

Multi-Partner Project: Multi-GPU Performance Portability Analysis for CFD Simulations at Scale

Panagiotis-Eleftherios Eleftherakis*, George Anagnostopoulos*, Anastassis Kapetanakis*, Mohammad Umair†, Jean-Yves Vet‡, Konstantinos Iliakis*, Jonathan Vincent†, Jing Gong†, Akshay Patil§, Clara García-Sánchez§, Gerardo Zampino†, Ricardo Vinuesa¶, Sotirios Xydis*

*National Technical University of Athens, Greece

†KTH Royal Institute of Technology, Sweden

‡Hewlett Packard Enterprise (HPE), France

§Technical University of Delft, Netherlands,

¶University of Michigan, USA

Abstract—As heterogeneous supercomputing architectures leveraging GPUs become increasingly central to high-performance computing (HPC), it is crucial for computational fluid dynamics (CFD) simulations, a de-facto HPC workload, to efficiently utilize such hardware. One of the key challenges of HPC codes is performance portability, i.e. the ability to maintain near-optimal performance across different accelerators. In the context of the REFMAP project, which targets scalable, GPU-enabled multi-fidelity CFD for urban airflow prediction, this paper analyzes the performance portability of SOD2D, a state-of-the-art Spectral Elements simulation framework across AMD and NVIDIA GPU architectures. We first discuss the physical and numerical models underlying SOD2D, highlighting its computational hotspots. Then, we examine its performance and scalability in a multi-level manner, i.e. defining and characterizing an extensive full-stack design space spanning across application, software and hardware infrastructure related parameters. Single-GPU performance characterization across server-grade NVIDIA and AMD GPU architectures and vendor-specific compiler stacks, show the potential as well as the diverse effect of memory access optimizations, i.e. $0.69\times$ - $3.91\times$ deviations in acceleration speedup. Performance variability of SOD2D at scale is further examined on the LUMI multi-GPU cluster, where profiling reveals similar throughput variations, highlighting the limits of performance projections and the need for multi-level, informed tuning.

Index Terms—Performance portability, CFD, Spectral Finite Element Method (FEM), high-fidelity simulation, multi-GPU acceleration, design space exploration, scalability analysis

I. INTRODUCTION

Computational Fluid Dynamics (CFD) is an essential tool for modern engineering, enabling detailed analysis of fluid behavior in aerospace, automotive, maritime, and environmental systems without costly physical prototypes. Recently, CFD-generated datasets have also become central to training physics-informed neural models, which embed physical laws into machine learning models for improved predictive accuracy [1], [2]. At its core, CFD solves the Navier–Stokes equations, expressing conservation of mass, momentum, and energy in fluid flows. Analytical solutions are rare, so numerical methods such

as the Finite Difference Method (FDM) and Finite Element Method (FEM) are employed. FDM is efficient for structured grids, while FEM better handles complex geometries at higher computational cost.

High-Performance Computing (HPC) has greatly advanced CFD, enabling high-fidelity simulations of complex phenomena. Heterogeneous architectures, particularly those leveraging Graphics Processing Units (GPUs), have driven progress toward exascale computing, supported by sophisticated solvers such as OpenFOAM, NekRS, and others [3], [4]. However, optimizing CFD solvers for diverse hardware remains challenging due to differing GPU architectures, compiler ecosystems, and programming models. Solvers optimized for one platform often underperform on another, hindering performance portability—i.e., sustained efficiency across architectures.

Within this context, the EU funded REFMAP project (*Reducing Environmental Footprint through transformative Multi-scale Aviation Planning*) targets low-latency, multi-fidelity urban-airflow prediction to support sustainable Innovative Air Mobility in urban environments. Its high-fidelity CFD component couples GPU-enabled simulations with machine-learning surrogates. High-resolution Direct Numerical Simulation (DNS) data train compact predictors deployable at the edge or in the cloud. Since such workloads are computationally intensive, REFMAP integrates GPU-accelerated CFD and autotuning to overcome data-generation bottlenecks. The Spectral Element solver SOD2D [5] is ported to GPUs using OpenACC [6] and MPI, enabling scalable, turbulence-resolving simulations across NVIDIA and AMD systems. Achieving performance portability across these heterogeneous HPC platforms is critical to REFMAP’s data-driven workflow, ensuring scalable, cost-effective surrogate training and deployment.

