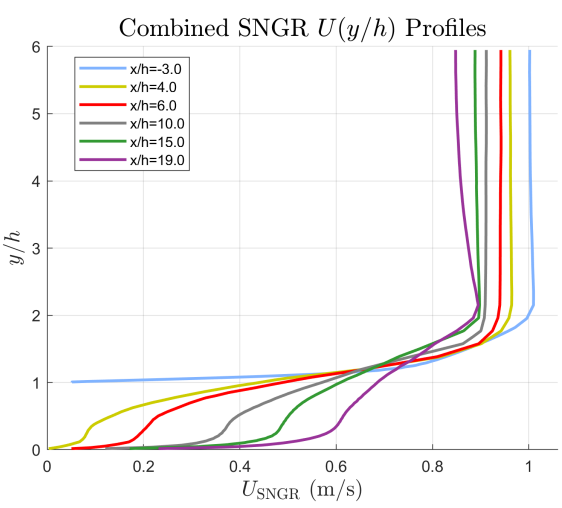

(a) Combined SNGR velocity profiles across all slices

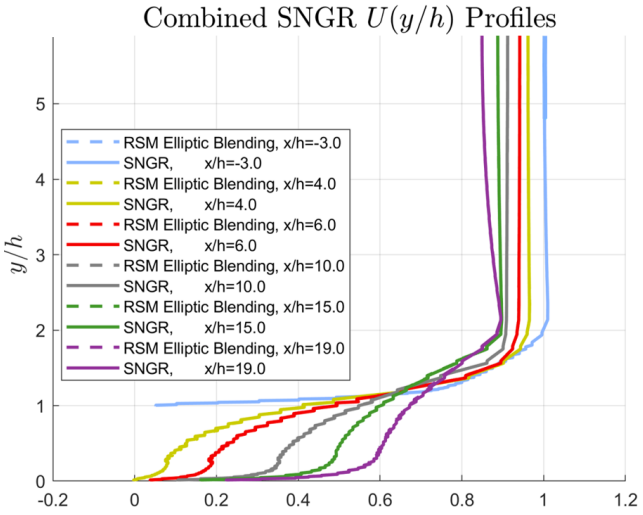

(b) Spanwise-averaged velocity profiles: CFD vs SNGR

Figure 42: (a) Combined SNGR velocity profiles across all slices. (b) Spanwise-averaged velocity profiles comparing CFD (dashed lines) and SNGR (solid lines); colors indicate the streamwise slice locations $x/h = -3.0$, 4.0, 6.0, 10.0, 15.0, 19.0.

Figure 42 shows that SNGR (represented by using solid lines in 42a) and CFD (exhibiting its data by means of dashed lines, as observed in Figure 42b) profiles agree closely across all slices. At $x/h = -3.0$ - represented using a blue line - both methods match almost perfectly, confirming consistent inlet conditions. At $x/h = 4.0$ and $x/h = 6.0$ - exhibited in yellow and red lines, respectively - both capture the strong velocity deficit and reversed flow near the wall, although SNGR slightly underpredicts the overshoot compared to CFD. At $x/h = 10.0$ - corresponding to the grey line - the agreement remains strong, with only minor differences in the outer region. Farther downstream at $x/h = 15.0$ and $x/h = 19.0$ - looking at the dark green and purple lines, respectively - both methods predict the recovery toward a turbulent boundary layer, though SNGR shows a small lag in near-wall velocity compared to CFD. Overall, the color-coded comparison confirms that SNGR reproduces the main flow features with high fidelity, with only minor deviations in overshoot and recovery rate.

With respect to the comparison between SNGR and DNS methods, Figure 43 shows that the predictions of both approaches are very similar above the step for most slices, with the notable exception of the slice closest to the outlet at $x/h = 19.0$. Here, the light continuous blue line representing the SNGR values shows moderately lower velocities compared to DNS. In the near-wall region ($y/h < 1$), differences are generally moderate. The closest agreement occurs at $x/h = 10.0$ and $x/h = 15.0$, where SNGR and DNS overlap closely. Moderate discrepancies appear immediately downstream of the step at $x/h = 4.0$ and $x/h = 6.0$. For the upstream slice at $x/h = -3.0$,

the profiles essentially overlap, indicating excellent agreement. The largest difference appears again at $x/h = 19.0$, reflecting a slight lag in the SNGR prediction of flow recovery in the far downstream region.

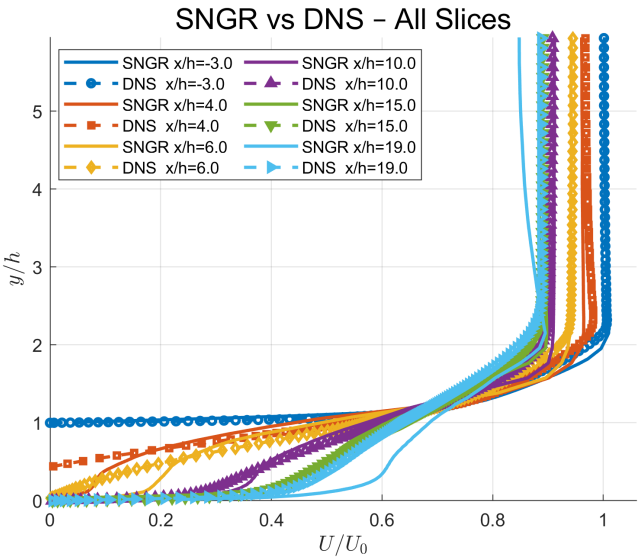

(a) All-slice velocity profile comparison: SNGR vs DNS

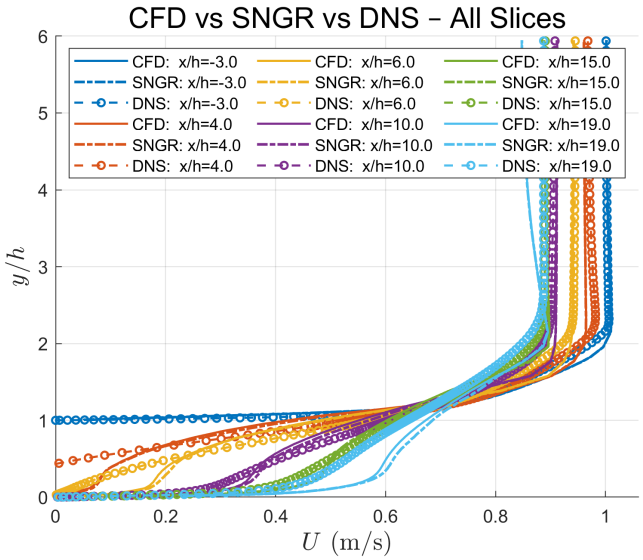

(b) All-slice velocity profile comparison: CFD vs SNGR vs DNS

Figure 43: Comparison of streamwise velocity profiles across all slices: (a) SNGR vs DNS, showing the agreement between the stochastic reconstruction and high-fidelity reference data; (b) CFD (RSM EB), SNGR, and DNS combined, highlighting the relative performance of each approach at multiple streamwise locations.

Results for the RANS k-omega SST

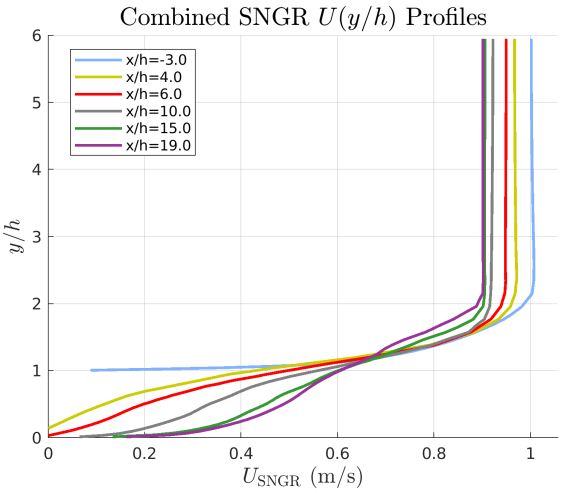

(a) Combined SNGR velocity profiles across all slices

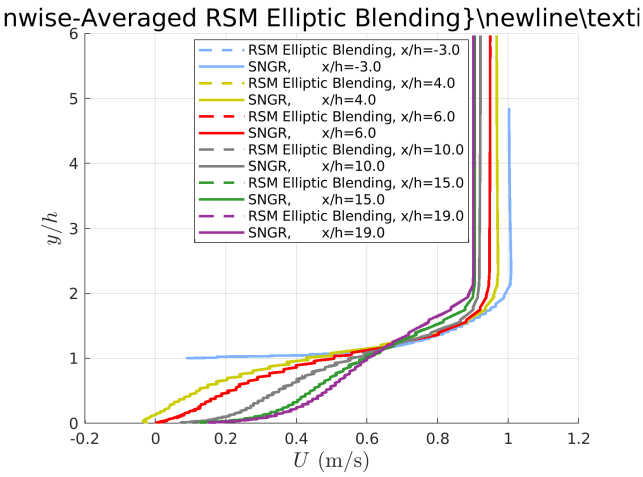

(b) Spanwise-averaged velocity profiles: CFD vs SNGR

Figure 44: (a) Combined SNGR velocity profiles across all slices. (b) Spanwise-averaged velocity profiles comparing CFD (k- SST, dashed lines) and SNGR (solid lines); colors indicate the streamwise slice locations.

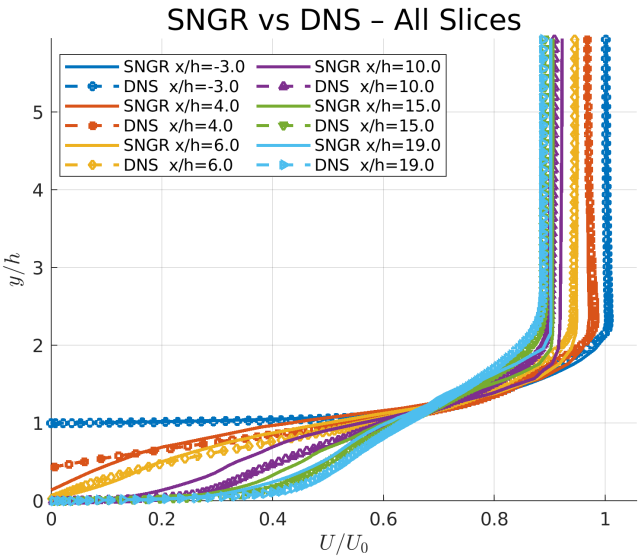

(a) All-slice velocity profile comparison: SNGR vs DNS

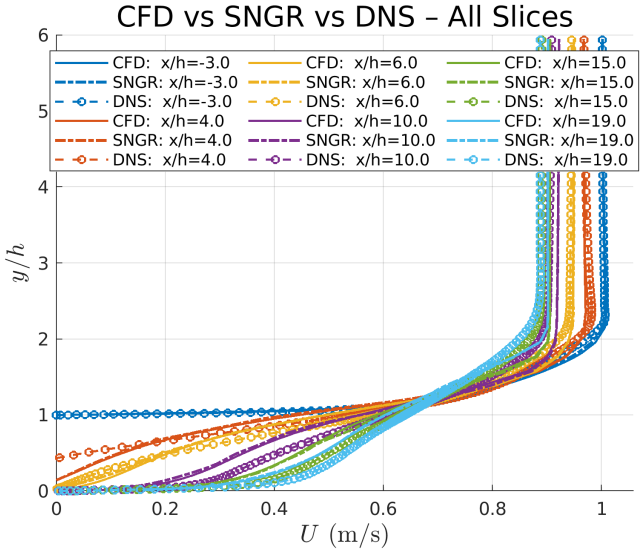

(b) All-slice velocity profile comparison: CFD vs SNGR vs DNS

Figure 45: Comparison of streamwise velocity profiles across all slices for the k- SST case: (a) SNGR vs DNS; (b) CFD (k- SST), SNGR, and DNS combined.

5.3 Velocity profiles in wall units: motivation and relevance

Non-dimensionalizing streamwise velocity profiles using wall units (y^+ and U^+) is a standard diagnostic for near-wall flow behaviour and turbulence modelling. The wall-normal coordinate y and the local mean velocity U are converted to wall units using the friction velocity u_τ and the kinematic viscosity ν, whose values are properly extracted from the RANS simulation run in *StarCCM+*. This is represented mathematically in Equation 71.

$$y^+ \equiv \frac{y\,u_\tau}{\nu}, \qquad U^+ \equiv \frac{U}{u_\tau}, \tag{71}$$

where $u_\tau = \sqrt{\tau_w/\rho}$ and τ_w is the local wall shear stress. These scaled variables collapse the near-wall region onto universal trends for canonical wall-bounded turbulent flows (viscous sublayer, buffer region,and logarithmic layer) and therefore make differences between methods immediately visible. Schlichting and Gersten [2000], Pope [2000]

Wall-unit plots provide a direct test of the near-wall shear representation. In the viscous sublayer the velocity follows $U^+ \approx y^+$ while in the logarithmic region it follows the classical law of the wall, as expressed in Equation 72.

$$U^+ = \frac{1}{\kappa}\ln(y^+) + B, \tag{72}$$

with κ the von Kármán constant and B an additive intercept (typical values $\kappa \approx 0.39$–0.41, $B \sim 4$–6 depending on flow and Reynolds number). Deviations from these trends reveal whether the wall shear is mispredicted (it might be caused due to coarse grid, wall-function usage, or RANS model limitations) and whether a reconstructed field (SNGR) preserves the physically required near-wall scaling inherited from the parent CFD data or DNS reference. Monin and Yaglom [1971], Marusic et al. [2010]

Besides, using wall units highlights model and resolution biases that are otherwise masked in dimensional plots. Two flows with different free-stream velocities, viscosities, or geometry lengths can exhibit very different dimensional profiles while collapsing in wall units if the near-wall dynamics are similar. Thus, this collapse isolates errors due to turbulence modelling or synthetic reconstruction (SNGR) from trivial dimensional differences, making it an indispensable tool when comparing DNS (which, as mentioned in Section 1.1, resolves the near-wall region), RANS (relying on wall functions or Reynolds-stress closures), and reduced-order/synthetic approaches such as SNGR. In separated flows as seen in this main study, shear and stress distributions change rapidly, and the wall-unit representation is particularly sensitive to incorrect wall stress or delayed reattachment. Marusic et al. [2010], Smits et al. [2011]

Moreover, wall-unit analysis provides practical guidance for discretization and validation. If CFD and SNGR profiles match DNS in wall units near the wall, chances that the wall shear and near-wall turbulence intensities are correctly represented increase, being this a relevant insight since many integral and local quantities (such as skin friction, wall-pressure fluctuations, and production of turbulent kinetic energy) depend nonlinearly on the near-wall shear. Conversely, systematic offsets in U^+ point to either insufficient grid resolution, mis-estimated τ_w, or intrinsic model bias (e.g. RANS closure limitations), all of which should be addressed before relying on the flow for acoustics or sensitivity studies.

Furthermore, there are direct implications for aeroacoustics and pressure-fluctuation metrics. Wall-pressure and near-wall velocity fluctuations are primary contributors to surface-based acoustic sources; errors in u_τ or in the shape of $U^+(y^+)$ propagate into incorrect estimates of Reynolds stresses and turbulent kinetic energy near the wall, and hence into the predicted acoustic source amplitude and spectral content. Evaluating profiles in wall units therefore links flow validation to subsequent noise-prediction fidelity and should be reported alongside dimensional profiles for a complete assessment.

For these reasons, presenting $U(y)$ both in dimensional form and in wall units $U^+(y^+)$ gives complementary information: dimensional plots show overall momentum recovery and wake behaviour, while wall-unit plots diagnose near-wall shear, model performance, and fidelity of synthetic turbulence reconstruction. It can be clarified that there two plotting conventions for wall-unit profiles are common in the literature. The standard orientation plots U^+ (vertical) against y^+ (horizontal); this convention is adopted throughout the main text because it facilitates visual verification of the viscous-sublayer slope ($U^+ = y^+$) and the logarithmic law ($U^+ = \frac{1}{\kappa}\ln(y^+) + B$). An inverted orientation (y^+ vertical, U^+ horizontal) could occasionally be used to emphasize the vertical extent of the inner layer. To ensure comprehensive documentation, the SNGR post-processing pipeline has been designed to automatically generate and archive plots in both conventions.

5.3.1 Velocity profiles in wall units: CFD vs. DNS

Before comparing the SNGR results, it is essential to establish the baseline performance of the underlying RANS simulation. Figure 46 compares the CFD (RSM Elliptic Blending) velocity profiles against the DNS reference data in wall units. This comparison quantifies the inherent accuracy of the input data that is fed into the SNGR pipeline.

