

HORIZON-EUROHPC-JU-2021-COE-01




The European Centre of Excellence for Engineering
Applications

Project Number: 101092621

D2.15

Updated Report on the STREAMS Application Use Case



EuroHPC
 Joint Undertaking

The EXCELLERAT P2 project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101092621. The JU receives support from the European Union’s Horizon Europe research and innovation programme and Germany, Italy, Slovenia, Spain, Sweden and France.

Work Package:	2	Use-Case Execution
Author(s):	Sergio Pirozzoli	URMLS
	Giulio Soldati	URMLS
	Francesco Salvatore	CINECA
Approved by	Executive Centre Management	2.06.2025
Reviewer	Gregor Weiss	HLRS
Reviewer	Tuan Anh Dao	KTH
Dissemination Level	Public	

Date	Author	Comments	Version	Status
15.05.2025	Giulio Soldati, Sergio Pirozzoli	Full document drafting	V0.1	Draft
19.05.2025	Francesco Salvatore	Document revision and finalisation	V1.0	Final
22.05.2025	Francesco Salvatore	Implementation of first internal review	V1.1	Final
30.05.2025	Francesco Salvatore	Document with final revisions	V1.2	Final

List of abbreviations

CFD	Computational Fluid Dynamics
CPU	Central Processing Unit
DNS	Direct Numerical Simulation
DR	Drag Reduction
GPU	Graphical Processing Unit
HIP	Heterogeneous-computing Interface for Portability
HPC	High-Performance Computing
MPI	Message Passing Interface
Q	Quarter
RANS	Reynolds-Averaged Navier-Stokes
RMS	Root Mean Square
StTW	Streamwise Travelling Waves
UC	Use Case
VTK	Visualization ToolKit
WP	Work Package

Executive Summary

This document summarises the progress made since M12 in the STREAMS use case (UC-6) on active flow control for drag reduction (DR) of transonic airfoils. During this period, the previous FLEW solver was refactored and merged into the modular, object-oriented STREAMS-2 code leading to STREAMS-2.1 version. The porting to multi-GPU architectures has been completed to enable large-scale Direct Numerical Simulation (DNS) on heterogeneous architectures. We considered and evaluated both NVIDIA and AMD Graphics Processing Unit (GPUs) from EuroHPC clusters. In addition, code suitable for Intel GPUs was also prepared and tested. Early runs exposed the need to improve grid generation, DNS initialisation, and far-field boundary treatments to suppress spurious acoustic reflections; these enhancements were implemented and validated. Building on Quadrio et al. [4], we completed a comprehensive parametric study of the proposed control technique at low Reynolds number ($Re = 3 \times 10^5$) comprising 26 controlled flow cases, as reported in a recently published paper [8]. Preliminary uncontrolled simulations have been conducted at higher Reynolds number ($Re = 6 \times 10^5$), and controlled cases are also planned. Finally, we are preparing an ambitious DNS at $Re = 1.2 \times 10^6$, which would result - to our knowledge - in the highest Reynolds number simulated by a DNS of a controlled airfoil. We aim to target at least one uncontrolled and one controlled run, with the possibility of additional exploratory cases. Over 1.1 M node-hours have been obtained through EuroHPC and other calls, so that UC-6 is positioned to deliver its high-Reynolds objectives.

Table of Contents

1	Use Case Introduction	7
2	Objectives of the Use Case.....	8
3	Workflow Description.....	8
4	Progress achieved since M12	10
5	Next Steps	12
6	References	13

Table of Figures

Figure 1: Wall shear stress on the suction side of the V2C airfoil. Applying StTW control (bottom panel) reduces near-wall turbulence intensity and shifts the shock foot downstream (red line).....	7
Figure 2: Schematic of the automated workflow. Each box represents a SLURM job; arrows denote de-pendencies managed by the Python driver.	9
Figure 3: Instantaneous numerical Schlieren for turbulent flow at $Re = 3 \times 10^5$, $M = 0.7$, $\alpha = 4^\circ$ over the V2C airfoil.....	10
Figure 4: Friction coefficient (top panel) and pressure coefficient (bottom panel) for turbulent flow at $Re = 3 \times 10^5$, $M = 0.7$, $\alpha = 4^\circ$ over the V2C airfoil. Black lines and dots refer to the baseline condition, red line and dots refer to the controlled condition using the same forcing parameter as in Quadrio et al. [4].	11

1 Use Case Introduction

Environmental concerns over CO₂ emissions have made aerodynamic efficiency a central goal in civil aviation. The International Civil Aviation Organization projects estimates that aviation’s carbon footprint could triple by 2050, consuming a quarter of the world’s remaining carbon budget [1]. In cruise flight, fuel consumption is nearly proportional to total drag about half of which arises from skin-friction drag on the wing. Consequently, reducing turbulent skin friction promises significant environmental and economic benefits. However, active drag-reduction techniques must overcome the challenges posed by high-Reynolds-number, compressible, shock-laden flows typical of modern transonic transport.

