ParCFD2024
35th International Conference on Parallel Computational Fluid Dynamics
Sep. 02-04, 2024, Bonn, Germany

In-Situ Visualization With Ascent and NekRS for Large-Scale CFD Problems on GPUs

Jens Göbbert[a,*], Damian Alvarez[a], Mathis Bode[a], Paul Fischer[b,c], Christos Emmanouil Frouzakis[d], Joseph A. Insley[c], Yu-Hsiang Lan[b,c], Victor A. Mateevitsi[c], Misun Min[c], Michael E. Papka[c], Silvio Rizzi[c], Roshan J. Samuel[e] and Jörg Schumacher[e]

[a] *Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, Wilhelm-Johnen-Straße, 52425 Jülich, Germany*
[b] *University of Illinois at Urbana-Champaign, 601 E. John Street, IL 61820 Urbana-Champaign, IL, USA*
[c] *Argonne National Laboratory, 9700 S Cass Ave, IL 60439 Lemont, IL, USA*
[d] *ETH Zurich, Sonneggstr. 3, CH-8092 Zurich, Switzerland*
[e] *TU Ilmenau, Ehrenbergstraße 29, 98693 Ilmenau, Germany*

ARTICLE INFO[†]

Keywords:
Computational Fluid Dynamics Workflows;
NekRS;
Graphics Processing Unit;
Visualization

ABSTRACT

This paper discusses in-situ visualization in the context of high-order GPU-based CFD solvers. For this purpose, ASCENT was coupled with NekRS and successfully used on JUWELS Booster and Frontier to visualize CFD applications, a Rayleigh-Bénard convection case and a Pebble Bed reactor, at scale. The setup allowed time-resolved visualization at low additional computational cost.

1. Introduction

Exascale supercomputers make it possible to address larger scientific problems using numerical methods and/or to solve them in a shorter time. This enables, for example, the simulation of flow problems with ever increasing scale separation, on the one hand as a result of turbulence or on the other hand due to multi-physics effects such as combustion, or the training of ever larger networks with ever more data. One energy-efficient way to realize an exascale supercomputer is by exploiting GPUs. Systems such as Frontier, the first US exascale supercomputer, or JUPITER, probably the first European exascale supercomputer, obtain the majority of their FLOP/s (floating point operations per second) from GPUs. The resulting heterogeneous systems with CPUs and GPUs increase the level of complexity by a further level, which makes data handling in particular more difficult. In the negative case, separate CPU and GPU memory forces a large number of expensive copy operations. Conversely, applications are generally only computationally efficient if they avoid copy operations to the HDD and limit copy operations between CPU and GPU memory.

Irrespective of the efficient supercomputer utilization described above, the size of the data generated by exascale simulations can be a further obstacle to successfully solving the scientific problem on which a simulation is based. Checkpoint files from exascale simulations can be several 100 TB in size, which is often too large for conventional tools for analysis and visualization. This means that writing the results is not only very expensive, which limits the number of checkpoint files that can be written, but the handling of checkpoint files in post-processing is also difficult.

In-situ analysis and visualization, i.e. the immediate processing of data while it is still in the computing memory, can both reduce copying operations and save the time-consuming handling of checkpoint files. Furthermore, they enable real-time information on the simulation, which can be used for simulation monitoring. Hence, in-situ analysis and visualization give researchers immediate access to visualizations even in the exascale era, enabling faster decision-making, and a deeper understanding of the phenomena under investigation.

Fluid mechanics often features transient or fluctuating processes and is therefore an ideal field of application for in-situ analysis and visualization. Especially in the context of flow solvers, which mainly run on GPUs, enormous

[†] This paper is part of the ParCFD 2024 Proceedings. A recording of the presentation is available on YouTube. The DOI of this document is 10.34734/FZJ-2025-02500 and of the Proceedings 10.34734/FZJ-2025-02175.