Results for the Elliptic Blending Reynolds Stress Model

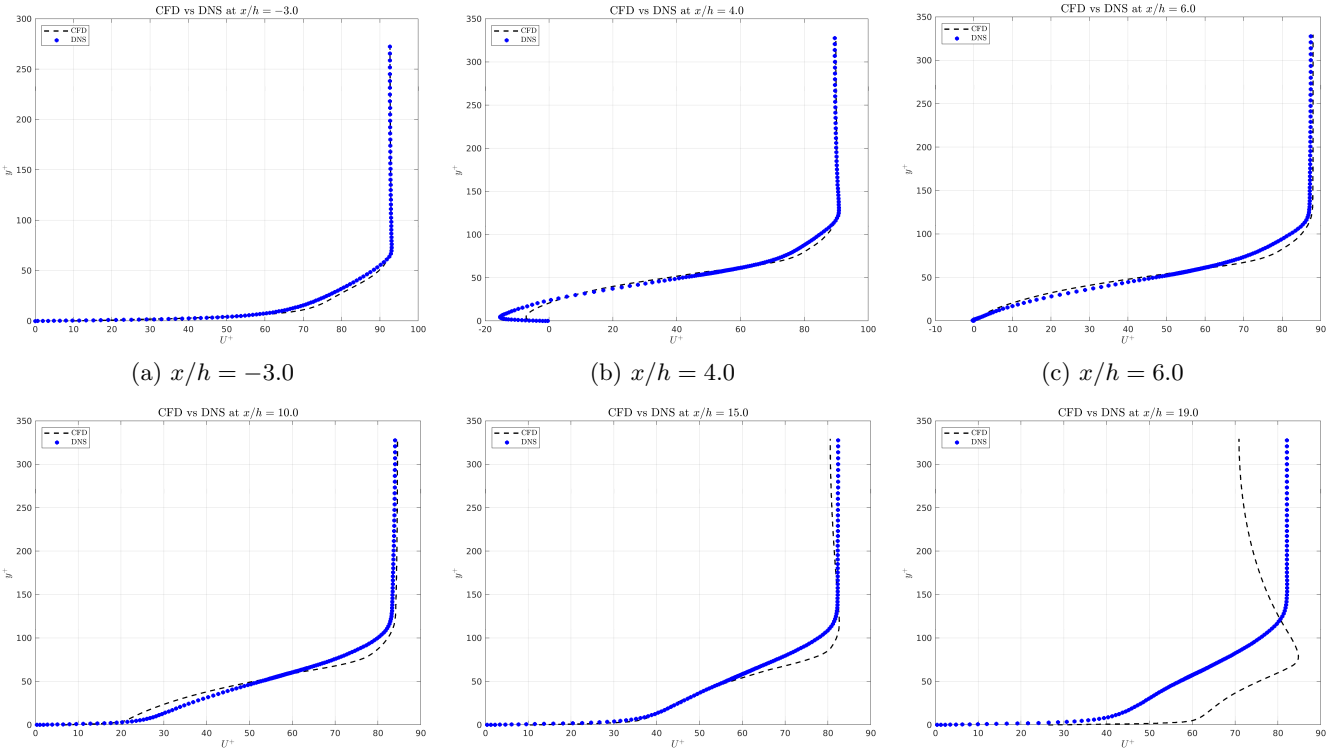

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 46: Streamwise velocity profiles in wall units: CFD (RSM EB) compared to DNS at six streamwise locations.

It can be observed that the RSM-EB model predicts pretty well the velocity profiles up to $x/h = 10.0$, as represented in Figure 46. More in detail, the streamwise velocity profile computed upstream of the step by means of Star-CCM+ matches very well the values estimated by DNS, as seen in Figure 46a. A recirculation region is clearly appreciated in Figure 46b, even though DNS estimates a higher intensity there. Additionally, the free-stream velocity of both methods overlaps in value until $x/h = 15.0$, with a clear gap due to the smoother transition expected according to DNS. Finally, Figure 46f shows a clear mismatch between CFD and DNS: the RSM–EB model overestimates the streamwise velocity below the step and underpredicts it in the region above the step.

With respect to the RANS k-omega SST model, Figure 47 shows an overall better agreement between CFD and DNS. The recirculation zone is described with overlapping lines in Figure 47b, even though CFD overpredicts this effect a bit further downstream the step, at $x/h = 6.0$, as seen in Figure 47c. Regarding the flow from $x/h = 10.0$, both CFD and DNS shows fully developed profiles, with very similar values at all slices, with a clear improvement at $x/h = 19.0$ compared to the representation made by RSM-EB, as seen in Figure 47f.

Results for the RANS k-omega SST

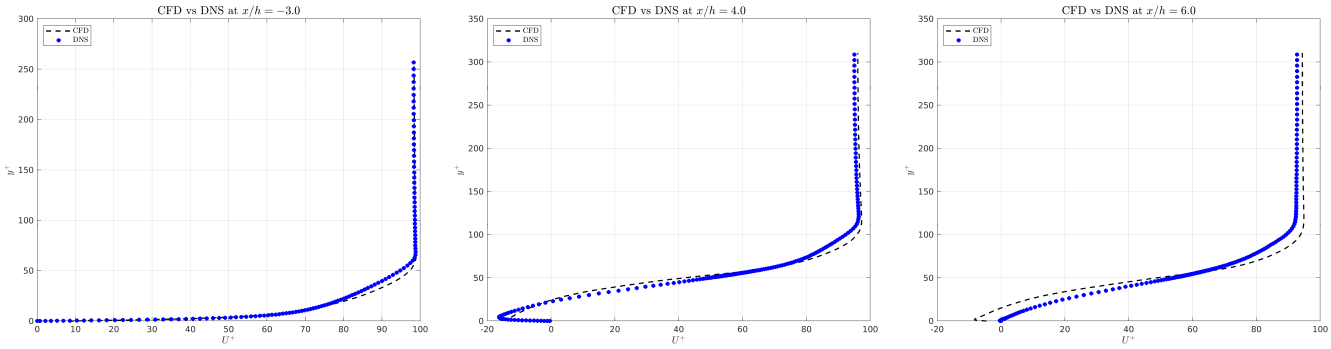

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

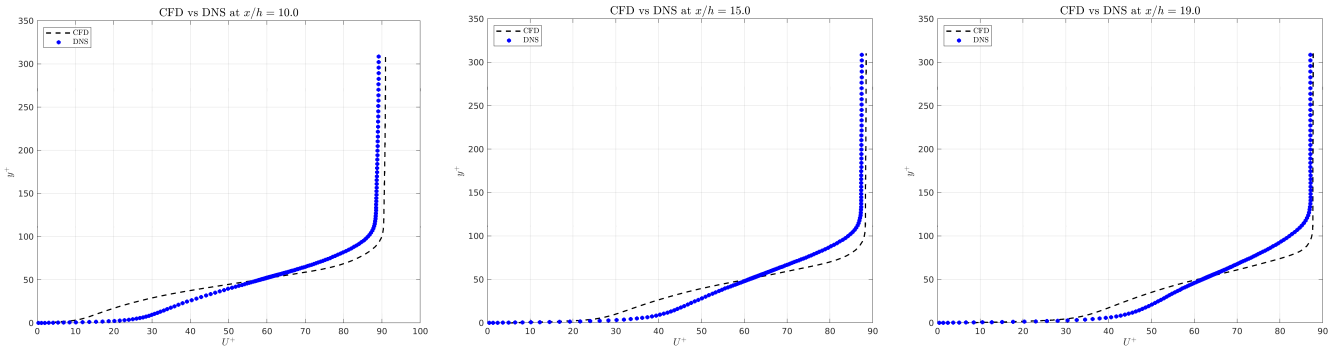

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 47: Streamwise velocity profiles in wall units: CFD (RANS k-omega SST) compared to DNS at six streamwise locations.

5.3.2 Velocity profiles in wall units: SNGR ensemble-averaged vs. CFD

Similarly, Figure 48 compares the ensemble-averaged SNGR velocity profile to the CFD input. In theory, if the ensemble average of the velocity fluctuations $\langle u' \rangle_{ens}$ is zero, this profile should be identical to the CFD profile. Any small deviations observed here are due to residual mean fluctuations from the finite number of ensemble samples, serving as a check on statistical convergence. The number of independent realizations considered for these SNGR simulations in MATLAB was 150.

Results for the Elliptic Blending Reynolds Stress Model

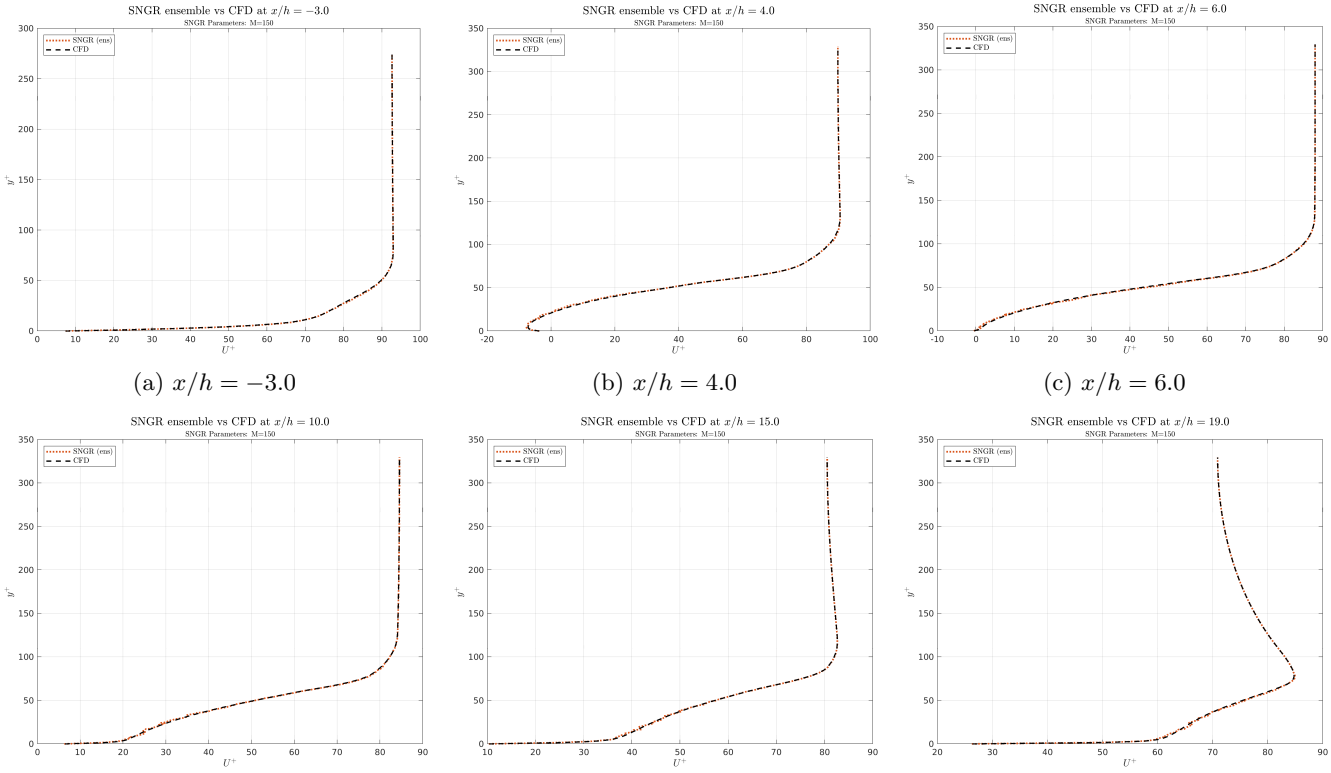

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 48: Streamwise velocity profiles in wall units: SNGR (ensemble-averaged) compared to its CFD input (RSM-EB) at six streamwise locations.

Figure 48 shows that the ensemble average reproduces exactly the streamwise velocity values as imported from CFD, indicating that the number of independent realizations is enough to successfully prove that these values are properly used as input data for the SNGR methodology.

Due to the nature of this ensemble-averaged SNGR approach, any artifacts present in the flow according to the CFD simulation are also represented using this SNGR estimation, with a clear underprediction of the free-stream velocity at $x/h = 19.0$, as indicated in Figure 48f.

Concerning the RANS k-omega SST model, all six plots described in Figure 49 show overlapping lines, as CFD data is properly used as the input values for the ensemble-averaged SNGR approach. Thus, this turbulence model represents - as said before - a clearer recirculation region, and smoother flow near the outlet, as shown in Figure 49b and Figure 49f, respectively.

Results for the RANS k-omega SST

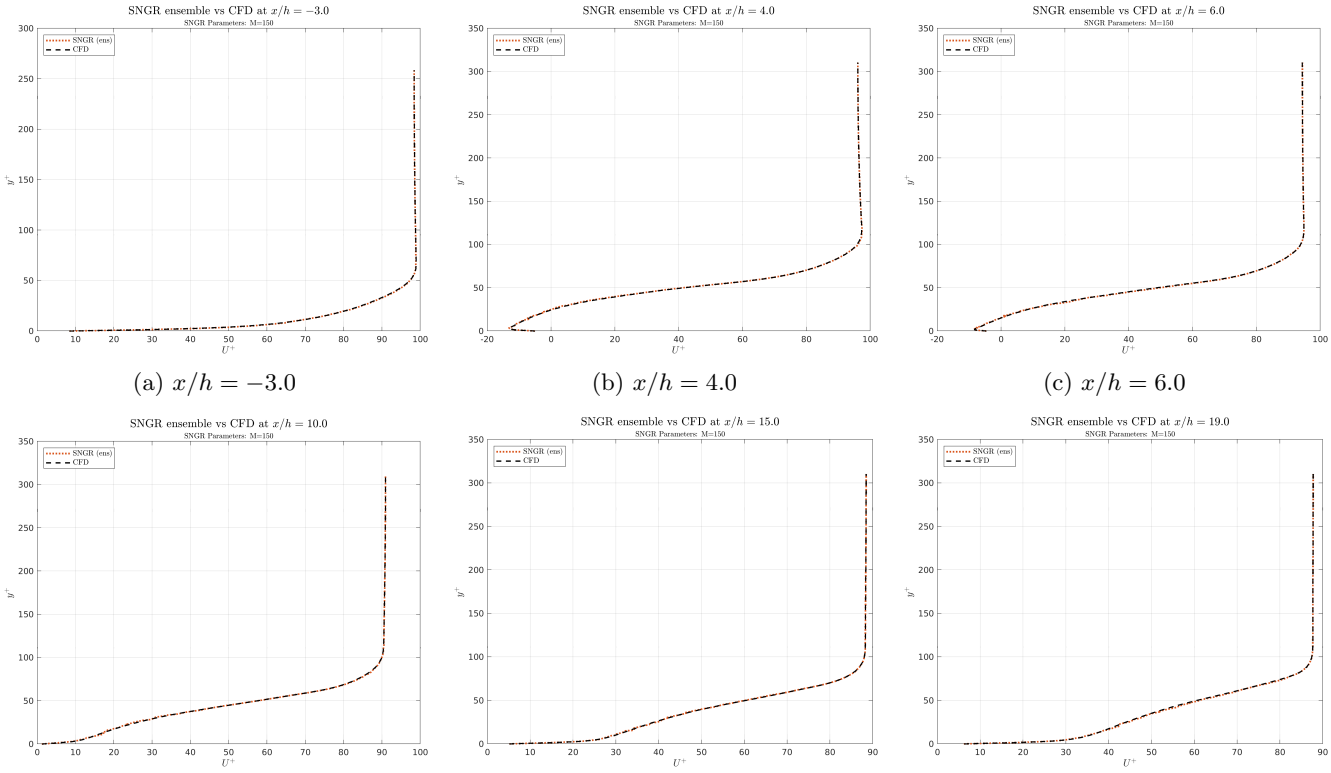

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 49: Streamwise velocity profiles in wall units: SNGR (ensemble-averaged) compared to its CFD input (RANS k-omega SST) at six streamwise locations.

5.3.3 Velocity profiles in wall units: SNGR time-averaged vs. CFD

This section validates the fidelity of the SNGR time-marching reconstruction against its input data. Figure 50 shows the comparison between the time-averaged SNGR velocity profile and the original CFD profile. A close match indicates that the SNGR process, including the addition of time-averaged fluctuations $\langle u' \rangle$, does not significantly distort the mean flow field provided by the RANS simulation. The number of iterations considered for the SNGR simulations run in MATLAB was 150, the same number that was used for the independent realizations from the ensemble-averaged approach.

Results for the Elliptic Blending Reynolds Stress Model

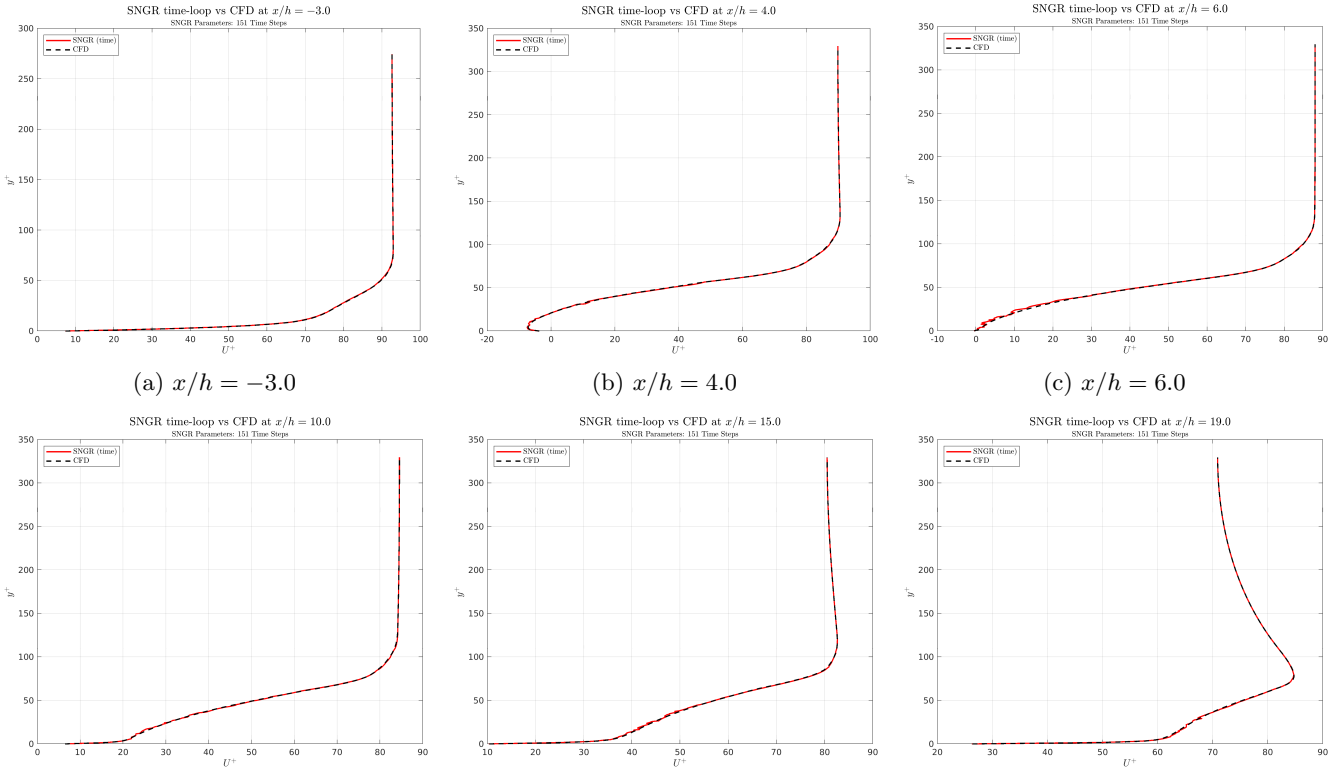

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 50: Streamwise velocity profiles in wall units: SNGR (time-marching) compared to its CFD input (RSM-EB)at six streamwise locations.