CFD solvers adopting high-order numerical schemes and multi-architecture support—such as URANOS [7], STREAMS [8], and PyFR [9]—have achieved significant GPU speedups, though cross-vendor analysis remains limited. Reguly and Mudalige [10] provide a broad overview of performance portability strategies, noting that few works (e.g., CloverLeaf [11], TeaLeaf [12], Nekbone [13]) use OpenACC across multiple platforms, with limited direct cross-vendor compar-

This work has been partially funded by the EU Horizon 2022 program under grant agreement No 101096698 REFMAP (website: <https://www.refmap.eu>). We acknowledge the EuroHPC Joint Undertaking for awarding this project access to the EuroHPC supercomputer LUMI, hosted by CSC (Finland) and the LUMI consortium through a EuroHPC Regular Access call.

ison. More recent efforts [14] similarly omit direct cross-vendor GPU evaluation. To our knowledge, this work is among the first to perform a detailed, full-stack cross-vendor study of GPU-accelerated CFD using OpenACC, spanning from solver configuration to hardware-level memory optimization.

II. THE REFMAP PROJECT OUTLINE

REFMAP targets low-latency, multi-fidelity urban-airflow prediction to support sustainable Innovative Air Mobility. As seen in Fig. 1, the low-fidelity branch is driven by the OpenFOAM CFD simulator (and accelerated by the NVIDIA AmgX solver) executing a model that aims at evaluating the neighbourhood-scale wind-associated risk [15]. Regarding the high-fidelity framework, the SOD2D simulator [5] is employed (a Fortran-based SEM code with OpenACC offload and MPI partitioning, where key kernels —i.e. convection, diffusion— are parallelized across elements and nodes) enabling the coupling of GPU-enabled CFD with data-driven surrogates. High-resolution LES/DNS simulations generate turbulence-resolving data to train compact models [16] that undergo extensive optimizations such as compression and pruning in order to deploy at the edge or in the cloud, enabling UAV trajectory optimization and high-fidelity sensor placement, thus providing a robust substrate for efficient and safe UAV navigation. On CPU-only systems, such cases require weeks to compute, limiting the datasets for surrogate training. REFMAP overcomes this bottleneck by extensive autotuning and performance exploration of SOD2D. Initial studies have shown performance sensitivity when freely varying gang numbers and vector lengths using OpenTuner’s AUC-Bandit search technique, obtaining 1.2–1.5× speedup [17], [18]. Moving towards REFMAP’s mature phase, multiple infrastructures were employed necessitating multi-level performance exploration [19] - i.e. spanning across hardware, software and application-specific parameters. *On the surrogate model side*, REFMAP employs fully convolutional networks (FCNs) that map wall-shear-stress and pressure inputs to velocity-fluctuation fields $\hat{u}, \hat{v}, \hat{w}$ at selected wall-normal positions. Models are trained on DNS data ($Re_\tau=180$) with 6,000 snapshots per $y^+ = 15, 30, 50, 100$, using MSE loss and normalized targets to balance component errors. While additional GAN and transformer models are explored, FCNs remain the main surrogate coupled with accelerated CFD for wall-based flow reconstruction. Leveraging structured pruning, quantization, and compiler-level optimizations, the final turbulence-prediction models achieve up to 50× compression with negligible loss in accuracy, significantly improving latency and throughput. These advances enable fast, reliable flow prediction and establish a foundation for integrating high-performance simulation and data-driven surrogates into real-time UAV navigation and urban risk assessment workflows. The rest of this paper will focus on the performance portability analysis that enables efficient generation of datasets for surrogate model training within the REFMAP workflow.

III. SOD2D: SPECTRAL FEM FOR COMPRESSIBLE SCALE-RESOLVING SIMULATIONS

A. Implementation and governing equations

The physical model considered is composed of the Navier–Stokes equations governing mass, momentum, and energy conservation in a compressible fluid [5]. In conservative form:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad (1)$$

$$\frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \otimes \mathbf{u}) - \nabla \cdot \boldsymbol{\tau} = -\nabla p + \mathbf{f} \quad (2)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot ((E + p)\mathbf{u}) - \nabla \cdot (\boldsymbol{\tau} \mathbf{u}) - \nabla \cdot (\kappa \nabla T) = S \quad (3)$$

Here ρ , \mathbf{u} , p , T , and E denote density, velocity, pressure, temperature, and total energy. The system is closed with the ideal gas law. The colored terms highlight **Temporal Evolution**, **Convection**, and **Diffusion**. The equations are solved by separating temporal and spatial discretization. Spatial terms use the Finite Element Method (FEM) and temporal integration uses a fourth-order Runge–Kutta (RK4) scheme.