Active flow-control methods can yield larger net savings than passive devices—such as riblets—provided that the actuation energy remains modest. Among these, spanwise wall forcing via streamwise-travelling waves (StTW) of spanwise velocity has demonstrated particular promise. In this approach, the wall velocity is prescribed as

$$w_w(x, t) = A \sin(\kappa_x x - \omega t),$$

where A is the forcing amplitude, κ_x the streamwise wavenumber, and ω the angular frequency. The resulting Stokes layer interacts with near-wall turbulence structures, disrupting the regeneration cycle of momentum streaks [5]. Early DNS on flat-plate flows reported up to 30% net skin-friction reduction [2], and subsequent experimental and numerical studies have confirmed its efficacy at Reynolds numbers relevant to aeronautical applications [3]. Recent DNS of a supercritical airfoil section ($Re = 3 \times 10^5$, $M_\infty = 0.7$, $\alpha = 4^\circ$) applied this control over a finite suction-side segment and observed not only decreased skin friction but also increased lift. The lift gain was traced to a downstream displacement of the λ -shock foot (see Figure 1), raising the local Mach number on the suction side [4]. Extrapolated to a full-aircraft context, these results imply potential cruise-drag savings on the order of 9% for a power penalty of approximately 1%.




Figure 1: Wall shear stress on the suction side of the V2C airfoil. Applying StTW control (bottom panel) reduces near-wall turbulence intensity and shifts the shock foot downstream (red line)

2 Objectives of the Use Case

Previous studies of streamwise-travelling waves of spanwise wall velocity (StTWs) have demonstrated significant skin-friction reductions in canonical, low-Reynolds-number flows—most prominently on flat plates and in simple duct geometries—but their impact on the complex aerodynamics of a transport-relevant wing remains insufficiently quantified. In transonic airfoils, curved surfaces, non-uniform pressure gradients and shock–boundary-layer interactions all contribute to drag and may interact nonlinearly with active control. This use case therefore targets a systematic assessment of StTW actuation on a supercritical airfoil under conditions representative of cruise flight, with three tiers of investigation:

1. **Baseline parameter exploration at low Reynolds number ($Re = 3 \times 10^5$).** We have completed a comprehensive sweep of the control parameters—oscillation frequency ω , streamwise wavenumber κ_x and forcing amplitude A —at chord-based $Re = 3 \times 10^5$, $M_\infty = 0.7$ and $\alpha = 4^\circ$. 28 DNS cases (2 uncontrolled and 26 controlled) were conducted using a previous DNS code [8]. These results establish a detailed map of the ω – κ_x – A space and demonstrate both friction reduction and lift enhancement via downstream shock displacement. A reference case at the same Reynolds number has been simulated using the new application code (STREAmS-2.1) proving excellent agreement of results while delivering a significant improvement under the High-Performance Computing (HPC) perspective.
2. **Moderate-Reynolds number evaluation at $Re = 6 \times 10^5$.** Building on insights from the DNS campaign at $Re = 3 \times 10^5$, we executed preliminary, uncontrolled DNS at $Re = 6 \times 10^5$ to further assess the accuracy and performance of the newly developed solver STREAmS-2 as well as mesh adequacy. Based on these baselines, we plan to perform two to three controlled StTW simulations, selecting parameter sets that showed the greatest net benefit at $Re = 3 \times 10^5$ (after suitable scaling of the parameters).
3. **High-Reynolds demonstration at $Re = 1.2 \times 10^6$.** To push the boundaries of DNS in realistic transport regimes, we will undertake the first simulations—one uncontrolled and one StTW-controlled—at chord-based Reynolds number $Re = 1.2 \times 10^6$, each on a mesh of approximately 2×10^{10} cells. These runs represent the largest DNS of active control to date and will critically test the scale-up of both performance gains and computational infrastructure. Subject to available resources, up to two additional exploratory cases may be added.