Figure 50 shows that the SNGR time-loop approach also describes extremely similar velocity profiles in wall-units compared to the original CFD data - considering the RSM-EB as the turbulence model - used as the input for the very first iteration of the simulation. Compared to the ensemble-averaged approach, the SNGR profiles are slightly less smooth, indicating that the length of the simulation must be increased to obtain better results. Nonetheless, the overall results are satisfying.

The same explanation can be told with reference to the RANS k-omega SST values, shown in Figure 51. Once again, the recirculation zones are properly recreated by the single-time-loop SNGR approach, taking into account that it is matching CFD values that already have some innacuracies compared to DNS, as observed in Section 5.3.1.

Results for the RANS k-omega SST

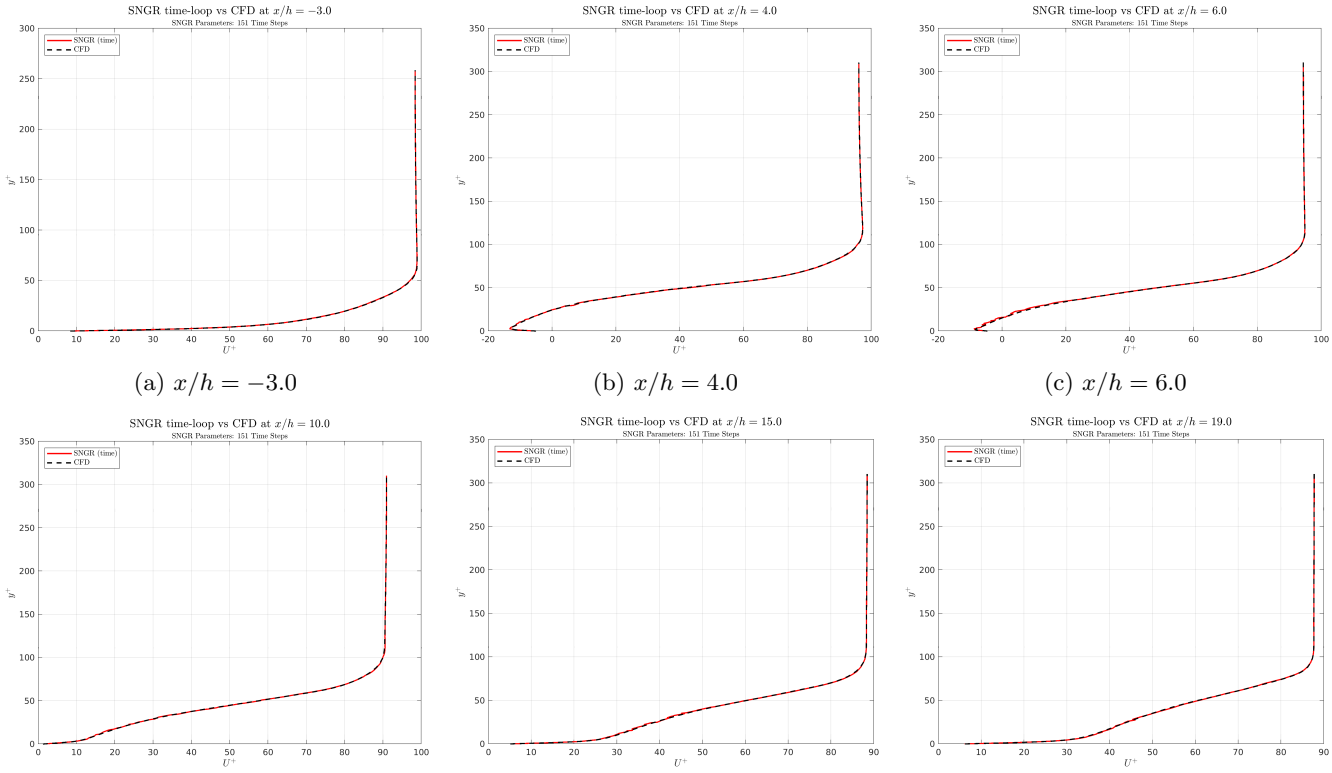

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 51: Streamwise velocity profiles in wall units: SNGR (time-marching) compared to its CFD input (RANS k-omega SST) at six streamwise locations.

5.3.4 Velocity profiles in wall units: SNGR time-averaged vs. SNGR ensemble-averaged

Having first established the baseline agreement between CFD and DNS (described in Section 5.3.1) and verified that SNGR reproduces its CFD input (as seen in Section 5.3.2 and in Section 5.3.3), the purpose of plotting the SNGR time-averaged and ensemble-averaged profiles in wall units is to isolate near-wall sampling and reconstruction effects. Wall units collapse the inner layer where shear production, wall stress and small-scale structure control both force- and noise-producing mechanisms. A comparison of $SNGR_{\text{time}}$ and $SNGR_{\text{ens}}$ in $U^+(y^+)$ therefore checks that the synthesized fluctuations preserve the RANS u_τ scaling and do not bias the mean, reveals whether observed discrepancies are due to finite-time sampling (time-loop length) or to systematic deficiencies in the reconstruction (anisotropy, interpolation or divergence errors).

Results for the Elliptic Blending Reynolds Stress Model

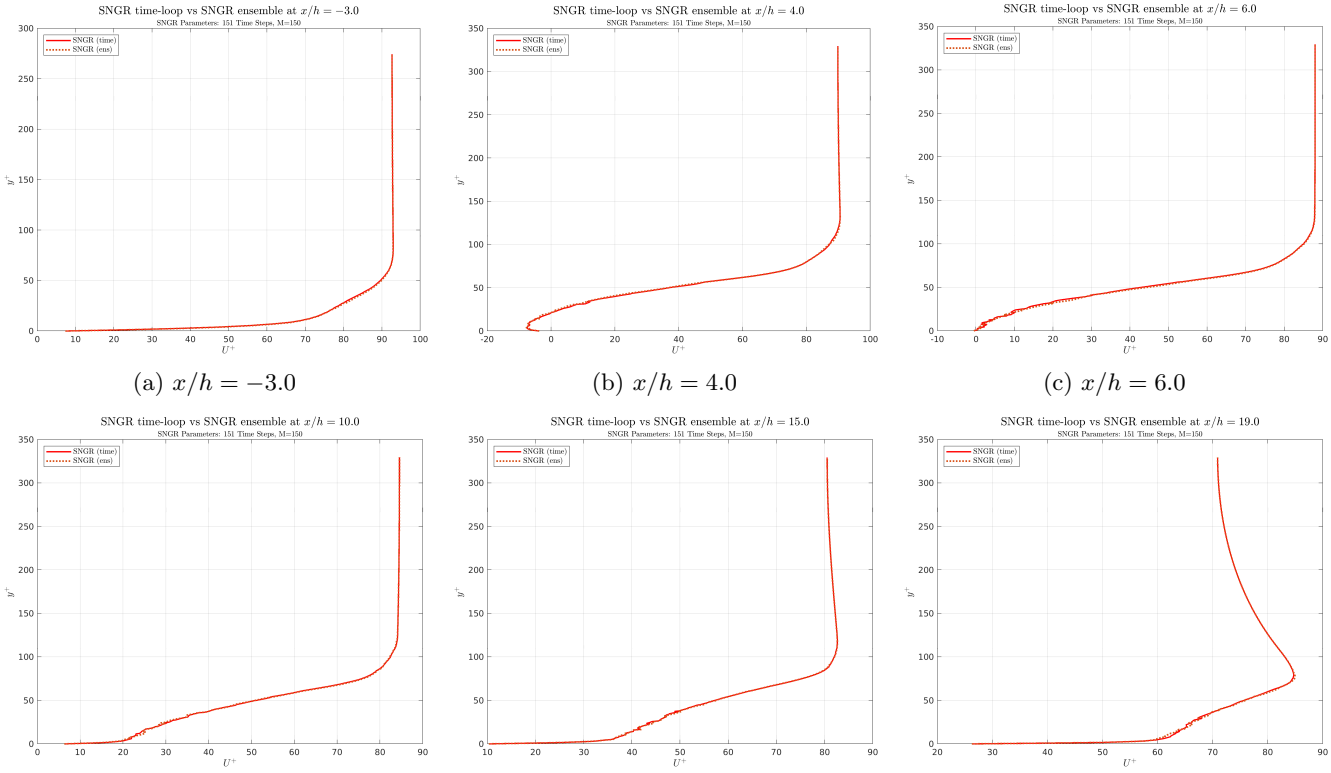

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 52: Streamwise velocity profiles in wall units: both SNGR methods (time-marching and ensemble average) are compared at six streamwise locations.

Figure 52 shows the direct comparison between the two SNGR variants only — the single time-loop and the ensemble-averaged results. Since no DNS curves are included in these plots, the purpose here is not to assess physical accuracy but to verify that both SNGR approaches produce consistent mean profiles under the same input conditions. For the RSM–Elliptic Blending case, the two curves overlap almost perfectly across all y^+ values, indicating that the time-loop simulation has reached statistical convergence for the chosen simulation length.

The same behaviour is observed for the RANS k–ω SST input (Fig. 53). Because the turbulence model only changes the CFD input field—not the SNGR algorithm—the agreement between $SNGR_{\text{time}}$ and $SNGR_{\text{ens}}$ reflects the convergence properties of SNGR itself.

Results for the RANS k-omega SST

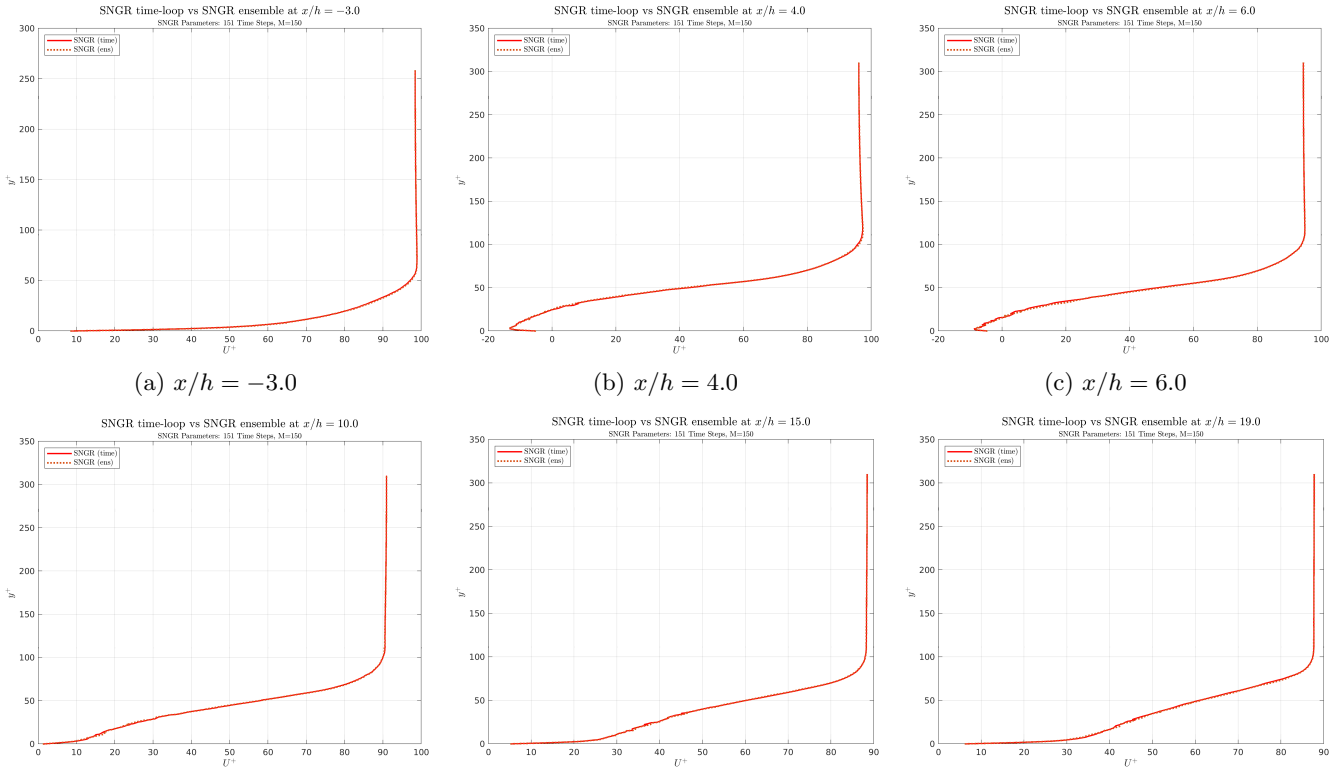

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 53: Streamwise velocity profiles in wall units: both SNGR methods (time-marching and ensemble average) are compared at six streamwise locations.

5.3.5 Velocity profiles in wall units: SNGR time-loop, SNGR ensemble, CFD and DNS

Results for the Elliptic Blending Reynolds Stress Model

Figure 54 show $U^+(y^+)$ at the same streamwise stations as seen before. Introducing CFD in the wall-unit comparison makes it possible to distinguish which differences are inherited from the RANS input (CFD vs DNS) and which are introduced or amplified by the synthetic reconstruction (SNGR vs CFD/DNS).

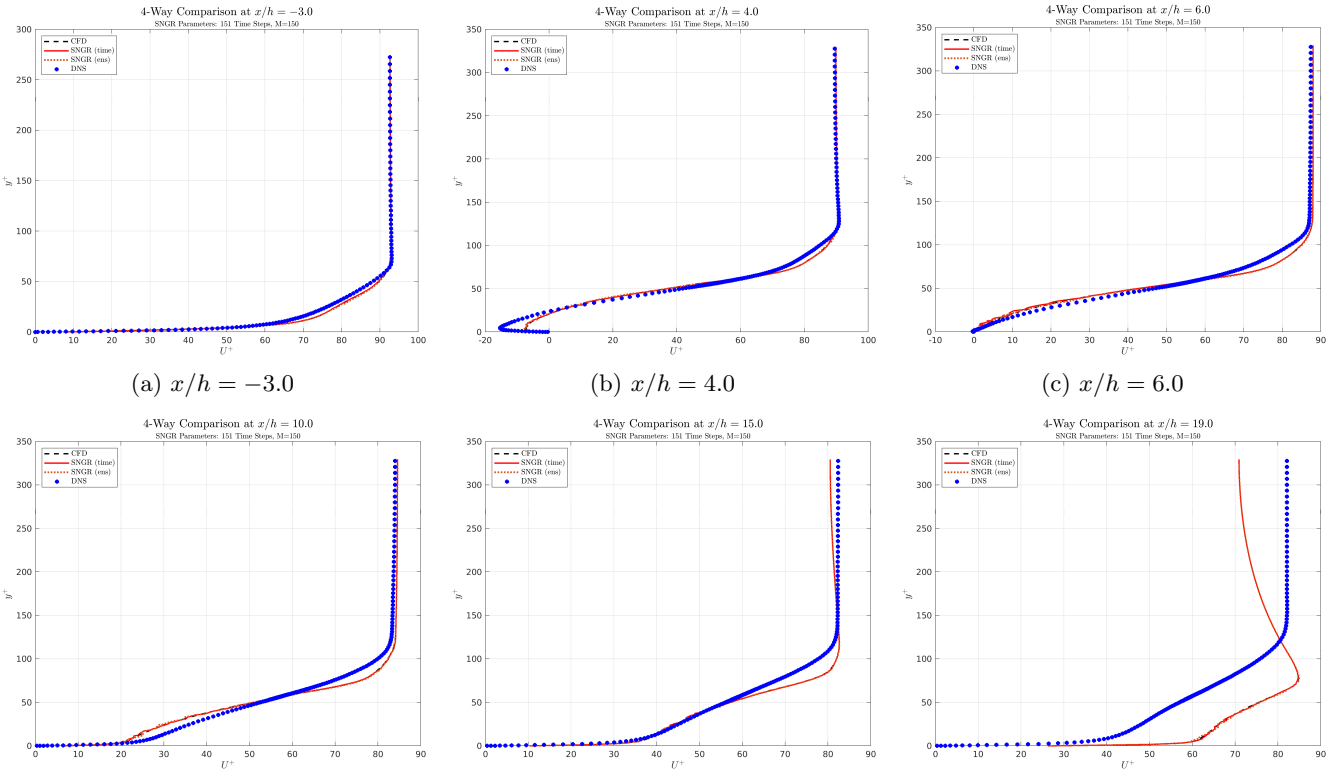

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 54: Streamwise velocity profiles in wall units: SNGR (time-marching), SNGR (ensemble), CFD (RSM EB) and DNS compared at six streamwise locations.

Upstream of the step - at $x/h = -3.0$ - all four curves collapse very closely in wall units across the viscous sublayer and buffer region, as shown in Figure 54a. This strong agreement indicates that the inlet/wall-stress estimate used in CFD is consistent with DNS, and that SNGR faithfully reproduces the near-wall scaling when the parent CFD statistics are accurate. Practically, this establishes a baseline: discrepancies observed downstream are unlikely to originate from a mismatched normalization u_τ at the inlet.

Immediately downstream of the step DNS shows a bigger recirculation region compared to CFD and SNGR methods near the wall, while still matching fairly well in the outer region, as presented in Figure 54b. DNS shows the strongest overshoot in the inner and the most extreme negative near-wall values, indicating the largest reversed flow magnitude. CFD follows DNS reasonably well, and SNGR follows CFD closely along the entire profile.

At $x/h = 6.0$, DNS presents a slightly sharper transition from near-wall region to outer free-stream region, whereas CFD and SNGR match closely. The four-way comparison clarifies that most of the large-scale differences with DNS are already present in CFD, inheriting the SNGR method those differences and adding a small extra smoothing due to finite modal content, as represented in Figure 54c.

