ParCFD 2026
37th International Conference on Parallel Computational Fluid Dynamics
Nov 03-06 2026, Nantes, FRANCE

VALIDATING THE ROOFPLANE PERFORMANCE MODEL FOR PARALLEL CFD ACROSS HPC ARCHITECTURES

M. Mosqueda-Otero[1], X. Álvarez-Farré[2], À. Alsalti-Baldellou[1,3], A. Gorobets[4], F. X. Trias[1] AND A. Oliva[1]

[1] Universitat Politècnica de Catalunya - BarcelonaTech, CTTC,
Carrer de Colom 11, 08222 Terrassa, Spain
e-mail: marcial.francisco.mosqueda@upc.edu, web page: https://www.cttc.upc.edu

[2] High-Performance Computing Team - SURF,
Science Park 140, 1098 XG Amsterdam, The Netherlands

[3] University of Padova, Department of Information Engineering,
Via Giovanni Gradenigo 6b, 35131 Padova PD, Italy

[4] Keldysh Institute of Applied Mathematics,
4A, Miusskaya Sq., Moscow 125047, Russia

Key words: Performance modelling, Roofplane, Roofline, Algebra-based CFD, Strong scaling, Heterogeneous HPC

Abstract. The growing gap between peak floating-point throughput and the memory and network bandwidths of modern supercomputers confines scientific computing applications to a small fraction of their theoretical performance. The classical roofline model captures this limitation at the device level, yet disregards the network as a third performance boundary in distributed-memory systems. The recently proposed roofplane model closes this gap with a third axis for network traffic and bandwidth, jointly visualising the three bottlenecks and supplying an analytical strong-scaling predictor, the parallel potential. The present work delivers the first cross-architecture empirical validation of the roofplane: a strong-scaling sweep of HPL, HPCG and the in-house TFA+HPC^2 algebra-based CFD code spans three current-generation systems—Deucalion (A64FX), MareNostrum 5 GPP (Intel Sapphire Rapids) and Snellius (AMD)—with consistent measurements of compute work, main-memory and network traffic that map each workload as a moving point in the roofplane and confront the predicted speed-up against the measured one.

1 INTRODUCTION

Distributed-memory CFD codes face a widening gap between peak floating-point throughput and the communication bandwidths, from two sides simultaneously: the memory-bandwidth wall at the device level and the network-bandwidth wall at the system level. Strong-scaling diagnoses have traditionally relied on experimental speed-up plots, which provide an empirical view but lack an analytical framework that anticipates which boundary the application is hitting. The classical roofline model [1] introduced a single-device upper bound; the roofplane model [2], a recent extension developed at our group, adds a third axis for network traffic and yields a parallel

potential, σ, that predicts whether a given workload distribution can hide its communication overhead.

However, the roofplane has so far been validated analytically on a sparse matrix-vector product and experimentally on a single CFD code on the Snellius supercomputer [2]. The present work extends that validation along two complementary directions: (i) a workload triplet that spans the three boundaries of the model, comprising HPL [3] as the compute-bound anchor, HPCG [4] as the memory- and network-sensitive benchmark, and our in-house TFA+HPC^2 algebra-dominant CFD code [5, 6], and (ii) three current-generation HPC architectures covering the most prevalent families in the Top500—ARM (A64FX), x86-Intel (Sapphire Rapids) and x86-AMD (Genoa).

2 THE ROOFPLANE MODEL

For a self-contained reading, the model is briefly recalled. Considering a parallel application that performs compute work W, accesses Q bytes of main memory and exchanges Q_x bytes through the interconnect, the achievable performance on a distributed-memory system reads:

$$R^{RP} = \min\left(\pi,\, I\,\beta,\, I_x\,\beta_x\right), \tag{1}$$

where π is the per-node peak floating-point throughput, β the main-memory bandwidth, β_x the effective network bandwidth, $I = W/Q$ the operational intensity and $I_x = W/Q_x$ the distributed intensity introduced by Álvarez-Farré *et al.* [2]. The projection of (1) onto the IR plane collapses to the classical roofline.

Furthermore, the model defines the parallel potential as:

$$\sigma = \max\left(1,\, \rho/I\right) \cdot I_x/\rho_x, \tag{2}$$

with $\rho = \pi/\beta$ and $\rho_x = \pi/\beta_x$. The condition $\sigma \geq 1$ delimits the ideal-scaling region, where data exchanges are entirely hidden behind bulk computations; the region $\sigma < 1$, in turn, identifies network-bound configurations whose speed-up is sub-linear regardless of the resources added. Hence, σ becomes an analytical strong-scaling predictor expressed solely in terms of system and application parameters.