These stages aim to identify robust actuation settings that deliver net drag savings without compromising lift, while elucidating the underlying physics of shock–boundary-layer–control interactions across a realistic range of flight conditions.

3 Workflow Description

Preparation and execution of simulations presents challenges from both a computational and configuration perspective. To minimise resource management challenges on the user side, STREAmS has preconfigured makefile specifications on a variety of HPC clusters, particularly from the EuroHPC context. As for the simulation configuration, we have developed a Python workflow management tool capable of automating the steps required to prepare the computational grid, initial and boundary conditions. Workflow automation integrated into the HPC

environment represents a significant element towards the usability of the code, especially in view of parametric simulation sequences. The main workflow steps are explained below.

Once provided with three inputs—the airfoil surface coordinates, a configuration file specifying flow conditions (Re , M_∞ , α) and control parameters (A , κ_x , ω), and a template for SLURM resource allocation—the driver proceeds as follows. First, it invokes our patched Construct2D module to generate a curvilinear mesh tuned for a preliminary 2D Reynolds-Averaged Navier-Stokes (RANS) simulation. Upon convergence (typically within a few hours), the driver automatically analyses flow statistics, computing local boundary-layer thickness, viscous length, and Kolmogorov length in the wake region. These turbulence metrics are fed back into the mesh template to guide DNS grid generation.

Next, the workflow transitions to DNS preparation. The enhanced version of Construct2D produces a body-fitted C-grid around the airfoil, with wall-normal and wake spacings prescribed by the RANS-derived scales. The driver then interpolates the converged RANS solution onto this refined grid, ensuring a physically consistent initial field that mitigates spurious transients. A SLURM job is submitted for STREAMS-2—leveraging either CUDA or HIP GPU backends as configured by the user—using precise dependency flags so that the DNS run commences only once the interpolation stage completes. During the DNS, Catalyst2 in-situ routines can be activated, which export Visualization ToolKit (VTK) snapshots of shock location, boundary-layer development, and spanwise wave penetration, allowing early quality checks.

Figure 2 illustrates the end-to-end workflow, with each box representing an automated stage—from geometry input and RANS mesh to DNS execution and results aggregation—and arrows denoting the SLURM dependencies enforced by the Python driver. This automated pipeline stands ready to support the high- Re frontier simulations, ensuring a minimum of manual overhead across various EuroHPC platforms.

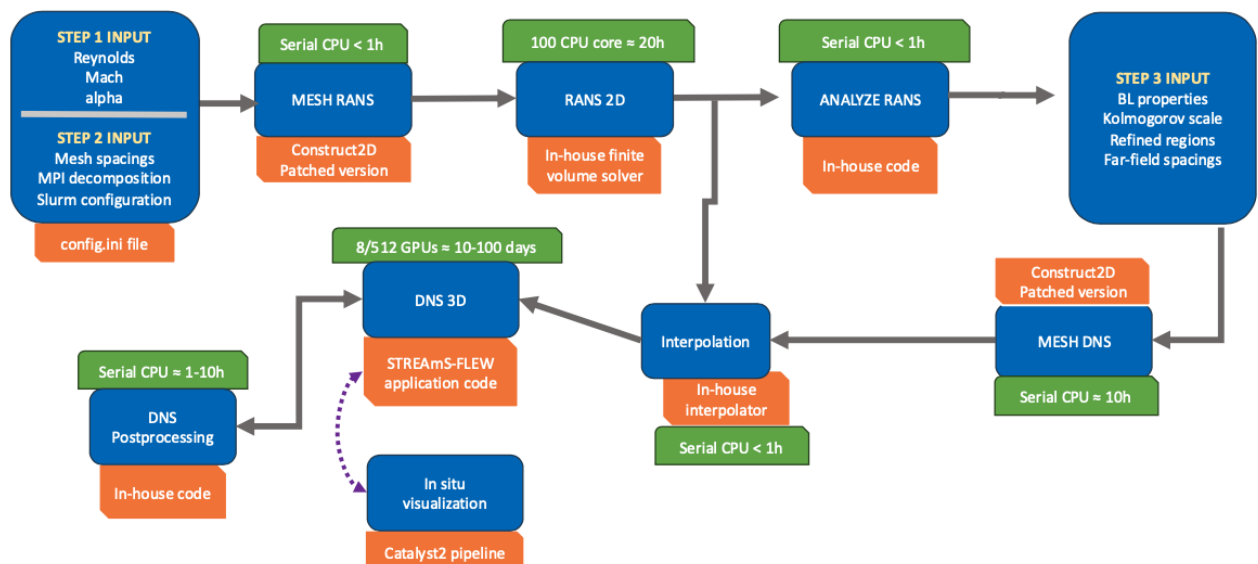


Figure 2: Schematic of the automated workflow. Each box represents a SLURM job; arrows denote dependencies managed by the Python driver