* Corresponding author

j.goebbert@fz-juelich.de (J. Göbbert); d.alvarez@fz-juelich.de (D. Alvarez); m.bode@fz-juelich.de (M. Bode); fischerp@illinois.edu (P. Fischer); cfrouzakis@ethz.ch (C.E. Frouzakis); insley@anl.gov (J.A. Insley); lan14@illinois.edu (Y. Lan); vmateevitsi@anl.gov (V.A. Mateevitsi); mmin@anl.gov (M. Min); papka@anl.gov (M.E. Papka); srizzi@anl.gov (S. Rizzi); roshan-rohn.samuel@tu-ilmenau.de (R.J. Samuel); joerg.schumacher@tu-ilmenau.de (J. Schumacher)

ORCID(s): 0000-0002-3807-6137 (J. Göbbert); 0000-0002-8455-3598 (D. Alvarez); 0000-0001-9922-9742 (M. Bode); 0000-0002-6506-4502 (P. Fischer); 0009-0006-0963-4650 (C.E. Frouzakis); 0000-0002-6955-869X (J.A. Insley); 0000-0002-1680-675X (Y. Lan); 0000-0002-6677-7520 (V.A. Mateevitsi); 0000-0002-5646-5689 (M. Min); 0000-0002-6418-5767 (M.E. Papka); 0000-0002-3804-2471 (S. Rizzi); 0000-0002-1280-9881 (R.J. Samuel); 0000-0002-1359-4536 (J. Schumacher)

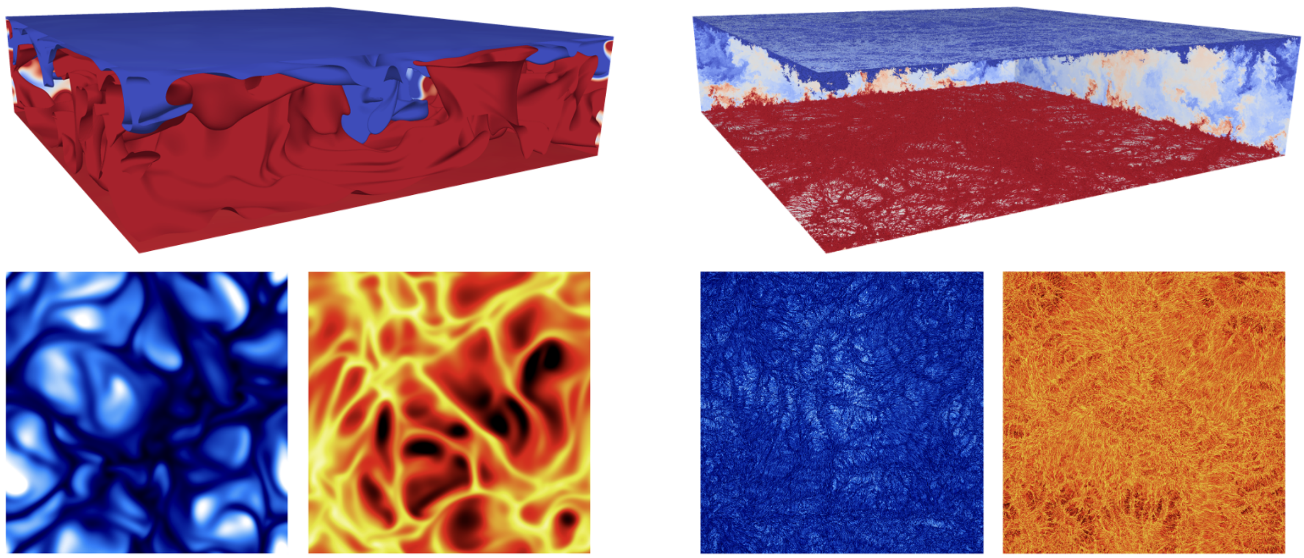

Figure 1: RBC overview for $Ra = 10^6$ (left) and $Ra = 10^{12}$ (right) in the top line. Bottom line shows cut planes through the boundary layers at the top and bottom.

advantages can arise. This is demonstrated in this paper using the example of NekRS [1] for two applications.

In the following, it is first explained how in-situ visualization can be integrated into NekRS using ASCENT. Subsequently, measurements for a Rayleigh-Bénard convection (RBC) case on JUWELS Booster and a pebble bed reactor (PBR) on Summit and Frontier are presented. This work finishes with conclusions.