Further downstream, at $x/h = 10.0$ and $x/h = 15.0$, there is no longer any recirculation and the four methods converge strongly: CFD and both SNGR approaches obtain very similar values compared to DNS in the viscous sublayer, buffer and logarithmic overlap, indicating that the boundary layer is re-developing and that modelling errors are diminishing. The overall close alignment at these stations is reassuring for aeroacoustic purposes because correct near-wall scaling directly affects estimates of near-surface pressure fluctuations and TKE production.

At last, the farthest downstream slice shows the largest remaining spread: DNS exhibits a clearly different outer-layer shape (and possibly a different log-intercept), while CFD and SNGR methods have overlapping values, as observed for all previous comparisons in wall-units, but both deviate modestly from DNS in an intermediate y^+ band. Because CFD and SNGR are matching here, the discrepancy is most likely dominated by RANS model bias or by small errors in the computed wall shear τ_w used for normalization, since SNGR does not introduce a large additional error beyond the CFD input at this station.

Overall, including CFD in the wall-unit comparison separates inherited model bias (CFD vs DNS) from reconstruction error (SNGR vs CFD). In Figure 54, it is clearly observed that where CFD matches DNS, SNGR also matches (SNGR faithfully inherits accurate RANS statistics), and where CFD is different than DNS, SNGR largely reproduces the same departure and may add at best a small additional smoothing.

Results for the RANS k-omega SST

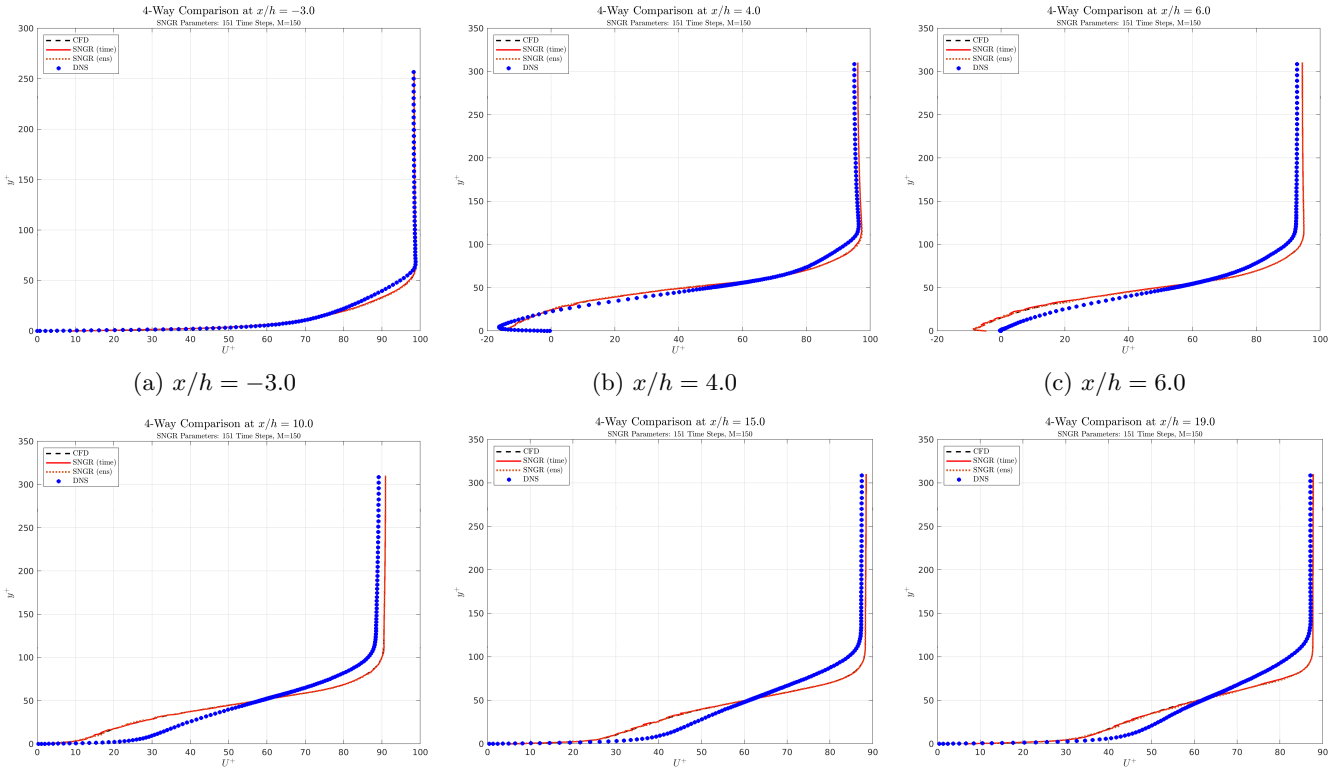

(a) $x/h = -3.0$ (b) $x/h = 4.0$ (c) $x/h = 6.0$

(d) $x/h = 10.0$ (e) $x/h = 15.0$ (f) $x/h = 19.0$

Figure 55: Streamwise velocity profiles in wall units: SNGR (time-marching), SNGR (ensemble), CFD (k–ω SST) and DNS compared at six streamwise locations.

A similar analysis may be described for the RANS k-omega SST comparison, evaluating four methods and showing the results in Figure 55. As thoroughly explained in previous subsections, SNGR single time-loop and ensemble averaged approaches match CFD values, which is in general in good agreement with DNS values.

5.4 Turbulent Kinetic Energy along the streamwise midplane and spanwise-averaged

The evaluation of Turbulent Kinetic Energy (from now on, TKE) is crucial in the SNGR method as TKE directly influences the intensity and spectral characteristics of the synthetic turbulence, which in turn affects the acoustic field generated by turbulent flows. TKE measures the kinetic energy per unit mass contained in the velocity fluctuations, quantifying the energy contained in turbulent velocity fluctuations. It is defined as shown in Equation 73:

$$k \;=\; \tfrac{1}{2}\langle u_i' u_i' \rangle \;=\; \tfrac{1}{2}\big(\langle u'u' \rangle + \langle v'v' \rangle + \langle w'w' \rangle\big) \tag{73}$$

where u', v', w' are the fluctuating velocity components and angle brackets denote an appropriate average (time or ensemble).

In SNGR, this energy is used to generate synthetic velocity fluctuations that mimic real turbulence, being TKE a primary amplitude scale: the synthetic velocity fluctuations are scaled so that their component variances are consistent with the desired total fluctuating energy. For an isotropic assumption each component variance would be:

$$\langle u'u' \rangle \;=\; \langle v'v' \rangle \;=\; \langle w'w' \rangle \;=\; \tfrac{2}{3}k \tag{74}$$

so the root-mean-square of any component is $u'_{\text{rms}} = \sqrt{2k/3}$. Thus an error in k directly changes the RMS fluctuation amplitude of the synthetic field. Because acoustic source amplitudes depend on the fluctuation amplitude, if TKE is inaccurate, the synthetic turbulence will not correctly represent the physical flow, leading to incorrect noise predictions.

Controls Noise Source Amplitude Higher TKE means stronger fluctuations, which typically produce more intense sound radiation.

The evaluation of Turbulent Kinetic Energy (TKE) is crucial for assessing the overall energy content of the synthetic turbulence generated by the SNGR method. TKE quantifies the mean kinetic energy per unit mass associated with the velocity fluctuations and is defined as:

$$k = \frac{1}{2}\left(\langle u'u' \rangle + \langle v'v' \rangle + \langle w'w' \rangle\right) \tag{75}$$

In the SNGR method, TKE serves as the primary scaling parameter for the amplitude of the generated fluctuations. While TKE sets the overall energy level, its ability to create a physically realistic field depends on its interplay with other key parameters, as discussed in detail in Section 3.2. This section now focuses on evaluating the spatial distribution of the TKE reconstructed by the SNGR method, comparing it against the input CFD data under various mesh and inlet conditions.

SNGR aims to simulate realistic turbulent noise sources without resolving all turbulent scales directly. TKE provides a physically grounded way to inject energy into the synthetic field, ensuring that the generated noise is consistent with the underlying flow physics, being physically realistic.

Note on Spanwise Averaging While a single spanwise-averaged streamwise velocity profile was previously shown in Figure 42b to provide an integrated comparison between SNGR and CFD, this section marks the first systematic use of spanwise averaging in the analysis. The motivation for applying spanwise averaging to the turbulent kinetic energy (TKE) field lies in its ability to reduce statistical noise and highlight large-scale trends in turbulence intensity. Unlike velocity profiles, which benefit from preserving local three-dimensional features for detailed comparisons with DNS and CFD at specific slices, TKE is a scalar quantity that is more meaningfully interpreted in an averaged form. This approach facilitates clearer visualization of the spatial distribution of turbulent energy and supports robust evaluation of the SNGR method under varying mesh and inlet conditions.

5.4.1 Considering the Reynolds Stress Model Elliptic Blending - Coarse Mesh

This section evaluates the TKE distribution on a coarse mesh configuration, using the same initial and boundary conditions that have been described in previous sections. It is important to note that this mesh resolution differs from the refined mesh used in the primary validation sections (Section 5.1 to 5.3). The purpose of this comparison is to evaluate the SNGR method's performance and robustness on grids with lower spatial resolution, representing SNGR-based reconstructions of TKE along the streamwise mid-plane and spanwise-averaged distribution, alongside a reference prediction from STAR-CCM+ using the Reynolds Stress Model with Elliptic Blending.

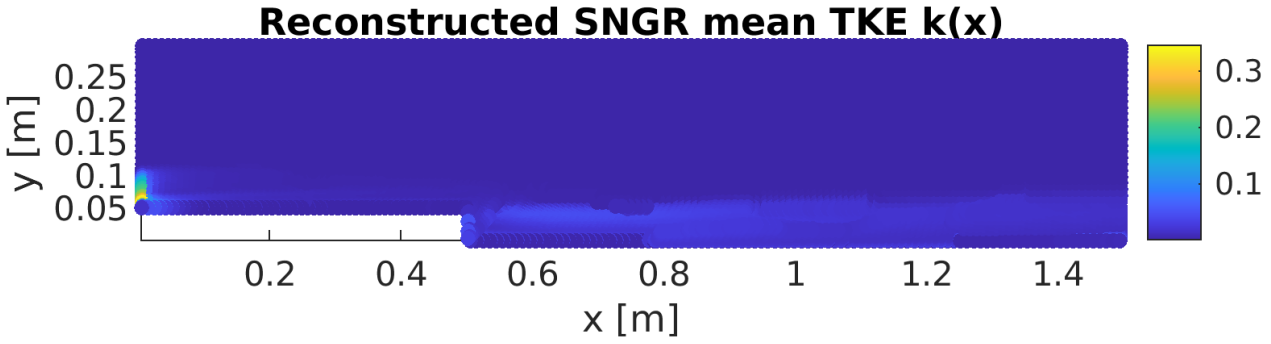

(a) TKE values along the streamwise mid-plane of the coarse BFS model

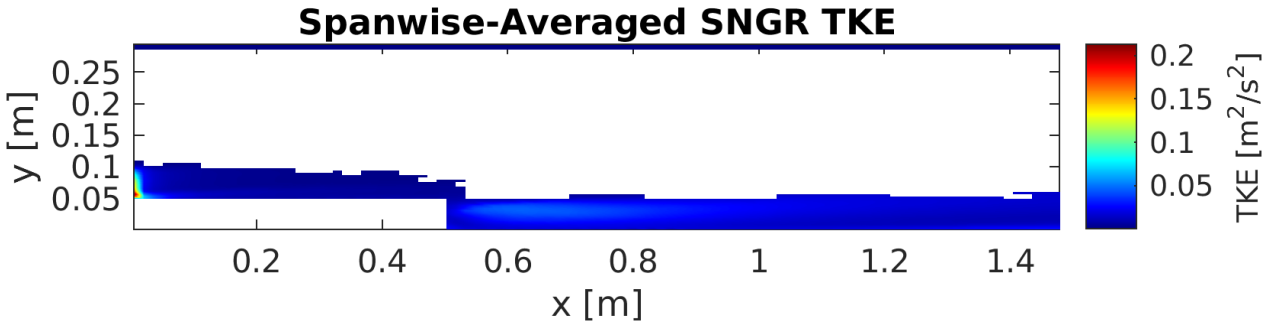

(b) TKE values considering the spanwise average of the coarse BFS model

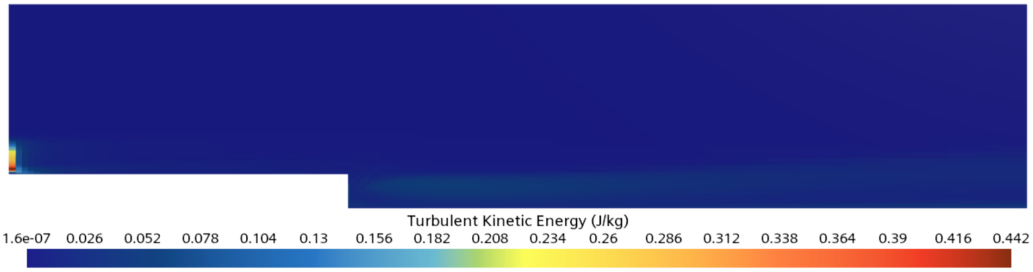

(c) TKE values predicted in StarCCM+ considering the RSM Elliptic Blending as the turbulence model

Figure 56: Estimation of the turbulent kinetic energy (TKE) using the single time-loop approach from the SNGR method: (a) distribution on the streamwise midplane, (b) spanwise-averaged distribution (also averaged along the streamwise direction, where applicable), and (c) reference from the CFD simulation. All results correspond to a coarse mesh and the RSM Elliptic Blending turbulence model.

The mean turbulent kinetic energy (TKE) reconstructed using the SNGR method along the streamwise mid-plane reveals a localized distribution and sharp peak near the step, with values reaching around 0.35, as observed in Figure 56a. The turbulence appears more concentrated and less diffused, which is typical of coarse grids that may under-resolve smaller turbulent structures but still capture the dominant flow features.

The spanwise-averaged TKE from the SNGR reconstruction corresponding to this case is described in Figure 56b. The averaging smooths out the localized peaks and provides a bit clearer and more accurate view of the overall turbulence distribution. The peak values are about 40% lower than in the mid-plane view, and the energy appears more confined to the near-wall region downstream of the step. This suggests that the coarse mesh, while sufficient for capturing large-scale turbulence, may limit the spread of turbulent energy across the span.

The TKE distribution provided by StarCCM+, represented in Figure 56c, aligns well with the SNGR reconstructions, especially in the location and shape of the high-TKE region. However, the peak values from SNGR are lower, indicating that the model may be more dissipative estimating turbulence intensity. Overall, the agreement between

the SNGR method and the CFD prediction supports the validity of both approaches under coarse mesh conditions.

5.4.2 Considering the Reynolds Stress Model Elliptic Blending - Coarse Mesh and reduced turbulent flow at the inlet

To assess the sensitivity of the SNGR method to inlet boundary conditions, this section presents results using a deliberately reduced turbulence profile at the inlet. This case contrasts with the primary validation studies (which utilize the high-fidelity Spalart profile) and serves to demonstrate how the SNGR reconstruction responds to lower-energy input data. This approach is slightly different from the DNS paper, as the actual incoming turbulent boundary layer is set based on a reference paper, but it is useful to compare CFD input data and SNGR.

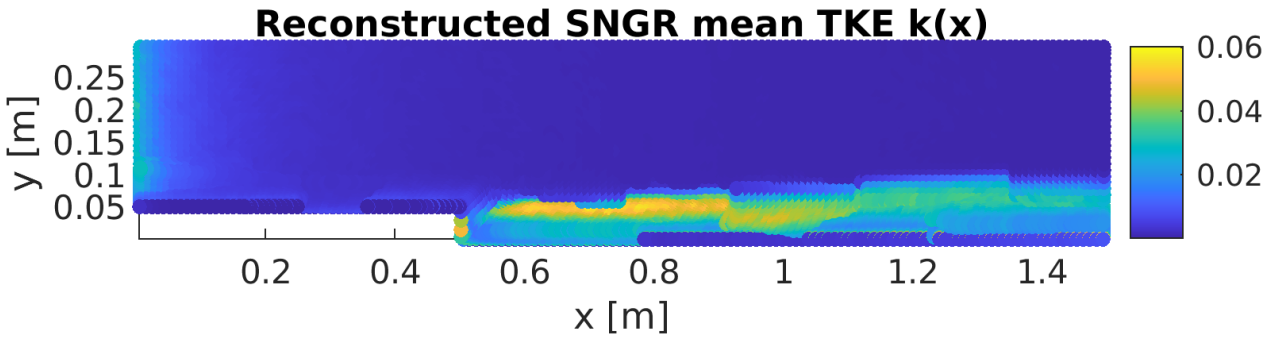

(a) TKE values along the streamwise mid-plane of the coarse BFS model

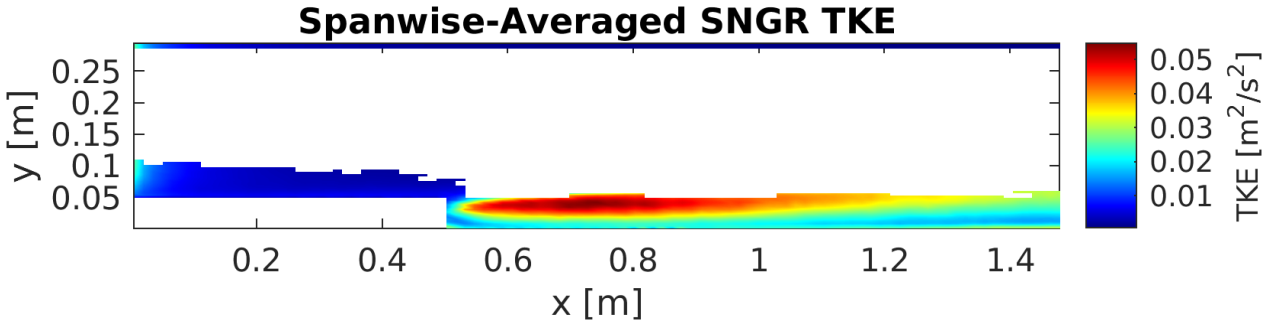

(b) TKE values predicted in StarCCM+ considering the RSM Elliptic Blending as the turbulence model

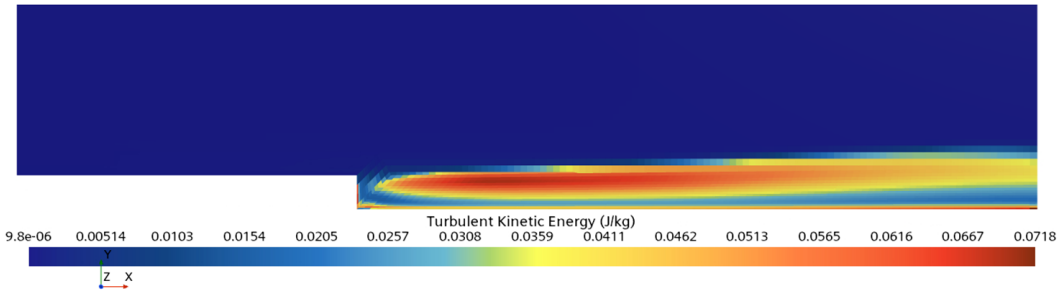

(c) TKE values considering the spanwise average of the coarse BFS model

Figure 57: Estimation of the TKE of the flow estimated using the single time-loop approach from SNGR, considering a refined mesh and RSM Elliptic Blending as the turbulence model

A massive difference can be noticed with respect to values obtained having considered the actual turbulent inlet profile, as maximum TKE values reach 0.06 m^2/s^2 in the recircultion region - approximately from $0.6m < x < 0.8m$, as can be seen in Figure 57. The lower range shown in this case allows for a visually clearer representation of the TKE computed by the SNGR method, both assuming a certain plane in the streamwise direction, as well as calculating the average in the spanwise direction.