4 Progress achieved since M12

Significant advances have been made on the three interlinked fronts of solver capability, workflow automation, and DNS execution. The previous FLEW code was refactored into the modular, object-oriented STREAMS-2 framework, enabling straightforward integration of multiple computational backends. In collaboration with Work Package (WP) 3, both CUDA (NVIDIA) and HIP (AMD) implementations were validated via preliminary DNS at $Re = 3 \times 10^5$. Additional computing backends were tested, i.e., based on CPUs and on Intel GPUs. The effort required for porting, testing, and maintaining the code across different computational backends is substantial, but the results demonstrate the solver's ability to achieve performance levels that reasonably reproduce the peak capabilities of diverse devices. This enables STREAMS to avoid being locked into a single vendor's ecosystem, allowing it to adapt to the most favourable market conditions and the availability of HPC systems across Europe and beyond. More details on code development are provided in Deliverable D3.2. Besides, in situ capabilities were substantially improved as detailed in D4.2. The results of simulations exposed the need for improved initialisation and boundary treatment: we now interpolate converged RANS solutions onto the DNS grid, enforce far-field relaxation zones to the RANS state, and embed a thin sponge layer to absorb acoustic disturbances and avoid spurious wave reflections.

A Python-driven workflow has been developed to drive every stage—mesh generation, preliminary RANS computation, flow statistics pre-processing, and DNS preparation and computation—with minimal manual effort. A bespoke routine reads RANS-derived viscous thickness and Kolmogorov length scales to prescribe wall-normal and wake-region spacings, and grid-generator code was updated to eliminate metric discontinuities and realign the finest C-mesh region with the actual flow wake. These refinements ensure consistent mesh quality and solver performance across all cases.

A comprehensive parametric study of the control parameter space has been completed at $Re = 3 \times 10^5$ [8]. A baseline simulation without control is ongoing at moderate Reynolds number ($Re = 6 \times 10^5$), and two to three spanwise-forcing cases are slated to validate drag-reduction mechanisms. Looking ahead, we are preparing the first DNS at $Re = 1.2 \times 10^6$, targeting at least one uncontrolled and one StTW-controlled run on $\sim 2 \times 10^{10}$ cell grids. These efforts confirm both the high-fidelity of STREAMS-2 and the robustness of our automated pipeline, and position UC-6 to deliver its final high-Reynolds milestones.

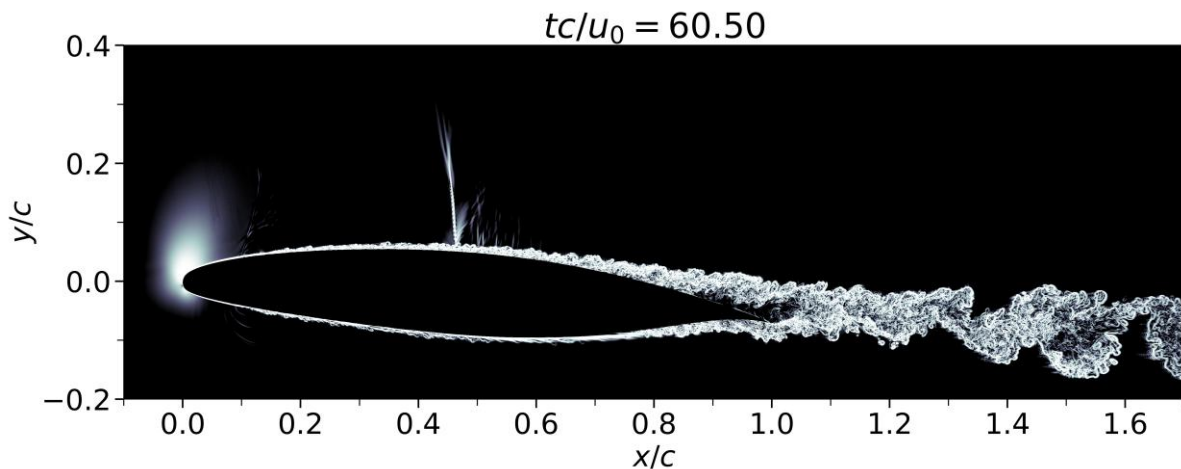


Figure 3: Instantaneous numerical Schlieren for turbulent flow at $Re = 3 \times 10^5$, $M = 0.7$, $\alpha = 4^\circ$ over the V2C airfoil