SOD2D [5] is an OpenACC-based CFD software leveraging MPI process partitioning techniques, thus targeting multiple platforms. It focuses on a variant of the Finite Element Method (FEM) known as the Spectral Element Method (SEM) and can handle unstructured meshes, thus modeling complex geometries. This method utilizes non-uniformly distributed nodes within each element while lowering the computational cost for high mesh resolutions. Additionally, the Gauss-Lobatto-Legendre quadrature rule is employed aligning the quadrature points with the element nodes in the discretized mesh.

The aforementioned design choices of SOD2D reflect its primary objective: to create a hardware-optimized CFD simulator capable of maximizing parallelism across multi-core CPUs as well as single- and multi-GPU architectures [5], [17]. The mathematical manipulation of the PDEs within the SEM framework results in a computationally efficient diagonal linear system. This structure eliminates the reliance on complex linear solvers, which are often a major bottleneck in traditional CFD applications [10].

B. SOD2D execution flow

Fig. 2 illustrates the execution flow of SOD2D’s various computational kernels, represented by the blue rectangles. The algorithm is structured around a loop over four Runge-Kutta steps, as described in [5]. Within each step, the necessary values for evaluating the Runge-Kutta (RK) function \mathbf{f} are computed. This is followed by the evaluation of the Diffusive and Convective term residuals at each point n within every mesh element e . Subsequently, the residuals of these two terms are accumulated at each node and the quantities’ values are updated at the current RK step. Finally, entropy viscosity is computed, a numerical technique used to stabilize fluid flow simulations, particularly in shock-resolving scenarios

The focus lies on the computation of the Diffusive and Convective terms, a process divided into three subparts: Preloading,

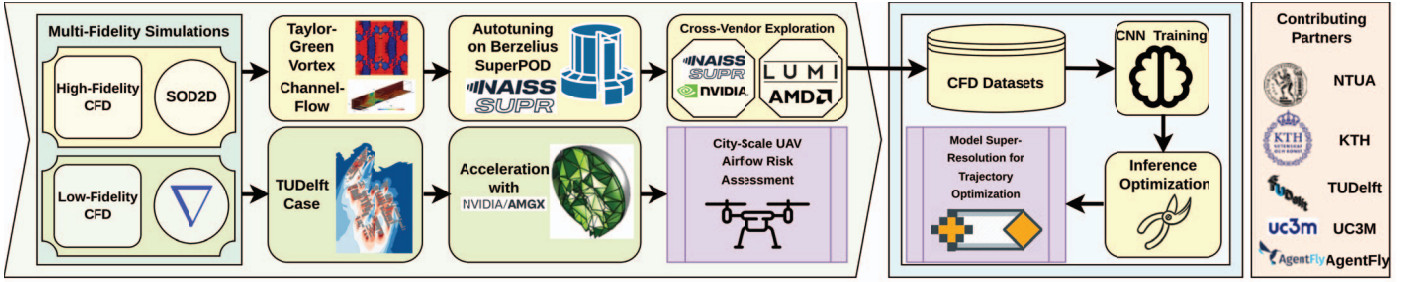


Fig. 1. Overview of REFMAP's HPC to ML workflow, with emphasis on CFD acceleration.