The distribution shown in Figure 57a reveals a more gradual buildup of turbulence downstream of the step, with noticeably lower peak values compared to the case with a correctly specified inlet profile, exhibited previously in Figure 56a. This reflects the reduced turbulence generation due to the weaker incoming boundary layer, and highlights the sensitivity of the flow development to inlet conditions.

With respect to the spanwise average, represented in Figure 57b, the predicted distribution is consistent with the SNGR reconstruction, showing low initial TKE and a delayed rise in turbulence intensity. The peak values remain modest (about 0.05 m^2/s^2), and the turbulence appears more confined to the near-wall region, suggesting that the

model captures the subdued turbulence dynamics reasonably well under these conditions, showing an akin profile as the one calculated in StarCCM+, shown in Figure 57c. In that plot, the averaging made by RANS simulation smooths out localized variations and confirms the overall low turbulence levels throughout the domain. Compared to previous cases, the energy distribution is more uniform and lacks strong gradients, reinforcing the impact of the reduced inlet turbulence on the flow field. This visualization is particularly useful for comparing with DNS or experimental data, as it emphasizes the broader flow behavior rather than localized features.

5.4.3 Considering the Reynolds Stress Model Elliptic Blending - Refined Mesh and very low turbulent flow at the inlet

The distribution of the turbulent kinetic energy is shown as well having considered the refined mesh, with about 10 times the amount of cells compared to the previous case, and a lower turbulent profile at the inlet, with the same magnitude as in in Section 5.4.2.

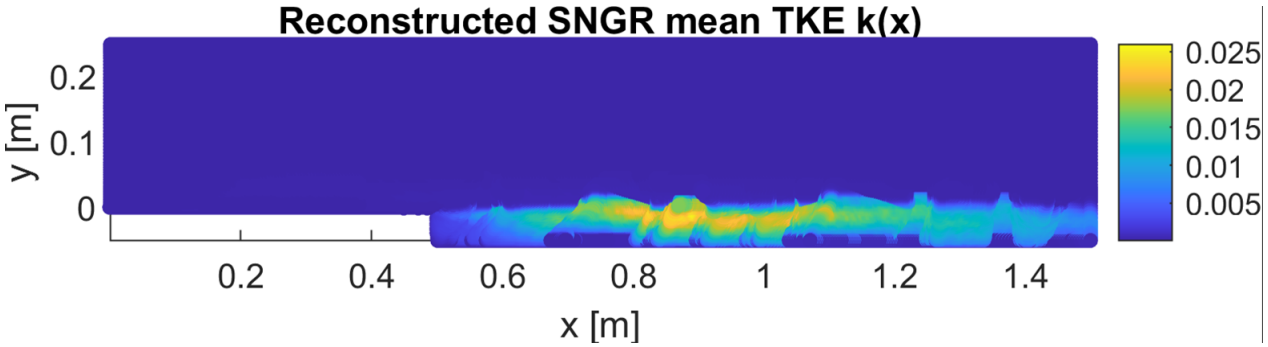

(a) TKE values along the streamwise mid-plane of the refined BFS model

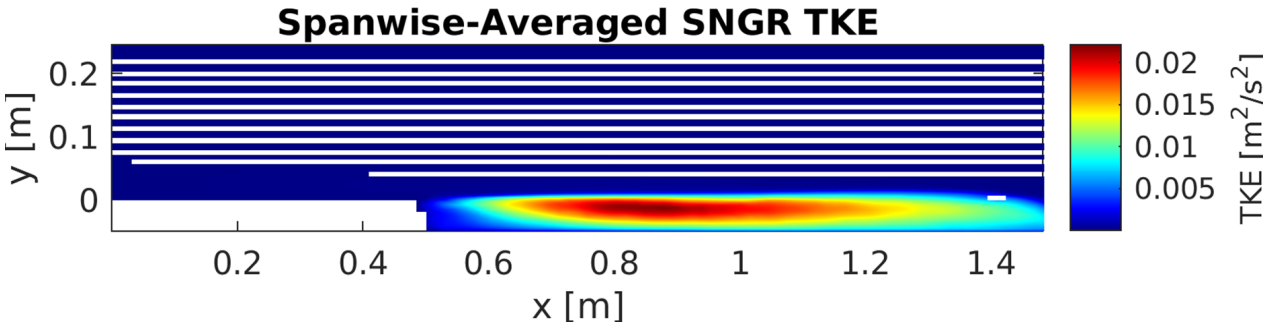

(b) TKE values considering the spanwise average of the refined BFS model

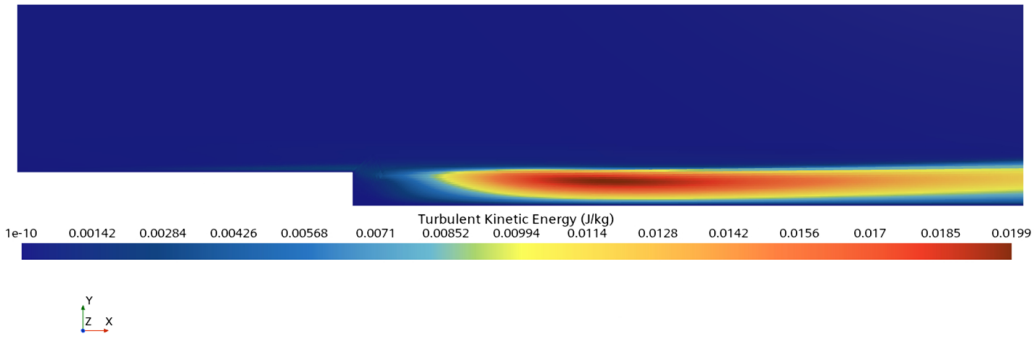

(c) TKE values predicted in StarCCM+ considering the RSM Elliptic Blending as the turbulence model

Figure 58: Estimation of the turbulent kinetic energy (TKE) using the single time-loop approach from the SNGR method: (a) distribution on the streamwise midplane, (b) spanwise-averaged distribution (also averaged along the streamwise direction, where applicable), and (c) reference from the CFD simulation. All results correspond to a refined mesh and the RSM Elliptic Blending turbulence model.

It is important to note the difference in representation between the subfigures. Figure 58a displays the TKE values extracted from the streamwise mid-plane by plotting the discrete cell centroids, which results in a granular, point-based appearance. In contrast, Figure 58b presents the spanwise-averaged TKE, where the data is integrated across the spanwise direction and rendered as a continuous contour map, yielding a smoother visual representation.

Figure 58a shows the mean turbulent kinetic energy (TKE) reconstructed using the SNGR method along the streamwise mid-plane of the backward-facing step geometry. The distribution reveals a clear increase in TKE just downstream of the step, particularly near the lower wall, which is consistent with the formation of a shear layer and the onset of recirculation. The peak values reach around 0.025, about 20% higher than maximum values observed in Figure 58c, even though the location of the peak matches pretty well the SNGR method.

The spanwise-averaged TKE, derived from the SNGR method, is presented in Figure 58b. This averaging smooths out local fluctuations and highlights the broader structure of the turbulent field. The TKE remains concentrated in the same downstream region, with slightly lower peak values of circa $0.02 m^2/s^2$, obtaining a very similar profile than the scene represented in StarCCM+, exhibited in Figure 58c. This view is particularly useful for comparing with RANS-based models or experimental data, as it emphasizes the overall energy distribution rather than localized peaks, proving that the SNGR method is properly computing the TKE read from the *cgns* file extracted from CFD.

Overall, the results computed by means of the SNGR time-loop approach show a similar spatial pattern to the CFD values, with turbulence concentrated in the recirculation zone and gradually decaying downstream. The maximum TKE is slightly lower than in the SNGR results, peaking at $0.0199 J/kg$, but the qualitative agreement is strong. This consistency supports the validity of the turbulence model and the refined mesh strategy used in the simulation.

5.5 Power Spectral Density Spectrum and analysis of SNGR velocity fluctuations

Having generated a time series of turbulent velocity fluctuations using SNGR, the next step is to verify that the synthetic field contains the correct frequency content. A Power Spectral Density (PSD) analysis answers the two key questions:

- Do the fluctuations contain the expected *energy-containing* (large-eddy) frequencies?
- Do they exhibit an *inertial subrange* with the Kolmogorov $-5/3$ slope and a subsequent high-frequency dissipation roll-off?

5.5.1 Definition & link to spatial spectra

Let $u'(t)$ be a velocity fluctuation recorded at a fixed probe location. The (one-sided) PSD is defined as the Fourier magnitude square of the (finite-time) signal and, in the continuous limit, as described in Equation 76.

$$S_u(f) \;=\; \lim_{T\to\infty} \frac{1}{T} \left| \int_0^T u'(t)\, e^{-i2\pi f t}\, dt \right|^2 . \tag{76}$$

In practice, $S_u(f)$ is estimated using with Welch's method (overlapped windowed segments), which reduces variance at the cost of frequency resolution.

Under Taylor's frozen turbulence hypothesis (valid when turbulent structures are convected past the probe with a convective speed U_c that is large compared with local fluctuation speeds), temporal frequency f and spatial wavenumber k are related by

$$f \;=\; \frac{U_c\, k}{2\pi}. \tag{77}$$

Consequently, a wavenumber spectrum $E(k) \propto k^{-5/3}$ implies a frequency spectrum

$$S_u(f) \;\propto\; f^{-5/3}, \tag{78}$$

so the Kolmogorov slope is preserved under the Taylor mapping provided U_c is chosen appropriately Kolmogorov [1941], Pope [2000].

5.5.2 Characteristic scales and expected frequency bounds

Three characteristic frequencies are used to compute and evaluate the frequency spectrum of the flow:

- **Integral (large-eddy) scale** L. A characteristic low frequency is

$$f_L \approx \frac{U_c}{L}. \tag{79}$$

  Frequencies $f \lesssim f_L$ belong to the energy-containing range; spectra here are typically flat or gently sloped.

- **Kolmogorov (dissipation) scale** $\eta = (\nu^3/\varepsilon)^{1/4}$. The corresponding dissipation frequency is

$$f_\eta \approx \frac{U_c}{2\pi\eta}. \tag{80}$$

  Frequencies $f \gtrsim f_\eta$ lie in the viscous dissipation range where the spectrum steepens.

- **Inertial subrange** is expected between f_L and f_η; the width of this range determines how clearly a $-5/3$ slope can be observed.

If ε is not directly available, a practical estimate is $\varepsilon \sim U'^3/L$ or using RANS k and ε fields where available: $T_{\text{int}} = k/\varepsilon$ and $L \sim U_c T_{\text{int}}$.

5.5.3 Recommended PSD estimation procedure

To obtain robust spectra and to avoid misleading slopes:

1. **Record length:** acquire records with total duration T_{tot} substantially longer than the integral time scale T_{int}. Frequency resolution is $\Delta f = 1/T_{\text{tot}}$; if Δf is too large low-frequency energy will not be resolved.
2. **Sampling frequency:** sampling rate $f_s = 1/\Delta t$ must be chosen so that the Nyquist frequency $f_N = f_s/2 \gg f_\eta$. Otherwise the high-frequency band is aliased or truncated.

3. **Welch settings:** use a Hanning window, segment length N_{seg} that balances frequency resolution and variance (typical: 50% overlap, window length between $T_{\text{int}}/5$ and $T_{\text{int}}/1$), and average several segments to reduce variance.

4. **Ensemble averaging:** where possible compute PSDs for multiple SNGR realizations and average them to separate sampling noise from model bias.

5. **Compensated spectrum:** plot $f^{5/3}S_u(f)$ (compensated spectrum). A plateau in the compensated plot is a strong indicator of a $-5/3$ inertial subrange.

5.5.4 Reasons Why slopes may not match expectations

Some simulations run during this study did not always show a clear $-5/3$ range, differing the low/high frequency behaviour differs from the main guidelines. Most likely causes are:

- **Insufficient low-frequency content / short record:** if T_{tot} is too short then f_L is not resolved and the low-frequency part may appear artificially steep.
- **SNGR generation limits:** finite number of Fourier modes ($N = 500$) and the modal sampling strategy control how well large scales are represented. If large-scale (low k) modes are undersampled the energy-containing range will be weak.
- **Time-correlation tuning (f_τ):** the OU time scales T_n (via f_τ) control high-frequency content. If f_τ is too large (longer correlation), high frequencies are damped and the dissipation roll-off shifts to lower f.
- **Convective velocity mismatch:** using an incorrect U_c when mapping spatial to temporal scales or when labelling reference frequencies will shift the apparent position of the inertial subrange.
- **Probe location / flow physics:** near-wall, corner or shear-layer probes may not show a clear inertial range (local anisotropy and production dominate). The -5/3 law is an isotropic inertial-range result — it may be harder to see in highly anisotropic regions.
- **Numerical dissipation / filtering:** if the SNGR field is passed through CFD or post-processing with numerical smoothing, high frequencies can be removed.
- **Sampling and aliasing:** insufficient sampling rate or lack of anti-alias filtering will corrupt high-frequency estimates.

Figure 59 presents the Power Spectral Density (PSD) of the streamwise, wall-normal, and spanwise velocity fluctuations (u) as a function of frequency, plotted on a logarithmic scale. The vertical axis represents the turbulence energy spectrum in units of $m^2/s^2/\text{Hz}$, while the horizontal axis shows frequency in Hz.

Each curve corresponds to a virtual probe placed at a distinct location within the domain, with coordinates (x, y, z) indicated in the legend. These probes were strategically positioned to capture different flow features. Two reference lines are included for comparison in the three following images, represented in Figure 59:

- A dashed line with slope $-5/3$, representing the inertial subrange predicted by Kolmogorov's theory.
- A dotted line with slope -4, indicating the dissipation range where viscosity dominates.

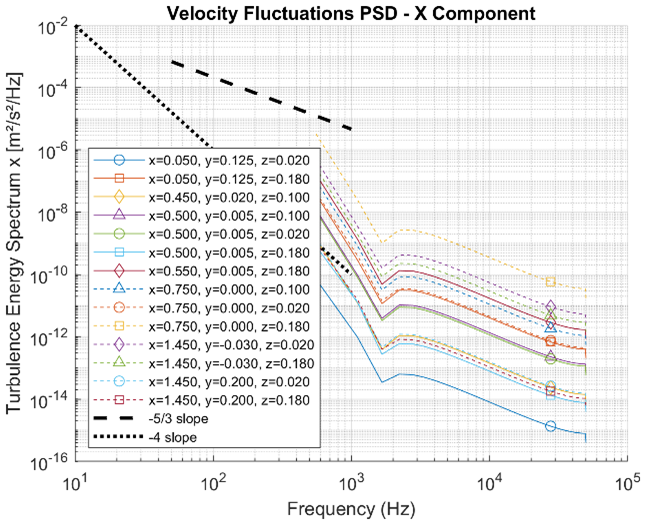

(a) PSD of the streamwise velocity component (u) along the mid-plane, highlighting the distribution of turbulent energy across frequencies.

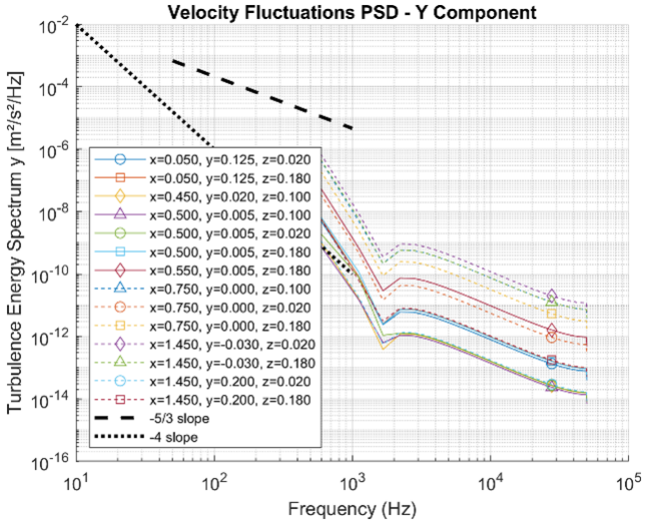

(b) PSD of the wall-normal velocity component (v), averaged spanwise and streamwise, showing the contribution of shear-layer and wake structures.