3 WORKLOADS, ARCHITECTURES AND METHODOLOGY

3.1 Workloads

The workload triplet is selected to span the three boundaries of the roofplane. HPL [3] performs a dense LU factorisation through BLAS-3-dominant kernels; its high operational intensity is expected to place it near the compute-bound plateau, anchoring the upper corner of the model. HPCG [4] solves a synthetic 27-point stencil with a preconditioned conjugate gradient, exhibiting a memory-bound operational intensity and frequent halo exchanges that drive its distributed intensity downwards under strong scaling. Finally, TFA+HPC^2—our in-house algebra-dominant CFD framework [5]—runs the symmetry-preserving discretisation [6] of a turbulent channel flow Poisson solve, the same configuration adopted as the ETMM15 baseline [7]. Its kernel set comprises `SpMV`, `axpy`, `axty` and `dot` products, delivering an equivalent operational intensity close to 0.125 flop/byte; its memory-bound character at single-node level has been established [7], but its placement in the (I, I_x) plane under strong scaling remains to be mapped.

3.2 HPC architectures

Three current-generation systems are targeted to encompass the most prevalent architectural families: Deucalion at MACC for ARM (Fujitsu A64FX, HBM2), MareNostrum 5 GPP (MN5-GPP) at BSC for x86-Intel (Sapphire Rapids 8480+, DDR5), and Snellius at SURF for x86-AMD (Genoa 9654, DDR5). Table 1 holds the characteristic parameters used to deploy the roofplane on each architecture.

Table 1: Characteristic per-node parameters of the target HPC architectures used to deploy the roofplane and to position the workloads of Fig. 1. Snellius nodes carry a single-port NDR ConnectX-7 adapter (200 Gb/s intra-rack, 100 Gb/s inter-rack); the inter-rack rate is adopted as the effective network bandwidth, since strong-scaling runs span multiple racks.

System	CPU / Acc.	π [TF/s]	β [GB/s]	β_x [GB/s]
Deucalion	Fujitsu A64FX	3.07	1024.0 (HBM2)	12.5 (HDR100)
MN5-GPP	2× Intel SPR 8480+	7.17	614.4 (DDR5)	25.0 (NDR200)
Snellius	2× AMD Genoa 9654	7.37	921.6 (DDR5)	12.5 (HDR100)

4 PRELIMINARY ANALYTICAL MAPPING

Figure 1 holds the (I, I_x)-plane projection of the roofplane for the three target architectures, with the **predicted** positions of HPL, HPCG and TFA+HPC^2 at increasing strong-scaling levels superimposed as discrete points. The σ-isolines, drawn for $\sigma = 2^n$ with integer n, partition the plane into the ideal-scaling ($\sigma \geq 1$) and limited-scaling ($\sigma < 1$) regions.

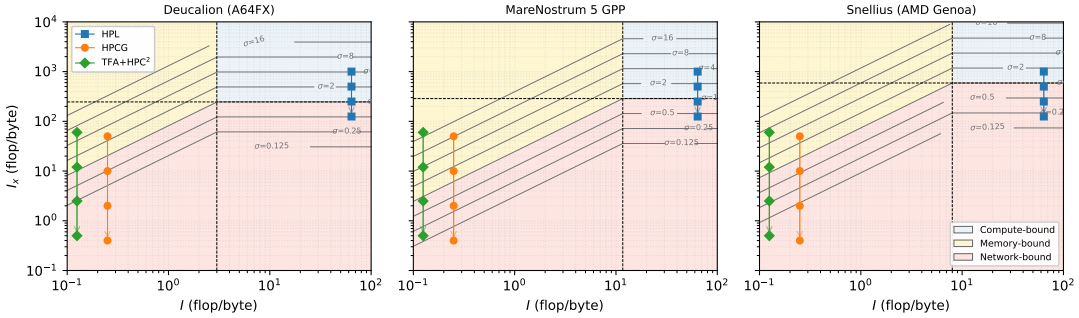

Figure 1: Predicted (I, I_x) mapping of HPL, HPCG and TFA+HPC^2 at four strong-scaling levels on the three target architectures. Best-guess system parameters; σ-isolines drawn for $\sigma = 2^n$.