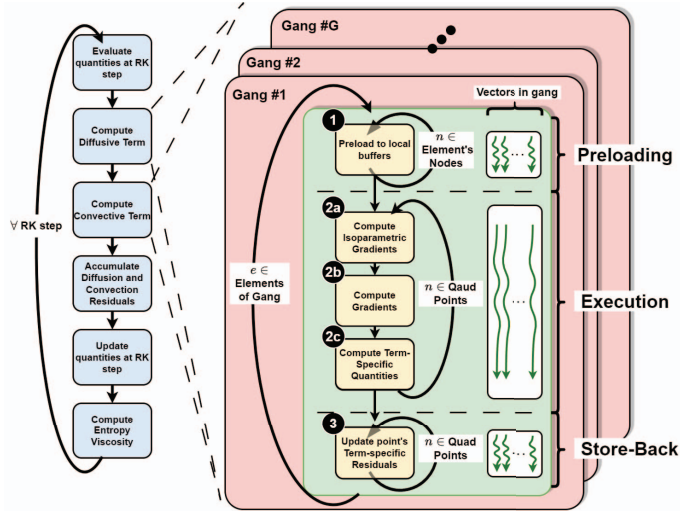


Fig. 2. High-level algorithmic overview of SOD2D

Execution and Store-Back. These subparts are executed independently, with synchronization barriers indicated by dashed lines. Each term computation iterates over all elements in the mesh. This element loop is distributed across fully parallel gangs, and vectors within gangs, as defined in [6], each with sizes defined by the kernel's execution parameters. In step 1, global quantities corresponding to the nodes of the current element are fetched into local arrays, thus the main execution part begins. The first operation involves calculating the isoparametric gradients 2a, which are gradients in the natural coordinate system. Those are then used to compute the gradients in the global coordinate system 2b by multiplying them with the inverse *Jacobian*. In the next step 2c, additional term-specific quantities are computed. This process is repeated for all quadrature points, which in the SEM framework coincide with the nodal points whose values need to be computed. The final step of the computation involves calculating each point's residual contribution for the current term 3. Since this step updates the global residual array, it encompasses multiple atomic updates to ensure proper memory synchronization, which serializes memory accesses and updates.

We consider analysis on two cases for turbulence simulations: a TGV (Taylor-Green Vortex) and a CF (Channel Flow) case configuration. Performance profiling reveals consistent hotspot behavior across both simulated cases, with

the `full_convect` (Convection) kernel dominating execution time -accounting for up to 50% (NVIDIA) and 70% (AMD) of total runtime— followed by the `full_diffusion` (Diffusion) kernel at 22% and 10% respectively. Given the strong similarity in hotspot distribution, subsequent analysis focuses on the Channel Flow (CF) case as more representative to airflow in urban environments, which describes the flow of fluid confined between two parallel surfaces/walls, typically forming a channel. The Convection kernel remains the primary computational bottleneck, showing notable sensitivity to floating-point precision and GPU vendor and mesh size, denoting the number of points in the unstructured mesh.

IV. MULTI-LEVEL DESIGN SPACE EXPLORATION

As illustrated in Fig. 3 we examine the application's behavior under all parameter space layers. We start from the configuration of the application's input parameters, extending down to the hardware infrastructure parameters, while also exploring code memory optimizations. The *Runtime Parameters layer* of our exploration framework considers single (FP32) and double (FP64) precision simulations. The Memory Access Optimizations section hierarchically fits between the Runtime Parameters and the Software Infrastructure space as it refers to core, manual code modifications elaborated on in subsection IV-A. The *Software Infrastructure layer* includes the compiler frameworks under examination, specifically NVHPC v24.5 for our local NVIDIA GPU and Cray Clang v17.0.1 for the LUMI cluster, utilizing AMD GPUs. With respect to the *Hardware Infrastructure layer*, a distinction between single and multi-GPU runs is made, as well as between different accelerator vendors.

Interestingly, initial profiling results demonstrated a $15.83\times$ slowdown of a single AMD MI250X Graphics Computing Die (GCD) compared to the NVIDIA Tesla V100 for 8M nodes, with single-precision arithmetic and without any code optimizations enabled. We attribute these disparities to four reasons: i) to the power cap on the GPUs of LUMI being just 400W compared to their nominal value of 500W, ii) to the use of a single GCD of the MI250X, iii) to the potentially unoptimized Cray compiler stack w.r.t. to the underlying hardware, iv) to the single-precision throughput of the MI250X faring much worse against that of the V100 compared to double-precision.