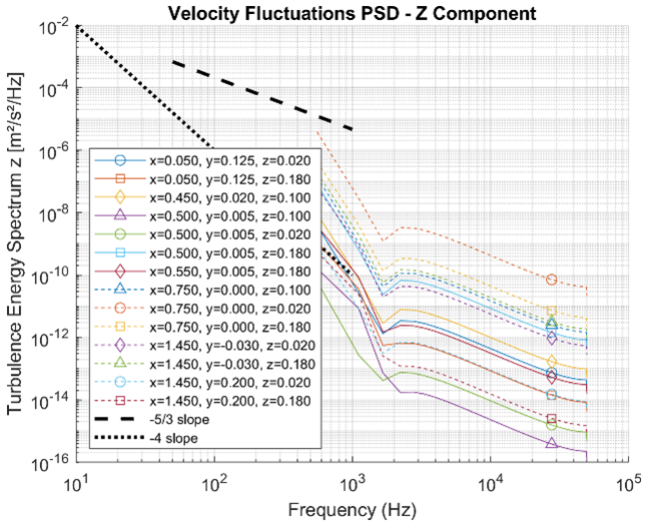

(c) PSD of the spanwise velocity component (w) from the RSM Elliptic Blending model in STAR-CCM+, used as a reference for comparison.

Figure 59: Power Spectral Density (PSD) analysis of turbulent velocity components generated using the SNGR method: (a) streamwise component along the mid-plane, (b) wall-normal component averaged across spanwise and streamwise directions, and (c) reference spanwise component from CFD simulation using the RSM Elliptic Blending turbulence model. All results correspond to a coarse mesh configuration.

Probes are grouped by physical location and the characteristic spectral behavior is described:

Upstream (incoming boundary layer): Probes placed upstream of the step (such as at $x = 0.050$ m, $y = 0.125$ m, $z = 0.020$ m and $z = 0.180$ m) lie within the incoming boundary layer and exhibit relatively flat spectra at low frequencies, consistent with the energy-containing range dominated by large eddies, indicating energy accumulation in large scales (energy-containing range) and limited fine-scale turbulence activity upstream of the expansion.

Pre-step / corner (shear and separation): Probes close to the step's location and corner ($x = 0.450$ m, $y = 0.005$ m, $z = 0.020$ m and $x = 0.500$ m, $y = 0.005$ m, $z = 0.100$ m) captures the transition region where shear-layer development begins to intensify, displaying elevated energy at intermediate and high frequencies, consistent with enhanced shear production and small-scale eddies caused by flow separation, being a clear sign of active turbulent mixing.

Near-wall locations: Near-wall probes (placed at $x = 0.500$ m, $y = 0.005$ m, $z = 0.020$ m and $x = 0.550$ m, $y = -0.001$ m, $z = 0.020$ m) are near the step corner, where high shear rates and flow separation occur, exhibiting stronger high-frequency content and an earlier steepening of the spectrum, reflecting viscous dissipation and

anisotropic small-scale motion in boundary layers, as well as indicating the presence of smaller eddies and enhanced turbulent mixing. This location captures near-wall turbulence, where fine-scale structures dominate. The spectrum steepens at high frequencies, approaching the expected -4 slope due to viscous dissipation, and being these effects strongest at low y. Overall, the PSD curves corresponding to these points show a clear inertial subrange with a slope close to $-5/3$, followed by a steep decay approaching -4 at higher frequencies.

Downstream / outlet: Downstream probes ($x = 0.750$ m to $x = 1.000$ m, various y and z) exhibit the clearest inertial subrange: a power-law segment whose slope is close to $-5/3$, followed by a dissipation roll-off near the -4 reference, reflecting these points the fully developed turbulent wake. The spectrum shows a well-defined inertial subrange with a slope close to $-5/3$, followed by a steep roll-off at higher frequencies. z: These probes capture the wake region downstream of the step. The spectra show well-defined inertial subranges with slopes close to $-5/3$, followed by steep roll-offs approaching -4, indicating a developed energy cascade in the wake.

In summary, the PSD plots shown in Figure 59 indicate that SNGR properly captures the qualitative frequency structure (energy-containing, inertial-like, and dissipation bands) across different flow regions.

5.6 Time marching approach: Turbulent Fluctuations

As previously explained in Section 5.4, the comparison of the TKE between CFD and SNGR is highly relevant. The backward-facing step model is now examined using the turbulent fluctuations represented and computed using the single time-loop approach from SNGR method, considering the same cutting planes that are visualized as well in STAR-CCM+. However, there is no direct means of visualizing instantaneous flow fields, as the CFD simulation employs a steady Reynolds-Averaged Navier–Stokes (RANS) approach with a Reynolds Stress Model (RSM). This modeling framework solves for time-averaged quantities, and therefore does not retain per-time-step snapshots of turbulent fluctuations.

Despite this limitation, STAR-CCM+ computes and stores the turbulent kinetic energy, k, at each computational cell. By plotting k on the same xy, xz, and yz planes, it is possible to obtain a spatial map indicating where and to what extent the model predicts the presence of turbulent fluctuations. Although the instantaneous structure of the turbulence cannot be visualized, the distribution of k provides valuable insight into the expected intensity and location of turbulent activity.

As seen in the following comparisons, the regions exhibiting the highest fluctuation intensities in the MATLAB visualizations correspond closely to areas with elevated TKE values in STAR-CCM+, thereby validating the consistency between the two representations.

To facilitate spatial correlation between the different planar visualizations of the turbulent kinetic energy (TKE) field, a consistent color-coding scheme has been employed across the subfigures. Each subfigure represents TKE values on a specific cutting plane — namely, the xy, xz, and yz planes, corresponding to the the streamwise, wall-normal, and spanwise direction, respectively — while also including colored lines that indicate the positions of the other two planes within the three-dimensional domain.

This is better shown in Figure 60, whose subfigure (a) displays the xy-plane at $z = 0.18\,\text{m}$, with a light blue contour outlining the extent of this plane. This same light blue line is superimposed in subfigures (b) and (c) to indicate the location of the $z = 0.18\,\text{m}$ slice within the xz and yz planes, respectively. Similarly, the yellow contour in subfigure (b), corresponding to the xz-plane at $y = 0.10\,\text{m}$, is also shown in subfigures (a) and (c) to mark the position of this plane. Finally, the red contour in subfigure (c), representing the yz-plane at $x = 0.25\,\text{m}$, is used in subfigures (a) and (b) to denote its spatial location.

This visual strategy enables a clear and intuitive understanding of how the different planes intersect within the flow domain, thereby enhancing the interpretability of the spatial distribution of TKE across the various cross-sections.

5.6.1 CFD - Low turbulent planes

First of all, the TKE values computed in relatively low turbulence regions of the flow are shown in Figure 60.

This subsection shows the RANS-computed TKE field on three low-turbulence cutting planes: the xy plane at $z = 0.18\,\text{m}$, the xz plane at $y = 0.10\,\text{m}$, and the yz plane at $x = 0.25\,\text{m}$. The views illustrate the spatial distribution of mean TKE and identify regions where the model predicts elevated fluctuation intensity (e.g., the shear layer immediately downstream of the step). Peak and representative values are reported to facilitate direct comparison with SNGR.

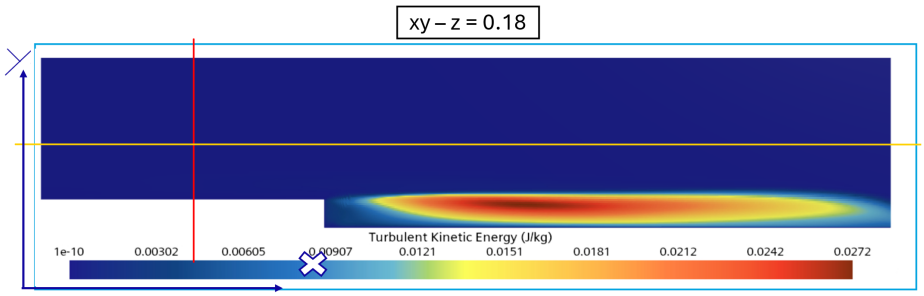

(a) TKE values along the streamwise direction at z = 0.18, estimated in StarCCM+ considering RSM EB as the turbulence model

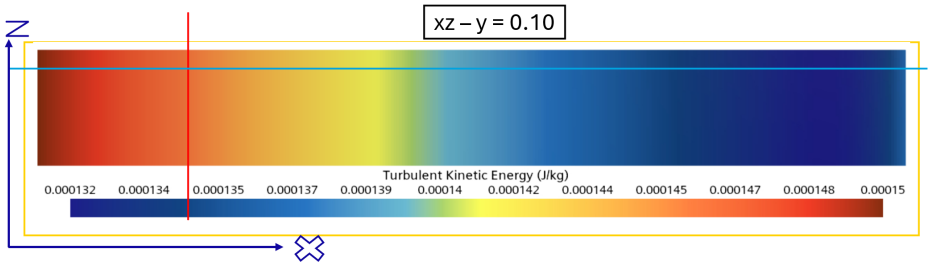

(b) TKE values along the streamwise direction at y = 0.10, estimated in StarCCM+ considering RSM EB as the turbulence model

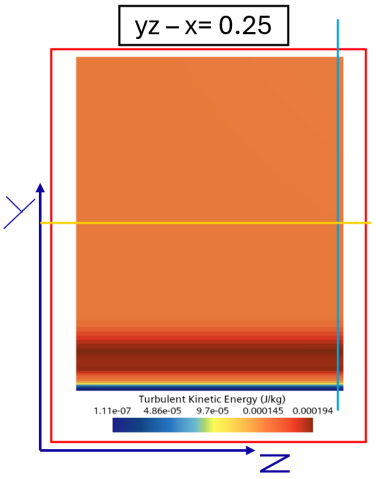

(c) TKE values along the streamwise direction at x = 0.25, estimated in StarCCM+ considering RSM EB as the turbulence model

Figure 60: TKE values along three planes of the BFS model, estimated in StarCCM+ considering RSM EB as the turbulence model.

The strongest turbulence is present after the step, at about slices $x = 4.0$ and $x = 6.0$, as seen in Figure 60a, reaching maximum values of 0.0272 J/kg, and extending far downstream with relatively high values until approximately $x = 15.0$, where at only 0.25 m from the outlet, TKE values are just a third smaller than at the peak (about 0.018 J/kg). At the slice located further downstream - at $x/h = 19.0$ - TKE is estimated to reach approximately 0.012 J/kg.

Figure 60b shows the xz plane, placed 0.10 m above the step. It is observed that the gradient of TKE values is almost negligible, where the difference between the lowest TKE values (close to the outley) and the highest (near the inlet) is 0.011%.

Furthermore, the low turbulence present at the inlet can be seen in Figure 60c, whose slice cuts a yz plane equidistant between the inlet and the step's corner. It shows very low turbulent values, in the order of 10^{-4}.

5.6.2 SNGR - Low turbulent planes

The same planes where the turbulent kinetic energy has been evaluated using CFD in Section 5.6.1, are shown according to the estimation made by means of the SNGR method, which provides images (that are saved as GIFs) of the evolution of the turbulent eddies, showing how they are convected and dissipated.

The plots presented in Figure 61 show the SNGR single time-loop estimation of turbulent fluctuations on the same three low-turbulence planes described in the previous page. In this case, SNGR produces instantaneous fluctuation fields and their time evolution, presenting representative frames and noting how the high-intensity regions compare spatially with CFD TKE contours.

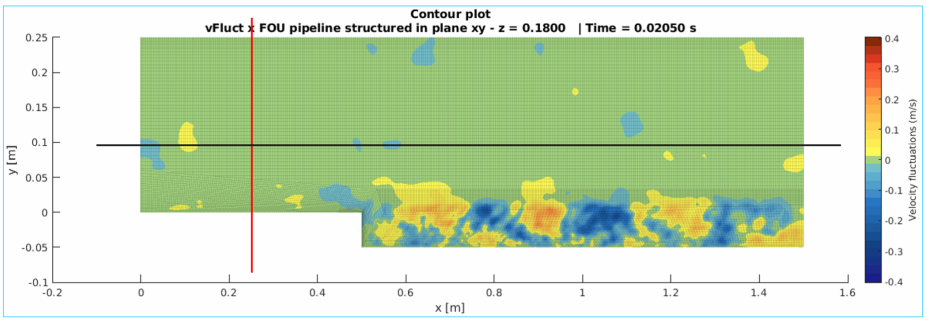

(a) TKE values along the streamwise direction at z = 0.18, estimated using the single time-loop approach in SNGR

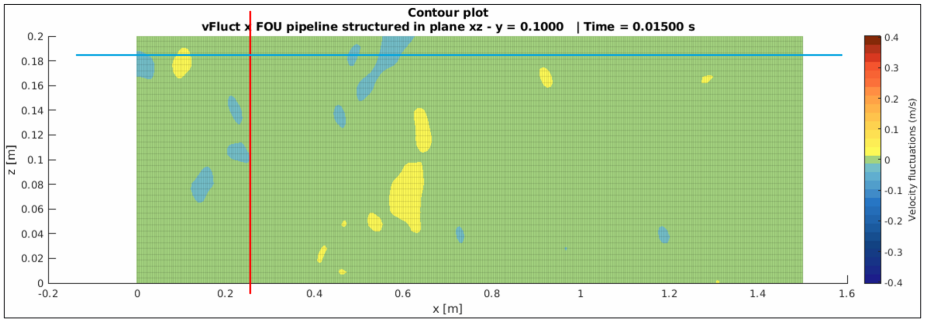

(b) TKE values along the streamwise direction at y = 0.10, estimated using the single time-loop approach in SNGR

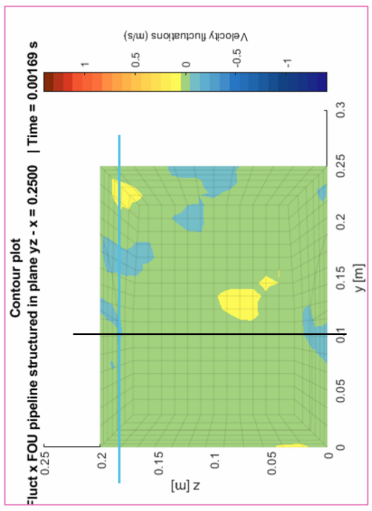

(c) TKE values along the streamwise direction at x = 0.25, estimated using the single time-loop approach in SNGR

Figure 61: TKE values along three planes of the BFS model, estimatedusing the single time-loop approach in SNGR.

As previously seen in the CFD analysis exhibited in Figure 60, a significant increase in fluctuation intensity is observed immediately downstream of the backward-facing step, corresponding to the development of a turbulent shear layer, as shown in Figure 61a. In this region, a wide range of eddy scales contributes to the signal. Further downstream and away from the step, individual vortical structures become discernible as they dissipate, illustrating the progressive breakdown of large-scale eddies during convection.

Above the step - at $y = 0.10$ m - turbulence is very low, as exhibited in Figure 61b, where most values are close to 0 m/s. The turbulent eddies represented by blue and yellow *bubbles* are convected and dissipated as they move towards the outlet.

5.6.3 CFD - High turbulent planes

The high-turbulence planes focus on regions close to the shear layer with stronger eddy activity (e.g., $z = 0.0025\,\text{m}$, $y = -0.0050\,\text{m}$, $x = 0.25\,\text{m}$). These RANS TKE maps highlight finer-scale energetic regions that are used to validate whether SNGR captures the expected increase in fluctuation intensity and eddy dispersion.

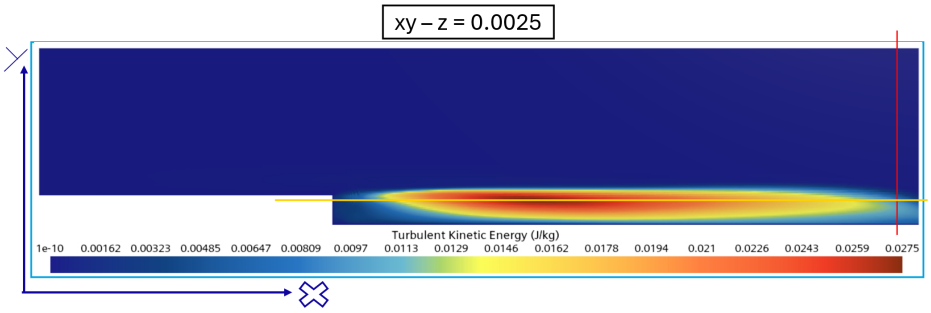

(a) TKE values along the streamwise direction at z = 0.0025, estimated in StarCCM+ considering RSM EB as the turbulence model

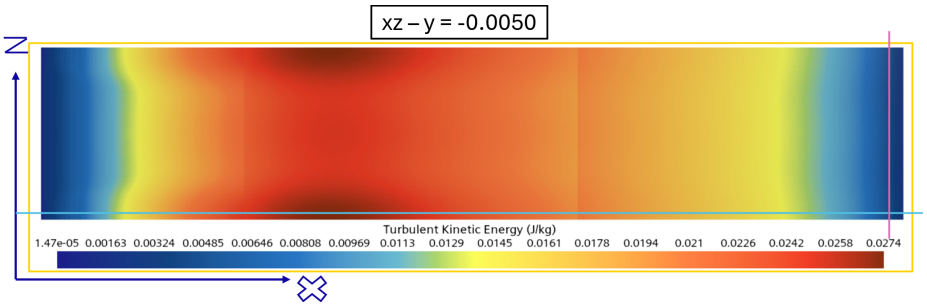

(b) TKE values along the streamwise direction at y = -0.0050, estimated in StarCCM+ considering RSM EB as the turbulence model

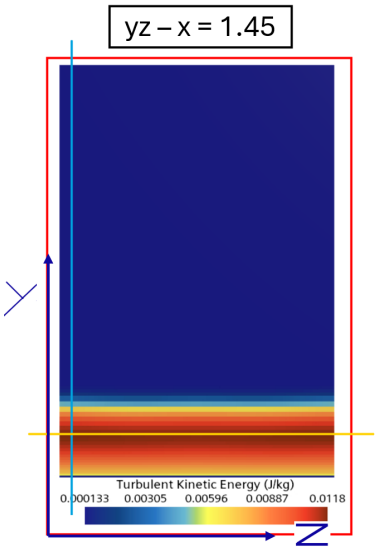

(c) TKE values along the streamwise direction at x = 1.45, estimated in StarCCM+ considering RSM EB as the turbulence model

Figure 62: TKE values along three planes of the BFS model, estimated in StarCCM+ considering RSM EB as the turbulence model.