To illustrate part of our progress, Figure 3 presents an instantaneous numerical Schlieren field from a $Re = 3 \times 10^5$ validation run, capturing fine turbulence structures downstream of the airfoil. Figure 4 compares our computed friction coefficient (c_f) and pressure coefficient (c_p) distributions with the reference data of Quadrio et al. [4], showing excellent agreement both in the baseline (black) and controlled (red) case, confirming the accuracy of STREAMS-2. Furthermore, various tests and validation cases on different HPC clusters have demonstrated STREAMS-2's readiness to run on different cluster architectures. The successes achieved were promptly disseminated, as demonstrated by the publications in peer-reviewed journals [6, 7, 8, 9].

In preparation for the next phase of the project requiring massive use of computational power, we participate as team members to the EuroHPC Extreme project FASTER - Flow control on Airfoils through Streamwise Traveling waves for Efficiency Rise (EHPC-EXT-2024E02-130). The project has been awarded 1 million node hours on LUMI-G. The allocation started on Monday, 7th of April 2025 and will end on Monday, 6th of April 2026. This project will then provide the computational resources for the most demanding simulations planned in UC-6.

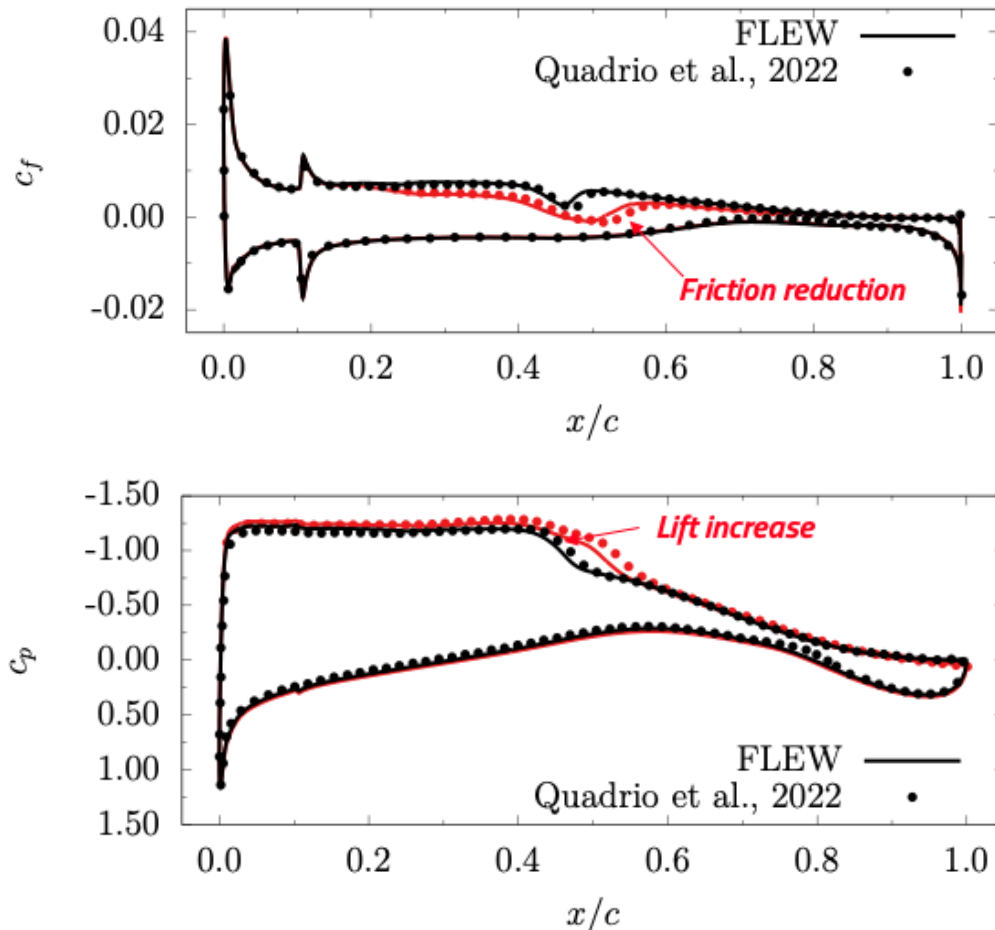


Figure 4: Friction coefficient (top panel) and pressure coefficient (bottom panel) for turbulent flow at $Re = 3 \times 10^5$, $M = 0.7$, $\alpha = 4^\circ$ over the V2C airfoil. Black lines and dots refer to the baseline condition, red line and dots refer to the controlled condition using the same forcing parameter as in Quadrio et al. [4]

5 Next Steps

In the coming period, we will concentrate on completing the remaining simulation tiers, finalising performance validation, and conducting the full analysis of the accumulated data.