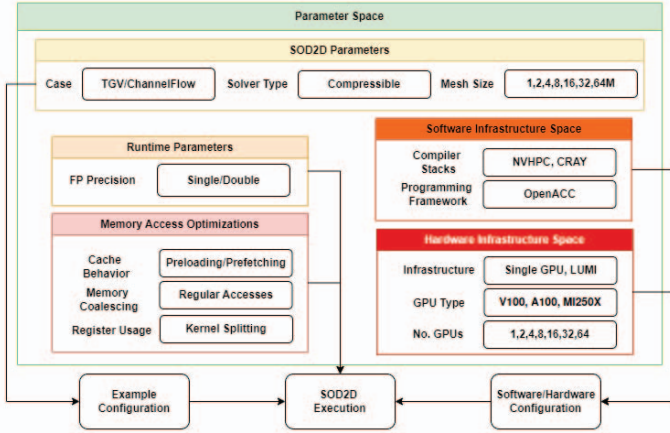


Fig. 3. Organization of the examined multi-level design space.

A. Memory Access Optimizations

Profiling results of the application on multi-GPU runs with Rocprofv3 and Omnitrace revealed maxed out local memory usage, along with register usage and inefficient coalescing of memory accesses in both shared and global memory. To this end, we parameterized the target kernels with memory access optimization strategies, applied with hand-crafted templates to the source code to alleviate each of the aforementioned issues. Details on the memory access optimizations follow:

Kernel Splitting: `full_convec`, the main hotspot of SOD2D, executes the steps ①–③ of Fig. 2 for Mass, Momentum and Energy equations concurrently, fetching global subarrays to per-element private buffers. The aforementioned steps are independent with regards to the partial results computed and the final output memory location. Due to the increased per-thread register pressure on the MI250X, *i.e.* *Average VGPR being equal to Maximum at 128, Average LDS Allocation being equal to Maximum at 4608* we opted to split the kernel into the three separate kernels corresponding to each equation. This provided an immediate $1.5\times$ speedup on `full_convec`.

Preloading: The preloading strategy is commonly employed in GPU-optimized codes to (1) improve data spatial locality, (2) potentially take advantage of shared memory within thread blocks, and (3) maximize memory bandwidth by coalescing memory requests to fetch multiple elements simultaneously.

The default version of SOD2D employs this technique [5], as illustrated in stage ① of Fig. 2. It is applied to global matrices that need to be accessed with an element index and are initially stored in an irregular manner. Since these matrices are in the Global Memory Space of the GPU and are accessed frequently in the Execution stage (②a–②c) of Fig. 2, they are first fetched to element-private local variables which removes the element node index, making memory accesses more regular and closer to computation.

Removing these preloaded variables and accessing global matrices directly provides moderate performance gains on many types of runs. We have attributed this effect to the sheer size of the local matrices, which leads to ineffective utilization of the lowest parts of the memory hierarchy.

Prefetching: Prefetching refers to the use of the `openacc`

cache directive at specific points in the `full_convec` subroutine. According to the OpenACC Standard 2.7, those directives are tied to loops located above or inside of and specify elements or ranges of subarrays that should be fetched into the highest level of the cache for the loop body.

Intuitively, the prefetching mechanism cannot be effective when applied to large global arrays with indirect access patterns. Range-based prefetching of global arrays was found to be ineffective. Therefore, our utilization of this directive is coupled with the preloaded local arrays described in subsection IV-A. Moreover, prefetching the local arrays for every quantity in the quadrature points loops is also ineffective. We have concluded that the optimal usage of this directive regards the quantities that participate in the deepest nested loops in the body of the quadrature loop. This provides a $1.24\times$ speedup with reference to just Preloading for the CF case with FP64, the split kernel and the V100 GPU.

Regular Accesses: A common way to improve GPU memory bandwidth is memory coalescing. Since sub-optimal coalescing ($6.5 \text{ sectors/request}$, $\max 16$) limits performance, we regularized global memory accesses in the “No Preloading” variant of SOD2D. In the quadrature loop (②a–②c in Fig. 2), accesses follow `rho(connec(ielem, invAtoIJK()))`, where `connec` is the mesh connectivity matrix and `invAtoIJK` an index linearization function. Such macros cause major performance penalties [8]. As `connec` and `invAtoIJK` depend only on static mesh topology, we precompute `connec(ielem, invAtoIJK())` and store it in a global matrix indexed by element node, quadrature point, and polynomial order. Subsequent kernel launches reuse it with perfect memory coalescing. Although `rho` accesses remain irregular, the costly linearization is removed. The matrices must fit in lower memory levels for efficiency.