Figure 62a shows the exact same distribution of TKE values that was presented in Section 5.6.1. This is because of the periodicity defined in the spanwise direction. The impact of considering planes within the recirculation zone can be seen in Figure 62b, whose xz plane lies just below the corner of the step (at $y = -0.0050$ m). Maximum values reach 0.0274 J/kg, which is the same largest magnitude observed in Figure 62a.

Concerning the yz plane - represented in Figure 62c - is located just before the outlet and shows, as the other two planes previously analized, insignificant TKE values above the step - at $y = 0.05$ m - and relatively high activity in the region below the step, with peak values of 0.0118 J/kg, matching the levels see in Figure 60a.

5.6.4 SNGR - High turbulent planes

SNGR results for the high-turbulence planes are shown here as still frames and GIF sequences to illustrate the formation, convection and dissipation of energetic eddies. The SNGR instantaneous structures are compared against the mean TKE maps from CFD - evaluated before in Section 5.6.3 to evaluate spatial agreement and the qualitative fidelity of eddy dynamics.

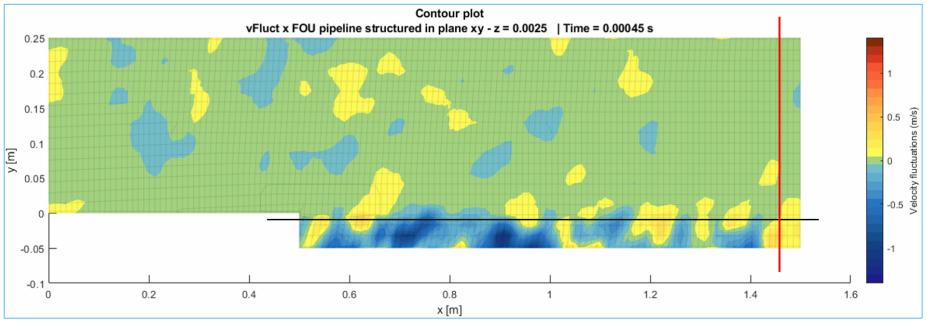

(a) TKE values along the streamwise direction at z = 0.0025, estimated in StarCCM+ considering RSM EB as the turbulence model

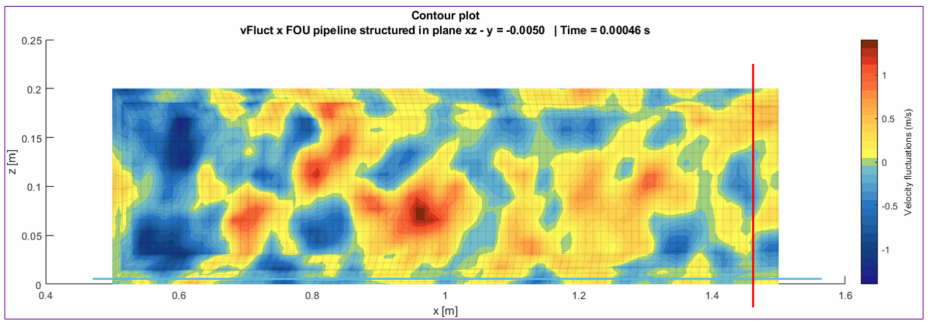

(b) TKE values along the streamwise direction at y = -0.0050, estimated in StarCCM+ considering RSM EB as the turbulence model

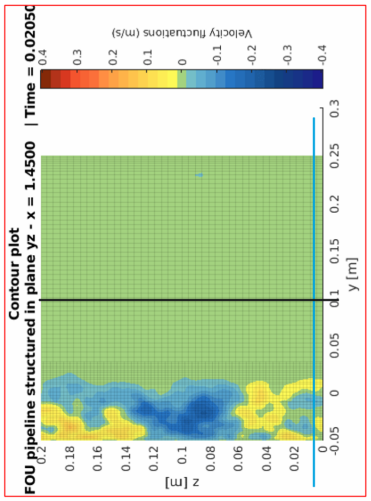

(c) TKE values along the streamwise direction at x = 1.45, estimated in StarCCM+ considering RSM EB as the turbulence model

Figure 63: TKE values along three planes of the BFS model, estimated in StarCCM+ considering RSM EB as the turbulence model.

Finally, the velocity fluctuations in the plane yz are exhibited in Figure 63c. It is observed that, compared to the CFD analysis at this slice, the turbulent fluctuations vary considerably along the spanwise direction.

In summary, SNGR's single time-loop approach provides instantaneous fluctuation fields that correlate well with the mean TKE distribution obtained from the RANS (RSM EB) simulation: the location and overall extent of energetic regions agree, while differences are concentrated in the magnitude of outer-layer overshoots and the smallest resolved scales.

6 Errors and uncertainties

A perfect measurement does not exist. All measurements inherently possess errors and uncertainties, regardless of efforts to mitigate them. Comprehending potential errors is a critical concern in any experimental science, as the conclusions extracted from data depend on how these errors and uncertainties are controlled Columbia University. There are two main types of uncertainties:

Systematic Uncertainties

Systematic errors and uncertainties are biases that introduce consistent deviations from the true value. As they are not randomly changing, the number of measurements does not help to reduce them, and they are often difficult to identify. The main advantage is that they can be corrected once identified. Thus, by identifying potential sources of systematic error in advance, it is often possible to design and conduct experiments with the right considerations to mitigate the limitations of the equipment.

Examples in the Analysis

- **Numerical Discretization Errors (StarCCM+ & MATLAB)**
  - The spatial and temporal discretizations considered for the DNS and in the SNGR method can introduce systematic errors. Low-order discretizations schemes (such as first-order upwind) introduce excessive numerical diffusion, leading to underestimating velocity gradients in these turbulent flows.
  - The grid resolution: to be as accurate as possible, all turbulence structures should be properly solved, meaning that. Otherwise, there would be a systematic underestimation of turbulence intensity and Reynolds stresses.
- **Turbulence Model Assumptions (StarCCM+ RANS Simulation)**
  - RANS models introduce systematic errors as they rely on assumptions that model the turbulent structures. The biggest impact is experienced in separated flows such as the reattachment region in backward-facing step flow.
  - First-order closures (like k-ϵ) assume isotropic turbulence, while higher-order models (such as Reynolds Stress Models Launder et al. [1975] or Bradshaw's anisotropy corrections Bradshaw [1973]) account for directional effects, introducing biases in the turbulent velocity fluctuations generated by means of the SNGR method, particularly in separated flows (Section 4.2.1).
  - RANS modelled k can be biased (over- or under-predicted) in separated flows; propagating such a bias into SNGR will have an influence on predicting the synthetic turbulence and the resulting noise. Besides, there could be a global scaling mismatch between the synthetic history and a reference Reynolds-stress field, and since $R_{ij} \propto u'^2$, a small multiplicative mismatch in u' produces a squared mismatch in the components.
- **Boundary Conditions and Domain Size**
  - Setting the parameters properly ensures that the comparison between the DNS study and the RANS simulation performed using StarCCM+ is fair and that the results correspond to the expected true values (limited by the method of solving the governing equations). Otherwise, it would lead to biased predictions of turbulence characteristics.
  - A truncated computational domain can cause artificial reflections at boundaries, influencing flow variables systematically, as they may cause disturbances to reflect in the domain.

Furthermore, systematic errors may arise if the SNGR method is applied to compressible flows without accounting for density fluctuations, as highlighted by Marvin Marvin [1977].

Random Errors

Random errors cause variability in results due to inherent noise, numerical precision limitations, or stochastic variations in experiments and simulations. These errors can be reduced by repeating simulations or measurements and analyzing statistical convergence. Contrary to systematic errors and uncertainties, random ones are unbiased, meaning that any value from a measurement is as likely to be too high as it is to be too low Columbia University.

Examples in this thesis report

- **Convergence Tolerance in StarCCM+**
  - The convergence criterion considered for the simulation must be large enough to make sure there are not persistent small fluctuations, which may lead to variations in results across different runs. For the coarse mesh, all residuals being evaluated reach at least values in the order of 10^-4 (with the exception of the Reynolds stresses cross-components $u'w'$ and $v'w'$). Likewise, for the refined mesh converge is considered to be reached when every residual - but the previously mentioned R_ij cross-components - are in the order of 10^-6 or lower.
- **Monte Carlo-Based Methods (MATLAB)**
  - When applying stochastic noise generation (SNGR) techniques, the random seed selection affects the resulting noise field. Different runs may yield slight variations in velocity fluctuations, even though the mean trends remain consistent.
- **Floating-Point Precision (MATLAB & StarCCM+)**
  - Due to finite precision in numerical computations, small rounding errors accumulate, leading to slight variations in results, particularly when handling very small or very large numerical values.

7 Conclusions

This thesis investigated the capabilities, accuracy, and limitations of the Stochastic Noise Generation and Radiation (SNGR) method when reconstructing turbulent fluctuations from steady RANS statistics. The work was structured around three research questions, each arising directly from the objectives of quantifying uncertainty, validating the method against DNS, and assessing how temporal correlation and convection influence SNGR accuracy. The conclusions are presented below.

RQ1 — How can the quantified uncertainties guide improvements to the SNGR code for internal flow applications?

The uncertainty analysis revealed that deviations between SNGR and DNS arise from two distinct sources: (1) **inherited model bias** from the RANS input, and (2) **reconstruction artifacts** within the SNGR algorithm. By isolating these sources, specific code improvements were identified.

First, the analysis of the Reynolds shear stress ($\langle u'v' \rangle$) uncovered a systematic failure to reproduce cross-correlations despite accurate normal stresses. This was traced to a flaw in the isotropic fluctuation generation step (detailed in Appendix B), where the initial field contained non-physical correlations that the anisotropic mapping could not correct. Fixing this initialization to ensure a truly uncorrelated seed field is the primary code improvement required to restore physical shear stresses.

Second, the analysis confirmed that SNGR accuracy is strictly bounded by the quality of the input. Therefore, improvements for internal flows must prioritize the robustness of the RANS ingestion pipeline (e.g., ensuring valid tensor exports and grid resolution of $y^+ < 1$) rather than solely tuning SNGR parameters.

RQ2 — How accurately does the SNGR code replicate DNS reference data for Backward-Facing Step flows?

The results demonstrate that SNGR successfully reconstructs the spatial distribution of turbulence with good agreement in the outer shear layer and free-stream regions. For the backward-facing step, SNGR captures the overall energy distribution of the separated shear layer and the mean flow topology as prescribed by the RANS input. The visualization of instantaneous structures confirms that SNGR reconstructs realistic eddy dynamics, including coherent convection downstream of the step.

However, accuracy in the recirculation zone is limited by the input data: because the RANS simulation (specifically the dataset exported for processing) predicted a shallow separation with no significant reverse flow, the SNGR reconstruction faithfully reproduced this profile and thus could not recover the strong recirculation bubble observed in DNS.

Despite this inherited limitation, the power spectral density analysis confirmed that the method correctly reproduces the inertial-range slopes ($-5/3$) and high-frequency roll-off, proving that the synthetic fluctuations are spectrally consistent with physical turbulence.

RQ3 — To what extent do temporal correlation and convective transport enhance the physical fidelity of SNGR, and what are the statistical convergence constraints associated with these mechanisms?

The study confirms that temporal correlation and convective transport are indispensable for achieving physical in the reconstructed field, but they introduce significant statistical convergence challenges.

- **Physical Fidelity:** The inclusion of convective transport allows the Single Time-Loop implementation to reproduce realistic eddy advection and spatial coherence that the Ensemble-Averaged approach lacks. This was evidenced by the correct reconstruction of the inertial subrange in the Power Spectral Density (PSD) analysis, which requires accurate temporal correlation to match Kolmogorov scaling. For applications such as flow-induced vibration or aeroacoustics, where frequency content is critical, these mechanisms are essential.
- **Convergence Constraints:** The analysis revealed a critical trade-off: while the Time-Loop approach captures dynamics, it suffers from finite-sampling artifacts (observed as "spikes" in the Reynolds stress profiles) when the simulation time T_{sim} is insufficient relative to the integral time scale T_{int}. Unlike the Ensemble method, which converges rapidly to a smooth mean, the Time-Loop method requires significantly longer integration times to achieve the same statistical smoothness. Therefore, while temporal mechanisms improve physical fidelity, they impose a higher computational cost to satisfy the effective sample size (N_{eff}) required for statistical convergence.

Final remarks

In conclusion, this study demonstrates that SNGR is a promising **intermediate-fidelity method** for embedding physically consistent turbulent fluctuations onto RANS mean fields. It successfully reproduces the dominant spectral content and spatial structures observed in DNS while maintaining low computational cost. However, some improvements must be made, particularly in **covariance preservation** (fixing the shear-stress reconstruction) and **sampling convergence**. With these refinements, SNGR offers a practical balance between accuracy and computational efficiency for industrial applications such as aeroacoustics and flow-induced vibration.

References

Reynolds-averaged navier-stokes equations for turbulence modeling. *Applied Mechanics Reviews*, July 2009. doi: 10.1115/1.3124648. URL https://www.researchgate.net/publication/245371007.

A. Acuña, T. Hauser, and C. Rumsey. Cfd general notation system (cgns): Standard for cfd data interchange. In *AIAA Paper 1999-3306*, 1999. doi: 10.2514/6.1999-3306.

B. F. Armaly, F. Durst, J. C. F. Pereira, and B. Schönung. Experimental and theoretical investigation of backward-facing step flow. *Journal of Fluid Mechanics*, 127:473–496, 1983.

ASML. A sustainability mindset. https://www.asml.com/en/news/stories/2024/a-sustainability-mindset, 2024. Accessed: 2025-03-10.

D. Barkley, M. G. M. Gomes, and R. D. Henderson. Three-dimensional instability in flow over a backward-facing step. *Journal of Fluid Mechanics*, 473:167–190, 2002.

Mustafa Barri, George K. El Khoury, Helge I. Andersson, and Bjørnar Pettersen. Dns of backward-facing step flow with fully turbulent inflow. *International Journal for Numerical Methods in Fluids*, 64(7):777–792, 2010. doi: 10.1002/fld.2176. Published online 16 September 2009.

W. Bechara et al. Stochastic approach to noise modeling for free turbulent flows. *AIAA Journal*, 32:455–463, March 1994. doi: 10.2514/3.12008.

L. Biferale and I. Procaccia. Anisotropy in turbulent flows and in turbulent transport. *Physics Reports*, 414:43–164, 2005. doi: 10.1016/j.physrep.2005.04.001.

M. Billson, L. Eriksson, et al. Modeling of synthesized anisotropic turbulence and its sound emission. *AIAA Journal*, 1, May 2004. doi: 10.2514/6.2004-2857.

Joseph Boussinesq. *Théorie de l'écoulement tourbillonnant et tumultueux des liquides dans les lits rectilignes à grande section*. Gauthier-Villars & Fils, Paris, 1897.

P. Bradshaw. The turbulence structure of equilibrium boundary layers. *Journal of Fluid Mechanics*, 29:625–645, 1967. doi: 10.1017/S0022112067001062.

P. Bradshaw. Effects of streamline curvature on turbulent flow. AGARDograph 169, AGARD, 1973. Anisotropy modeling.

T. Cebeci and A. M. O. Smith. *Analysis of Turbulent Boundary Layers*. Academic Press, 1974.

CGNS Steering Committee. *CFD General Notation System (CGNS): Standard Interface Data Structures (SIDS)*, 2022. Available at https://cgns.github.io.

T. P. Chiang and Tony W. H. Sheu. A numerical revisit of backward-facing step flow problem. *Physics of Fluids*, 11 (4):862–874, April 1999. doi: 10.1063/1.869958.

A. J. Chorin. Numerical solution of the navier–stokes equations for an incompressible fluid. *Mathematics of Computation*, 22(104):745–762, 1968. doi: 10.1090/S0025-5718-1968-0242392-2. URL https://math.berkeley.edu/~chorin/chorin68.pdf.

Columbia University. Introduction to error and uncertainty. https://www.physics.columbia.edu/files/content. Accessed: March 12, 2025.

H. Eckelmann. The structure of the viscous sublayer and the adjacent wall region in a turbulent channel flow. *Journal of Fluid Mechanics*, 65(3):439–459, 1974. doi: 10.1017/S0022112074001479.

Yuanxuan Fang and Yunfei He. Resolution technology of lithography machine. *Journal of Physics: Conference Series*, 2221:012041, 2022. doi: 10.1088/1742-6596/2221/1/012041. 2nd International Conference on Electronics, Circuits and Information Engineering (ECIE-2022), 07/01/2022–09/01/2022, Online.

Hervé Guillard and Cécile Viozat. On the behaviour of upwind schemes in the low mach number limit. *Computers & Fluids*, 28(1):63–86, 1999. URL https://hal.archives-ouvertes.fr/hal-01601857/document.

K. Hanjalic, M. Popovac, and M. Hadziabdic. A robust near-wall elliptic-relaxation eddy-viscosity turbulence model for cfd. *International Journal of Heat and Fluid Flow*, 25(6):1047–1051, 2004.

Fluent Inc. *FLUENT 6.3 User Guide*, September 29 2006. URL https://www.ansys.com. Modeling Turbulence: Standard, RNG, and Realizable $k - \epsilon$ Models.

Javier Jiménez. The contributions of a. n. kolmogorov to the theory of turbulence. *Arbor*, 178(704), 2004. doi: 10.3989/arbor.2004.i704.550.

W. P. Jones and B. E. Launder. The prediction of laminarization with a two-equation model of turbulence. *International Journal of Heat and Mass Transfer*, 15(2):301–314, 1972a. ISSN 0017-9310. doi: 10.1016/0017-9310(72)90076-2.

W. P. Jones and B. E. Launder. The prediction of laminarization with a two-equation model of turbulence. *International Journal of Heat and Mass Transfer*, 15:301–314, 1972b. doi: 10.1016/0017-9310(72)90076-2.

W. Kaiser. The evolvement of lithography optics towards advanced euv lithography: Enabling the continuation of moore's law for six decades. *IEEE Electron Devices Magazine*, 2(1):23–34, March 2024. doi: 10.1109/MED.2023.3343627.