First, the moderate-Reynolds study will be concluded by performing two to three additional StTW-controlled DNS at $Re = 6 \times 10^5$. These cases will be selected based on the drag-reduction trends observed in the completed baseline runs and the 26-case campaign at $Re = 3 \times 10^5$. The resulting data will possibly validate the persistence of control benefits under increased turbulence levels.

Second, we will execute the initial high-Reynolds simulations at $Re = 1.2 \times 10^6$. At minimum, one uncontrolled and one controlled DNS will be performed on meshes of approximately 2×10^{10} cells. We plan to complete both simulations by the end of 2025, with a preliminary estimate of completion in the third quarter (Q3) for the uncontrolled case and in the fourth quarter (Q4) for the controlled case. Subject to available node-hours, we may add one or two exploratory runs to probe adjacent regions of the control parameter space. These simulations represent the largest active-control DNS to date and will test both physical performance and computational robustness.

Concurrently, cross-platform tests will compare solver throughput and numerical consistency across NVIDIA and AMD GPUs—and, where relevant, CPU-only runs. This validation will ensure that STREAMS-2 operates predictably on the target EuroHPC systems and guide any final adjustments to Message Passing Interface (MPI) decomposition and I/O strategies.

Once the full set of DNS is complete, post-processing scripts will extract mean, root mean square (RMS), spectral statistics, and assemble aerodynamic coefficients into the project database. A systematic analysis will then map the control parameters (ω , κ_x , A) and flight conditions (Re , M_∞ , α) to performance outcomes, identifying robust operating points that maximise net drag savings without compromising lift.

A public release of the STREAMS-2 workflow repository, complete with tutorials and sample scripts has been recently released and is available on GitHub (<https://github.com/STREAMS-CFD/STREAMS-2>). Finally, our findings will be disseminated through peer-reviewed publications.

6 References

- [1] Graver, B., Zhang, K., & Rutherford, D. (2019) CO2 emissions from commercial aviation, 2018, *Tech. Rep. International Council on Clean Transportation*.
- [2] Quadrio, M., Ricco, P., & Viotti, C. (2009). Streamwise-travelling waves of spanwise wall velocity for turbulent drag reduction. *Journal of Fluid Mechanics*, 627, 161-178.
- [3] Gatti, D., & Quadrio, M. (2016). Reynolds-number dependence of turbulent skin-friction drag reduction induced by spanwise forcing. *Journal of Fluid Mechanics*, 802, 553-582.
- [4] Quadrio, M., Chiarini, A., Banchetti, J., Gatti, D., Memmolo, A., & Pirozzoli, S. (2022). Drag reduction on a transonic airfoil. *Journal of Fluid Mechanics*, 942:R2.
- [5] Quadrio, M., & Ricco, P. (2011) The laminar generalized Stokes layer and turbulent drag reduction. *Journal of Fluid Mechanics* 667, 135–157.
- [6] Soldati, G., Ceci, A., & Pirozzoli, S. (2024). FLEW: A DNS Solver for Compressible Flows in Generalized Curvilinear Coordinates. *Aerotecnica Missili & Spazio*, 1-13.
- [7] Salvatore, F., Soldati, G., Ceci, A., Rossi, G., Memmolo, A., Della Posta, G., ... & Pirozzoli, S. (2025). STREAMS-2.1: Supersonic turbulent accelerated Navier-Stokes solver version 2.1. *Computer Physics Communications*, 109652.
- [8] Berizzi, N., Gatti, D., Soldati, G., Pirozzoli, S., & Quadrio, M., (2025). Aerodynamic performance of a transonic airfoil with spanwise forcing. *Journal of Fluid Mechanics*. 1010:A18.
- [9] Sathyanarayana, S., Bernardini, M., Modesti, D., Pirozzoli, S., & Salvatore, F. (2025). High-speed turbulent flows towards the exascale: STREAMS-2 porting and performance. *Journal of Parallel and Distributed Computing*, 196, 104993.