V. EVALUATION

A. Experimental Setup

In this section, we present the results of our evaluation, providing insights on the behavior of SOD2D under different configurations of the parameter space. We examine the CF case with a mesh size of 1, 2, 4, 8, 16 million nodes with compressible solver. All input cases were run on both FP32 and FP64 modes of SOD2D. Regarding the SOD2D versioning, the application was compiled a total of 8 times. More specifically, the versions were SOD2D without preloading (**no_prl**), with preloading (**base**), preloading and prefetching (**prl_prf**) and regular memory accesses (**reg**), each of which were with (**split**) and without (**unified**) kernel splitting.

All cases and SOD2D versions were run on both NVIDIA and AMD systems. Specifically, a single NVIDIA Tesla V100-PCIE-32GB GPU in a local server as well as an NVIDIA A100 GPU from the Berzelius SuperPOD and a single Graphics Compute Die (GCD) of the AMD CDNA™ 2 MI250X-128GB courtesy of LUMI HPC cluster were utilized. Scalability experiments were run exclusively on LUMI.




Fig. 4. Single-GPU run times across all kernel optimizations for Channel Flow.

B. Single-GPU Exploration

Since near identical results are obtained between the TGV and CF cases in this section, we will focus on the more realistic CF case. The multi-level parameter space exploration and Cross-vendor single GPU comparative analysis reveals varying behavior of the main hotspot (`full_convvec`) when sweeping different parameters, highlighting the potential for the selection of locally optimal modifications.

As mentioned previously, the slowdown seen on a single GCD of the MI250X compared to the V100 is $15.83\times$ on 8M nodes and FP32 ($23.38\times$ for the A100). The respective slowdown value for the split kernel is $9.87\times$ and $8.66\times$ for the `prl_prf` version of SOD2D. For double-precision, the latter value is $5.23\times$. Regarding performance disparity between NVIDIA GPUs, the A100 provides a $1.47\times$ and a $3.05\times$ speedup over the V100 on FP32 and FP64 respectively.

Since AMD is consistently outperformed, the relative slowdowns compared to NVIDIA are not noted on Fig. 4. Instead, runtimes are evaluated across the same vendor and FP precision. Each bar corresponds to the average of 5 runs ranging from 1 to 16 million nodes, therefore the plot corresponds to a total of 240 runs of SOD2D.

We note that baseline SOD2D comes optimized with the Preloading technique detailed in Section IV-A. This process is expensive in terms of runtime since the program flow must go through the same loop bounds as **2a-2c** to preload global variables. Since inherent prefetching to caches in the `no_prl` version can be equally efficient, it is not obvious whether the benefit of Preloading compensates for the overhead of block **1**. As depicted in Fig. 4a, the `no_prl` version always outperforms the `base` (Preloading) version, with the exception of the split version with FP64 on the MI250X. This effect is greater for FP64 (DP) than FP32 (SP) and more prominent for the AMD MI250X compared to NVIDIA GPUs for FP32, while the opposite is true for FP64.

Preloading paired with Prefetching (`prl_prf`) consistently improves its performance (`base`). The effect is mostly minor, except for AMD MI250X on FP32, and NVIDIA V100 on FP64 regarding the split kernel. In the former case, a 42% and 16% speedup is provided for unified and split kernel respectively, while in the latter, a 63% and 25% speedup is observed,

rendering it the optimal code modification for AMD on FP32 and NVIDIA on FP64. We once again note the alternating effect of optimizations between NVIDIA-AMD and FP32-FP64.

Regular Accesses (`reg`) underperform on the different platforms due to the size of the regular matrix being $18\times$ that of `connec`, exceeding the capacity of the lowest parts of the memory hierarchy. Thus, irregular accesses to a smaller matrix are favorable to regular accesses to a larger one. However, `reg` is the only parameter that displays drastically different behavior across the cases run. Specifically for FP32, it fares better on the unified kernel. Once again, this behavior alternates on AMD. Such minor differences are consistent for FP64.

The impact of Kernel Splitting (`split`) is highly variable across vendors and floating point precisions as well, but mostly consistent across cases. For FP32, it speeds up the execution on AMD but incurs a slowdown for NVIDIA. Interestingly, the opposite behavior is displayed for FP64.

The behavior outlined above is consistent across all optimizations except Access Regularization. We attribute these differences to the differently sized Register Files for FP32 and FP64 across architectures.