2. Methodology

This work couples the state-of-the-art GPU-based CFD solver NekRS with the in-situ visualization library ASCENT to reduce the overall time-to-solution (incl. visualization) of large-scale CFD simulations. In selecting a suitable in-situ analysis and visualization library for this work, ASCENT emerged as the preferred choice due to several distinct advantages. First and foremost, its lightweight design and ease of integration into existing simulation code bases significantly reduce the implementation overhead. In addition, ASCENT supports zero-copy GPU-to-GPU data passing, which allows the direct transfer of device pointers between the simulation code and ASCENT, ensuring that data remains exclusively on the GPU. Such an approach effectively eliminates the need to pass data back to the CPU, eliminating a common bottleneck in high-performance computing applications.

The implications of this GPU-centric data handling strategy are profound, particularly in terms of minimizing the memory footprint of the in-situ visualization process. Traditional data movement to the CPU introduces significant memory overhead and duplication [2], and by bypassing this transfer, ASCENT facilitates more efficient use of computational resources.

3. Applications

The coupled in-situ workflow was successfully used for two applications. The focus in each case was to avoid data copying and to visualize directly on the GPUs. The visualization frequency was chosen to capture the smallest temporal scales and the computational overhead was measured. The first use case was Rayleigh-Bénard convection (RBC) on JUWELS Booster. JUWELS Booster has 936 GPU-accelerated nodes, is fully directly liquid cooled, and is based on Atos' Sequana XH2000 architecture. Each node relies on 2x AMD EPYC 7042 processors that orchestrate 4x NVIDIA A100 GPUs. These GPUs have the SXM4 form factor, and are integrated in a so-called Redstone board, in which the GPUs are interconnected in direct all-to-all NVLink3 fabric. The RBC simulation [3] used up to 900 nodes for RAYLEIGH numbers (Ra) up to 10^{12} that ran on 46.7 billion grid points. The resulting visualizations of two different RAYLEIGH numbers, 10^6 and 10^{12}, are presented in Fig. 1.

The second example is the full core of pebble bed reactor (PBR), illustrated in Fig. 2, which has 352,625 spherical pebbles and the fluid mesh comprising 98 million spectral elements of order 7, or 33.8 billions grid points. The PBR