M. Karweit et al. Simulation of the propagation of an acoustic wave through a turbulent velocity field: A study of phase variance. *The Journal of the Acoustical Society of America*, 89:52–61, January 1991. doi: 10.1121/1.400415.

D. Kazazis, J. G. Santaclara, J. van Schoot, et al. Extreme ultraviolet lithography. *Nature Reviews Methods Primers*, 4:84, 2024. doi: 10.1038/s43586-024-00361-z. Accepted 03 October 2024; Published 28 November 2024.

A. Khavaran. Role of anisotropy in turbulent mixing noise. *AIAA Journal*, 37(7):832–841, 1999. ISSN 0001-1452. doi: 10.2514/2.7531.

J. Kim, P. Moin, and R. D. Moser. Turbulence statistics in fully developed turbulent channel flow at low reynolds number. *Journal of Fluid Mechanics*, 177:133–166, 1987. doi: 10.1017/S0022112087000892.

A. N. Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Doklady Akademii Nauk SSSR*, 30:301–305, 1941. (English translation: Proc. R. Soc. A (1991) 434, 9–13).

B. E. Launder and D. B. Spalding. The numerical computation of turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, 3:269–289, 1974. doi: 10.1016/0045-7825(74)90029-2.

B. E. Launder, G. J. Reece, and W. Rodi. Progress in the development of a reynolds-stress turbulence closure. *Journal of Fluid Mechanics*, 68(3):537–566, 1975. doi: 10.1017/S0022112075001814. Higher-order closures.

H. Le and P. Moin. Direct numerical simulation of turbulent flow over a backward-facing step. Annual Research Briefs N94-12296, Center for Turbulence Research, Stanford University, 1992. Annual Research Briefs 1992.

Hung Le, Parviz Moin, and John Kim. Direct numerical simulation of turbulent flow over a backward-facing step. *Journal of Fluid Mechanics*, 330:349–374, 1997. doi: 10.1017/S0022112097003780.

Myoungkyu Lee and Robert D. Moser. Direct numerical simulation of turbulent channel flow up to $Re_\tau \approx 5200$. *Journal of Fluid Mechanics*, 774:395–415, 2015. doi: 10.1017/jfm.2015.160.

M. v. den Brink. Continued Scaling in Semiconductor Manufacturing Enabled by Advances in Lithography. In *2019 IEEE International Electron Devices Meeting (IEDM)*, pages 1.2.1–1.2.5, San Francisco, CA, USA, 2019. doi: 10.1109/IEDM19573.2019.8993590.

R. Manceau and K. Hanjalić. Elliptic blending reynolds–stress model of near wall turbulence. *Journal of Fluid Mechanics*, 441:199–222, 2002.

I. Marusic, B. J. McKeon, P. J. Monkewitz, H. M. Nagib, A. J. Smits, and K. R. Sreenivasan. Wall-bounded turbulent flows at high reynolds numbers: Recent advances and key issues. *Physics of Fluids*, 22(6):065103, 2010. doi: 10.1063/1.3453711.

J. G. Marvin. Turbulence modeling for compressible flows. Technical Memorandum NASA-TM-X-73188, NASA Ames Research Center, Moffett Field, CA, United States, 1977. URL https://ntrs.nasa.gov/citations/19770008398. Document ID: 19770008398.

Max Planck Institute for Dynamics and Self-Organization. Max planck institute for fluid dynamics. https://www.ds.mpg.de/, 2025. Accessed: 2025-03-11.

F. R. Menter. Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8):1598–1605, 1994. k-ω SST model development.

F. R. Menter, M. Kuntz, and R. Langtry. Ten years of industrial experience with the sst turbulence model. *Turbulence, Heat and Mass Transfer*, 4:625–632, 2003.

Florian R. Menter. Review of the shear-stress transport turbulence model experience from an industrial perspective. *International Journal of Computational Fluid Dynamics*, 23(4):305–316, 2009.

P. Moin and K. Mahesh. Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics*, 30:539–578, 1998a. doi: 10.1146/annurev.fluid.30.1.539.

Parviz Moin and Krishnan Mahesh. Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics*, 30(1):539–578, 1998b.

A. S. Monin and A. M. Yaglom. *Statistical Fluid Mechanics: Mechanics of Turbulence*, volume 1. MIT Press, Cambridge, MA, 1971.

G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38:114–117, 1965. Classic paper.

M. V. Morkovin. Effects of compressibility on turbulent flows. In *Mécanique de la Turbulence*. CNRS, 1962.

Neways Electronics. New. *NEW*, 2024. URL https://newayselectronics.com/new-magazine/. Accessed: 2025-03-11.

Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.

L. Prandtl. Zur turbulenten strömung in rohren und längs platten. *Ergebnisse der Aerodynamischen Versuchsanstalt Göttingen*, 1925.

L. Prandtl. Report on investigation of developed turbulence. Technical Report TM 1231, National Advisory Committee for Aeronautics (NACA), 1949. Translation from Zeitschrift fuer Angewandte Mathematik und Mechanik, vol. 5, no. 2, pp. 136–139.

P. G. Saffman. The large-scale structure of homogeneous turbulence. *Journal of Fluid Mechanics*, 41(2):327–343, 1970. doi: 10.1017/S002211207000064X.

Pierre Sagaut. *Large Eddy Simulation for Incompressible Flows: An Introduction*. Springer Science & Business Media, 2006.

P. Schlatter and R. Örlü. Assessment of direct numerical simulation data of turbulent boundary layers. *Journal of Fluid Mechanics*, 659:116–126, 2010. doi: 10.1017/S0022112010003113.

Heinz Schlichting and Klaus Gersten. *Boundary-Layer Theory*. Springer, Berlin / Heidelberg, 8 edition, 2000. ISBN 3-540-66270-7.

François G. Schmitt. About boussinesq's turbulent viscosity hypothesis: historical remarks and a direct evaluation of its validity. *Comptes Rendus Mécanique*, 335(9–10):617–627, 2007. doi: 10.1016/j.crme.2007.08.004.

Stefan Schoder. openCFS-Data: Implementation of the Stochastic Noise Generation and Radiation Model (SNGR). *arXiv preprint arXiv:2302.05622*, 2023. doi: 10.48550/arXiv.2302.05622. URL https://arxiv.org/abs/2302.05622.

M. Siba, W. Wanmahmood, M. Zaki Nuawi, R. Rasani, and M. Nassir. Flow-induced vibration in pipes: Challenges and solutions – a review. *Journal of Mechanical Engineering and Sciences*, 15(3):8131–8150, 2021. doi: 10.15282/jmes.15.3.2021.12.0636.

A. J. Smits, B. J. McKeon, and I. Marusic. High-reynolds-number wall turbulence. *Annual Review of Fluid Mechanics*, 43:353–375, 2011. doi: 10.1146/annurev-fluid-122109-160730.

P. R. Spalart and S. R. Allmaras. A one-equation turbulence model for aerodynamic flows. *AIAA Paper*, (92-0439), 1992.

Philippe R. Spalart. Direct simulation of a turbulent boundary layer up to $r_\theta = 1410$. Technical Report NASA-TM-89407, National Aeronautics and Space Administration, December 1986. N87-16030, CSCL 20D, Unclassified, 63/34 43301.

S. Tavoularis and S. Corrsin. Experiments in nearly homogeneous turbulent shear flow with a uniform mean temperature gradient. part 1. *Journal of Fluid Mechanics*, 104:311–347, 1981. doi: 10.1017/S0022112081003116.

G. I. Taylor. Statistical theory of turbulence. *Proceedings of the Royal Society A*, 151:421, 1935. doi: 10.1098/rspa.1935.0158.

H. Tennekes and J. L. Lumley. *A First Course in Turbulence*. MIT Press, Cambridge, MA, 1972a.

H. Tennekes and J. L. Lumley. *A First Course in Turbulence*. MIT Press, Cambridge, MA, 1972b.

J. van Schoot. Exposure tool development toward advanced euv lithography: a journey of 40 years driving moore's law. *IEEE Electron Devices Magazine*, 2:8–22, 2024.

S. Wallin and A. V. Johansson. An explicit algebraic reynolds stress model for incompressible and compressible turbulent flows. *Journal of Fluid Mechanics*, 403:89–132, 2000.

Michael Weitz, Stefan Schoder, and Manfred Kaltenbacher. A Stochastic Approach to Compute Cavity Noise using SNGR. In *Waves 2019*, pages 154–155. Technische Universität Wien, 2019. ISBN 978-3-200-06511-6.

David C. Wilcox. *Turbulence Modeling for CFD*. DCW Industries, La Canada, CA, USA, 3rd edition, 1998.

A Calculation of the pressure jump required to obtain a target friction Reynolds number Re_τ for the TCF model

This appendix shows the derivation and numerical evaluation used to select the pressure jump (streamwise pressure difference across the periodic domain) so that the fully-developed turbulent channel reaches a target friction Reynolds number $Re_\tau = 180$. All quantities are expressed in SI units. First, the governing relations are established, and then the specific values corresponding to this study are replaced into those equations.

For a plane channel of half-height δ the friction velocity u_τ, wall shear stress τ_w, and friction Reynolds number Re_τ are related by

$$u_\tau = \frac{Re_\tau\,\nu}{\delta}, \tag{81}$$

$$\tau_w = \rho\,u_\tau^2, \tag{82}$$

where ρ is the fluid density and ν the kinematic viscosity. For a fully-developed channel the streamwise pressure gradient balances the wall shear (per unit half-height), so

$$\left|\frac{dp}{dx}\right| \;=\; \frac{\tau_w}{\delta}. \tag{83}$$

The pressure jump Δp applied across the periodic domain of streamwise length L_x is then

$$\Delta p \;=\; \left|\frac{dp}{dx}\right| L_x \;=\; \frac{\rho\,Re_\tau^2\,\nu^2\,L_x}{\delta^3}. \tag{84}$$

Numeric example used in this work

Using the values specified in the TCF section of the report:

$$\rho = 1.225\ \mathrm{kg\,m^{-3}}, \qquad \nu = 5.0\times10^{-5}\ \mathrm{m^2\,s^{-1}}, \qquad \delta = 1.0\ \mathrm{m}, \qquad L_x = 8\pi\ \mathrm{m}, \qquad Re_\tau = 180,$$

the following value are obtained, as seen from Equation (85) to Equation (88).

$$u_\tau = \frac{Re_\tau\,\nu}{\delta} = \frac{180\cdot 5.0\times10^{-5}}{1.0} = 9.00\times10^{-3}\ \mathrm{m\,s^{-1}} \tag{85}$$

$$\tau_w = \rho\,u_\tau^2 = 1.225\cdot(9.00\times10^{-3})^2 \approx 9.9225\times10^{-5}\ \mathrm{Pa} \tag{86}$$

$$\left|\frac{dp}{dx}\right| = \frac{\tau_w}{\delta} \approx 9.9225\times10^{-5}\ \mathrm{Pa\,m^{-1}} \tag{87}$$

$$\Delta p = \left|\frac{dp}{dx}\right| L_x \approx 9.9225\times10^{-5}\cdot(8\pi) \approx 2.494\times10^{-3}\ \mathrm{Pa} \tag{88}$$

Thus, the analytic estimate for the pressure jump that produces $Re_\tau = 180$ in the specified domain is

$$\boxed{\Delta p \approx 2.49\times10^{-3}\ \mathrm{Pa}}.$$

Note on geometrical scaling

If a different half-height δ or streamwise domain length L_x is used, substitute those values into equation (84) to compute the corresponding Δp. In the BFS case the relevant length-scale is the step height h; replace δ by h (or otherwise use the appropriate characteristic length) when deriving the pressure jump for that geometry.

B Final Correction: Fixing the Generation of Isotropic Fluctuations

After implementing a correct anisotropy transformation, the persistence of a near-zero Reynolds shear stress ($\langle u'v' \rangle \approx 0$) pointed to a more fundamental issue. A detailed analysis of the code responsible for generating the initial velocity field revealed that the field was not truly isotropic and uncorrelated as required by the SNGR theory. This appendix details the flaw and presents the final, corrected implementation.

B.1 The Flaw in the Initial Fluctuation Generation

The root cause of the issue was traced to the function 'computeTurbulentVelocityField.m' and its sub-function 'getInputParamsUfluct.m'. The implementation created a strong, non-physical spatial correlation in the initial "isotropic" field.

1. **Calculation of a Single Direction Vector per Mode:** The function 'getInputParamsUfluct.m' calculates a single direction vector, σ_n, for each Fourier mode n. This vector is derived from the cross product of two random vectors, ζ_n and κ_n:

   ```
   ...
   sigma = cross(zetaVec, kappaVec, 2);
   sigma = sigma ./ vecnorm(sigma, 2, 2);
   ```

   The resulting 'sigma' is a matrix of size '[nModes x 3]', where each row 'sigma(n,:)' is a constant vector used for the entirety of the n-th mode.

2. **Uniform Application Across All Spatial Locations:** This constant direction vector is then passed to 'calcUfluctBatch.m', where the velocity fluctuation at each spatial point $\mathbf{x}$ is synthesized:

   $$\mathbf{u}'(\mathbf{x}) = \sum_{n=1}^{N} A_n(\mathbf{x}) \cdot \cos(\kappa_n \cdot \mathbf{x} + \psi_n) \cdot \sigma_n$$

   While the scalar part of this summation varies correctly with position $\mathbf{x}$ through the cosine term, the directional part, σ_n, remains fixed for all cells for a given mode.

3. **Creation of Artificial Correlation:** This leads to a rigid, non-physical correlation between velocity components. For any single mode n, the ratio of the generated velocity components is constant everywhere in space:

   $$\frac{u'_n(\mathbf{x})}{v'_n(\mathbf{x})} = \frac{\sigma_{n,x}}{\sigma_{n,y}} = \text{constant for all } \mathbf{x}$$

   When summed over all modes, this artificial correlation does not average out to zero, resulting in a 'vFluctHistory' field that is strongly, and incorrectly, correlated before the anisotropy step is even applied. The field is not isotropic.

4. **Why the Anisotropy Function Failed:** The corrected anisotropy function, 'apply target anisotropy simplified', assumes its input is an uncorrelated, isotropic field. When fed a field that is already contaminated with a strong, incorrect correlation, it cannot produce the correct anisotropic output. The final result is a superposition of the flawed initial correlation and the intended target correlation, which fails to reproduce the desired $\langle u'v' \rangle$.

B.2 The Correct Implementation: A New `computeTurbulentVelocityField.m`

To fix this, the entire fluctuation generation process must be replaced with a method that generates a truly isotropic, divergence-free field. The following code implements the standard projection method, which starts with three independent random fields and projects them to satisfy the divergence-free condition.

C Inlet Velocity Profile Implementation for the Backward-facing step from 1997

The inlet velocity profile for the backward-facing step simulation was extracted from the DNS data of **?**, where the mean streamwise velocity U^+ is provided in wall units (y^+). The profile corresponds to a turbulent boundary layer at $Re_h = 5100$ with $\delta_{99} = 1.2h$.

Conversion Methodology

The dimensionless quantities were converted to physical units using:

$$u_\tau = 0.0384 U_0, \quad \nu = \frac{U_0 h}{Re_h} \tag{89}$$

$$y = \frac{y^+ \nu}{u_\tau}, \quad U = U^+ u_\tau \tag{90}$$

where $U_0 = 1\,\mathrm{m\,s^{-1}}$ and $h = 1\,\mathrm{m}$.

Table 4: Complete velocity profile conversion from DNS wall units to physical quantities

y^+	U^+	y (m)	U ($\mathrm{m\,s^{-1}}$)
1.994	2.886	1.02×10^{-4}	0.111
2.373	3.254	1.21×10^{-4}	0.125
2.887	3.783	1.47×10^{-4}	0.145
3.513	4.450	1.79×10^{-4}	0.171
4.276	5.071	2.18×10^{-4}	0.195
5.147	5.854	2.63×10^{-4}	0.225
5.990	6.544	3.06×10^{-4}	0.251
7.845	7.717	4.01×10^{-4}	0.296
10.453	9.281	5.34×10^{-4}	0.356
14.732	11.006	7.53×10^{-4}	0.423
18.863	11.995	9.64×10^{-4}	0.461
23.613	12.869	1.21×10^{-3}	0.494
27.943	13.444	1.43×10^{-3}	0.516
33.252	14.019	1.70×10^{-3}	0.538
37.821	14.249	1.93×10^{-3}	0.547
43.998	14.594	2.25×10^{-3}	0.560
59.211	15.308	3.03×10^{-3}	0.588
80.586	16.067	4.12×10^{-3}	0.617
93.234	16.481	4.77×10^{-3}	0.633
104.904	16.941	5.36×10^{-3}	0.651
126.249	17.585	6.45×10^{-3}	0.675
152.771	18.160	7.81×10^{-3}	0.697
182.850	18.873	9.35×10^{-3}	0.725
222.571	19.655	1.14×10^{-2}	0.755
246.195	19.908	1.26×10^{-2}	0.764
267.784	20.115	1.37×10^{-2}	0.772
302.877	20.322	1.55×10^{-2}	0.780
354.168	20.391	1.81×10^{-2}	0.783
481.483	20.391	2.46×10^{-2}	0.783
719.743	20.391	3.68×10^{-2}	0.783
956.928	20.437	4.89×10^{-2}	0.785

STAR-CCM+ Implementation

The profile was implemented through the following steps:

1. The data from Table 4 was saved as a CSV file
2. Imported into STAR-CCM+ via `Tools` → `Table Data` → `New Table`

3. At the inlet boundary:
    - Velocity specification set to `Profile`
    - Table assigned to `Velocity Magnitude`
    - Direction vector: $(1, 0, 0)$
4. Turbulence quantities initialized as:

$$k = 1.5(IU_0)^2 = 6 \times 10^{-4}\,\mathrm{m^2\,s^{-2}}, \quad \omega = \frac{\sqrt{k}}{0.3L_t} = 0.97\,\mathrm{s^{-1}} \tag{91}$$

   where $I = 0.02$ and $L_t = 0.084\,\mathrm{m}$ $(0.07\delta_{99})$