C. Scalability Analysis

Fig. 5 represents data resulting from 160 weak scaling runs. For each configuration of the total number of GPUs, the size of the problem in terms of Mesh Nodes increases linearly, providing the same amount work per GPU ($\sim 1M$ Nodes).

In the context of scalability analysis of SOD2D, we evaluate efficiency as a metric of throughput. Thus, we introduce the *Giga Node Operations Per Second (GNOPS)* as:

$$\frac{\#TotalMeshNodes \times \#RungeKuttaSteps \times \#TotalTimeSteps}{SOD2DRuntime(ns)}$$

The *GNOPS* metric emerges from the homogeneous way the nodes are updated throughout the steps **1-3** in Fig. 2.

As observed in Fig. 5 that shows runs on FP32 (RP4), scaling is mostly consistent for all of the optimizations (FP64 is omitted as FP32 displays $1.22\times$ higher average GNOPS). However, we observe that CF demonstrates high inconsistency between scaling of different memory optimizations, with optimal configurations alternating between different scales (superlinear




Fig. 5. Throughput at scale for Channel Flow.

performance is observed due to partitioning inefficiencies at specific numbers of nodes [17]). We report that the Split kernel with Preloading and Prefetching (**prl_prf**) is the optimal across scales except for 32 GPUs (4 GPU nodes). We attribute this to redundant preloading operations at this degree of partitioning. For CF the improvements from worst to optimal configurations at maximum scale are 23.8% (FP32) and 14.6% (FP64) respectively. Therefore, as is apparent in Fig. 5, selecting an optimization for Channel Flow at 64 GPUs based on the optimum at 32 GPUs, would incur a 23.8% (FP32) and 8.3% (FP64) throughput penalty. Such disparities are absent in TGV, with 12.3% variation and consistent optimal configurations. The above analysis provides an important lesson for large scale accelerated compressible fluid dynamics, that *on a case-independent basis, performance prediction at different scales through simply extrapolating on low scale optimal memory access decisions is precarious, i.e. more sophisticated modeling is required at scale.*

VI. CONCLUSIONS

This paper explored the performance portability challenges in GPU-accelerated CFD simulations through a multi-level analysis of the SOD2D framework. We highlighted how architectural differences between GPU vendors significantly impact solver performance, revealing substantial disparities in scalability and performance portability, underscoring the need for parameter-space-informed optimizations. This study highlights the importance of multi-level exploration for developing performance-portable CFD frameworks, enabling efficient adaptation to heterogeneous HPC environments. Future work will extend this approach to new compiler and hardware infrastructure as well as other types of simulations and input meshes with an emphasis on further automating parameter optimization.