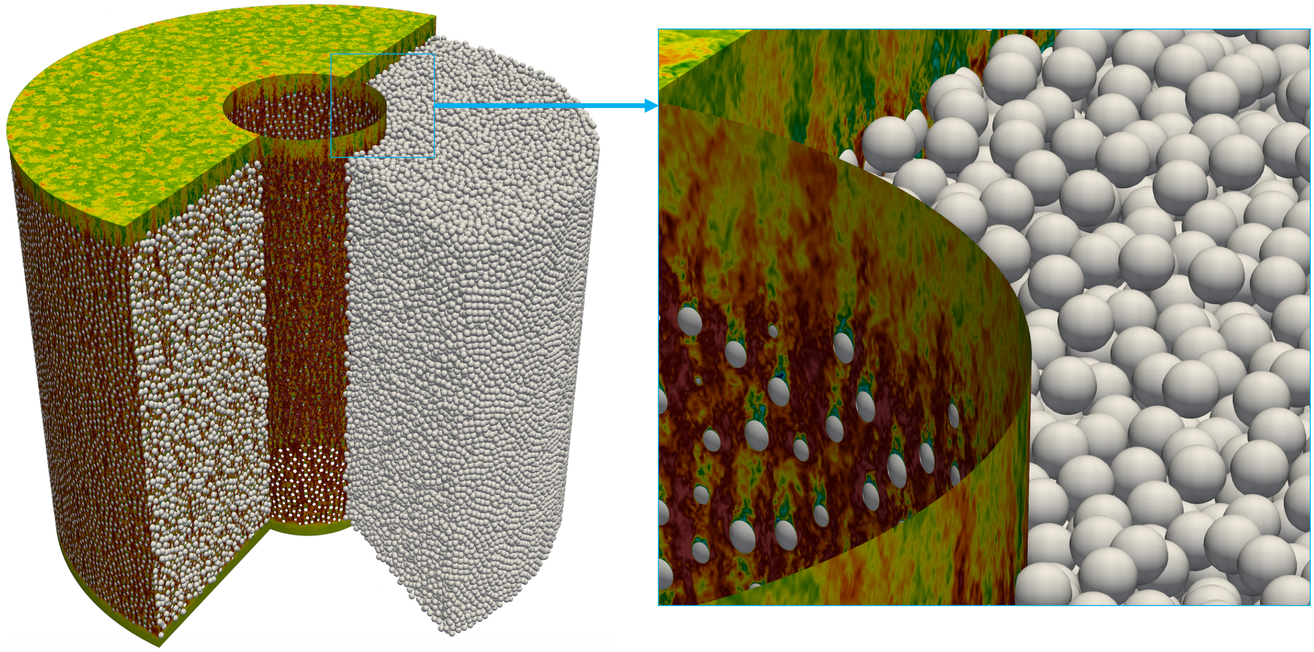

Figure 2: Full core LES of PBR with 352k pebbles and (right) close-up showing velocity distribution.

was the first wall-resolved large Eddy simulation (LES) for the full reactor core [4]. It has been demonstrated NekRS can simulate a single flow-through time in less than 6 hours using the entire Summit, a post-petascale machine in OLCF. The in-situ workflow makes it possible to visualize high fidelity transient solutions at a very high resolution which provides more insights in dynamic turbulent structure. Using 1,400 nodes on Frontier, the PBR case renders 8K images (8,192 x 8,192 pixels) every 20 timesteps and producing 710 images in just 4 hours.

4. Discussion and conclusions

The implemented in-situ workflow allows to reduce the overhead for a time-resolved visualization below 10 % for the RBC application and enables the realization of high-resolution images for the PBR case. In this work, the focus was on avoiding copying processes between GPU and CPU. In the future, it could be interesting to do the visualization in parallel on the otherwise idling CPUs, although this requires a periodic copy operation. For more complex visualizations and if the CPU is otherwise not used at all, this could be the most efficient solution overall.

Acknowledgements

This work was supported by and used resources of the Argonne Leadership Computing Facility, which is a U.S. Department of Energy Office of Science User Facility supported under Contract DE-AC02-06CH11357. This work was supported by Northern Illinois University. This work was supported in part by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract DE-AC02-06CH11357, through the grant "Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery", program manager Dr. Margaret Lenz. The authors from JSC acknowledge computing time grants for the project TurbulenceSL by the JARA-HPC Vergabegremium provided on the JARA-HPC Partition part of the supercomputer JURECA at Jülich Supercomputing Centre, Forschungszentrum Jülich, the Gauss Centre for Supercomputing e.V.[1] for funding this project by providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC), and funding from the European Union's Horizon 2020 research and innovation program under the Center of Excellence in Combustion (CoEC) project, grant agreement no. 952181. Support by the Joint Laboratory for Extreme Scale Computing (JLESC)[2] for traveling is acknowledged.

[1] https://www.gauss-centre.eu

[2] https://jlesc.github.io/

References

[1] P. Fischer, S. Kerkemeier, M. Min, Y.-H. Lan, M. Phillips, T. Rathnayake, E. Merzari, A. Tomboulides, A. Karakus, N. Chalmers, T. Warburton, NekRS, a GPU-accelerated spectral element Navier–Stokes solver, Parallel Computing 114 (2022) 102982. doi:10.1016/j.parco.2022.102982.

[2] V. A. Mateevitsi, M. Bode, N. Ferrier, P. Fischer, J. H. Göbbert, J. A. Insley, Y.-H. Lan, M. Min, M. E. Papka, S. Patel, S. Rizzi, J. Windgassen, Scaling Computational Fluid Dynamics: In Situ Visualization of NekRS using SENSEI, in: Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis, ACM, New York, NY, USA, 2023, pp. 862–867. doi:10.1145/3624062.3624159.

[3] R. J. Samuel, M. Bode, J. D. Scheel, K. R. Sreenivasan, J. Schumacher, No sustained mean velocity in the boundary region of plane thermal convection, Journal of Fluid Mechanics 996 (2024) A49. doi:10.1017/jfm.2024.853.

[4] M. Min, Y.-H. Lan, P. Fischer, E. Merzari, S. Kerkemeier, M. Phillips, T. Rathnayake, A. Novak, D. Gaston, N. Chalmers, T. Warburton, Optimization of Full-Core Reactor Simulations on Summit, in: SC22: International Conference for High Performance Computing, Networking, Storage and Analysis, IEEE, New York, NY, USA, 2022, pp. 1–11. doi:10.1109/SC41404.2022.00079.