Three observations follow from the analytical mapping. First, HPL is anticipated to remain in the compute-bound region across all strong-scaling levels on the three architectures, providing a stable anchor for the upper corner of the model. Second, HPCG and TFA+HPC^2 are expected to start in the memory-bound region at low node counts and to drift towards the network-bound region as the workload distribution increases the halo-to-bulk ratio; with the DDR-based systems exhibiting a wider ideal-scaling region than the HBM-based A64FX. Third, the relative ordering of HPCG and TFA+HPC^2 in the (I, I_x) plane is highly sensitive to the communication topology of the underlying solver.

5 CLOSING REMARKS

This work proposes the first cross-architecture empirical validation of the roofplane model on a workload triplet that spans its three boundaries—HPL, HPCG and the in-house TFA+HPC^2 CFD code—on Deucalion, MN5-GPP and Snellius. The harmonised measurement methodology, will be detailed in the full paper as a contribution towards portable performance characterisation in distributed-memory systems.

By the time of the conference, the full 3×3 measurement matrix and the comparison between the analytical parallel potential σ and the measured speed-up will be delivered. Further work will encompass refining the system-dependent overhead factor χ of the roofplane [2] to incorporate NUMA and topology effects, as well as extending the workload set towards SpMM-accelerated TFA+HPC^2 configurations [8].

ACKNOWLEDGEMENTS

M.MO, and F.X.T are supported by the Ministerio de Ciencia, Innovación y Universidades, Spain, SIMEX project (PID2022-142174OB-I00). M.MO is also supported by the Catalan Agency for Management of University and Research Grants (2024 FI-1 00684). Computational time on MareNostrum 5, Snellius and Deucalion was granted by BSC, SURF and MACC, respectively; the authors thankfully acknowledge these institutions.

REFERENCES

[1] S. Williams, A. Waterman, and D. Patterson, "Roofline: An insightful visual performance model for multicore architectures," *Communications of the ACM*, vol. 52, no. 4, pp. 65–76, 2009. [Online]. Available: https://doi.org/10.1145/1498765.1498785

[2] X. Álvarez-Farré *et al.*, "Roofplane: A visual performance model for high-performance distributed systems," 2026, manuscript in preparation.

[3] A. Petitet, R. C. Whaley, J. Dongarra, and A. Cleary, "HPL – a portable implementation of the high-performance linpack benchmark for distributed-memory computers," https://www.netlib.org/benchmark/hpl/, 2018.

[4] J. Dongarra, M. A. Heroux, and P. Luszczek, "High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems," *The International Journal of High Performance Computing Applications*, vol. 30, no. 1, pp. 3–10, 2016. [Online]. Available: https://doi.org/10.1177/1094342015593158

[5] X. Álvarez-Farré, A. Gorobets, F. X. Trias, R. Borrell, and G. Oyarzun, "HPC^2 – A fully portable algebra-based framework for heterogeneous computing. Application to CFD," *Computers & Fluids*, vol. 173, pp. 285–292, 2018. [Online]. Available: https://doi.org/10.1016/j.compfluid.2018.01.034

[6] F. X. Trias, O. Lehmkuhl, A. Oliva, C. D. Pérez-Segarra, and R. W. C. P. Verstappen, "Symmetry-preserving discretization of Navier–Stokes equations on collocated unstructured grids," *Journal of Computational Physics*, vol. 258, pp. 246–267, 2014. [Online]. Available: https://doi.org/10.1016/j.jcp.2013.10.031

[7] M. Mosqueda-Otero, A. Alsalti-Baldellou, J. Plana-Riu, X. Álvarez-Farré, G. Colomer, F. X. Trias, A. Gorobets, and A. Oliva, "A portable algebraic implementation for reliable industrial LES," in *15th International ERCOFTAC Symposium on Engineering, Turbulence, Modelling and Measurements (ETMM15)*, 2025, pp. 683–687. [Online]. Available: https://doi.org/10.5281/zenodo.17280519

[8] A. Alsalti-Baldellou, X. Álvarez-Farré, F. X. Trias, and A. Oliva, "Exploiting spatial symmetries for solving Poisson's equation," *Journal of Computational Physics*, vol. 486, p. 112133, 2023. [Online]. Available: https://doi.org/10.1016/j.jcp.2023.112133