REFERENCES

- [1] R. Vinuesa and S. L. Brunton, "Enhancing computational fluid dynamics with machine learning," *Nature Computational Science*, vol. 2, no. 6, p. 358–366, Jun. 2022. [Online]. Available: <http://dx.doi.org/10.1038/s43588-022-00264-7>
- [2] H. Eivazi, M. Tahani, P. Schlatter, and R. Vinuesa, "Physics-informed neural networks for solving reynolds-averaged navier–stokes equations," *Physics of Fluids*, vol. 34, no. 7, Jul. 2022. [Online]. Available: <http://dx.doi.org/10.1063/5.0095270>
- [3] P. Fischer, S. Kerkemeier, M. Min, Y.-H. Lan, M. Phillips, T. Rathnayake, E. Merzari, A. Tomboulides, A. Karakus, N. Chalmers *et al.*, "Nekrs, a gpu-accelerated spectral element navier–stokes solver," *Parallel Computing*, vol. 114, p. 102982, 2022.
- [4] H. Jasak, "Openfoam: Open source cfd in research and industry," *International journal of naval architecture and ocean engineering*, vol. 1, no. 2, pp. 89–94, 2009.
- [5] L. Gasparino, F. Spiga, and O. Lehmkuhl, "Sod2d: A gpu-enabled spectral finite elements method for compressible scale-resolving simulations," *Computer Physics Communications*, vol. 297, p. 109067, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010465523004125>
- [6] O. Organization, *OpenACC Application Programming Interface Version 2.7*, November 2018, <https://www.openacc.org/sites/default/files/inline-files/OpenACC.2.7.pdf>. [Online]. Available: <https://www.openacc.org/sites/default/files/inline-files/OpenACC.2.7.pdf>
- [7] F. De Vanna and G. Balzan, "Uranos-2.0: Improved performance, enhanced portability, and model extension towards exascale computing of high-speed engineering flows," *Computer Physics Communications*, vol. 303, p. 109285, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S001046552400208X>
- [8] S. Sathyanarayana, M. Bernardini, D. Modesti, S. Pirozzoli, and F. Salvatore, "High-speed turbulent flows towards the exascale: Stream-2 porting and performance," *Journal of Parallel and Distributed Computing*, vol. 196, p. 104993, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0743731524001576>
- [9] F. Witherden, A. Farrington, and P. Vincent, "Pyfr: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach," *Computer Physics Communications*, vol. 185, no. 11, pp. 3028–3040, 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010465514002549>
- [10] I. Z. Reguly and G. R. Mudalige, "Productivity, performance, and portability for computational fluid dynamics applications," *Computers & Fluids*, vol. 199, p. 104425, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0045793020300013>
- [11] S. McIntosh-Smith, "Performance portability across diverse computer architectures," in *P3MA: 4th International Workshop on Performance Portable Programming models for Manycore or Accelerators*, 2019.
- [12] R. O. Kirk, G. R. Mudalige, I. Z. Reguly, S. A. Wright, M. J. Martineau, and S. A. Jarvis, "Achieving performance portability for a heat conduction solver mini-application on modern multi-core systems," in *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2017, pp. 834–841.
- [13] J. Gong, S. Markidis, E. Laure, M. Otten, P. Fischer, and M. Min, "Nekbone performance on gpus with openacc and cuda fortran implementations," *J. Supercomput.*, vol. 72, no. 11, p. 4160–4180, Nov. 2016. [Online]. Available: <https://doi.org/10.1007/s11227-016-1744-5>
- [14] Z. Dai, L. Deng, Y. Che, M. Li, J. Zhang, and Y. Wang, "Evaluating performance portability of five shared-memory programming models using a high-order unstructured cfd solver," *Journal of Parallel and Distributed Computing*, vol. 187, p. 104831, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0743731523002010>
- [15] A. Patil and C. García-Sánchez, "Quantifying the impact of urban geometric detail for urban air mobility risk forecasting," *Sustainable Cities and Society*, vol. 132, p. 106750, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2210670725006249>
- [16] L. Guastoni, A. Güemes Jiménez, A. Ianiro, S. Discetti, P. Schlatter, H. Azizpour, and R. Vinuesa, "Convolutional-network models to predict wall-bounded turbulence from wall quantities," 06 2020.
- [17] K. Koliogeorgi, G. Anagnostopoulos, G. Zampino, M. Sanchis, R. Vinuesa, and S. Xydis, "Auto-tuning multi-gpu high-fidelity numerical simulations for urban air mobility," in *Design, Automation & Test in Europe Conference & Exhibition, DATE 2024, Valencia, Spain, March 25-27, 2024*. IEEE, 2024, pp. 1–6. [Online]. Available: <https://doi.org/10.23919/DATES58400.2024.10546549>
- [18] G. Anagnostopoulos, P. Eleftherakis, A. C. Maroudis, A. Kapetanakis, K. Iliakis, L. Guastoni, M. Umair, J. Vet, J. Vincent, J. Gong, R. Vinuesa, and S. Xydis, "Refmap: From HPC to high fidelity CFD simulations to optimized flow prediction models for next-gen urban airspace mobility," in *Embedded Computer Systems: Architectures, Modeling, and Simulation - 24th International Conference, SAMOS 2025, Samos, Greece, June 28 - July 3, 2024*, ser. Lecture Notes in Computer Science, L. Carro, F. Regazzoni, and C. Pilato, Eds. Springer, 2025, pp. 1–16.
- [19] P.-E. Eleftherakis, G. Anagnostopoulos, A. Kapetanakis, M. Umair, J.-Y. Vet, K. Iliakis, J. Vincent, J. Gong, R. Vinuesa, and S. Xydis, "Poster: Performance portability in gpu-accelerated spectral finite element fluid simulations: A cross-layer exploration approach," in *Proceedings of the 22nd ACM International Conference on Computing Frontiers*, ser. CF '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 228–229. [Online]. Available: <https://doi.org/10.1145/3719276.3729458>