Erlangen National High Performance Computing Center

Newsletter #6

February 2022

In This Issue...

Quick News

Node with Intel Optane nonvolatile memory

A node with 256 GB DDR4 RAM and 256 GB of Intel "Optane" nonvolatile memory is now available in the NHR@FAU test cluster. It is equipped with two 32-core Intel Xeon Platinum 8362 CPUs.

Recording of HPC Cafe on file systems available

A video recording of the January 18 HPC Cafe on how to avoid overload on shared file systems is now available at fau.tv/clip/id/40199. We strongly recommend it to all new customers of NHR@FAU.

PerfLab talk on PINNs

On February 15, Dr. Stefano Markidis (see right) will give a PerfLab seminar talk on physics-informed neural networks. See hpc.fau.de/?p=11051 for details.

Newsletter mailing list

You can now subscribe to a mailing list in order to get notified when a new NHR@FAU newsletter is out: lists.fau.de/cgi-bin/listinfo/nhr-newsletter

Missed a talk?

The NHR@FAU YouTube channel at youtube.com/NHRFAU provides recordings of some talks, lectures, and courses, as far as this is covered by copyright restrictions.

"Fritz" Cluster in Early User Mode

Fritz, the new CPU cluster at NHR@FAU, is fully powered on and in early user operations since the end of January. The 944 compute nodes are equipped with two 36-core Intel Xeon "Ice Lake" Platinum 8360Y each, amounting to 67968 cores in total. Due to IT supply shortages, only a part of the 944 nodes is currently equipped with an InfiniBand network interface. Many details concerning the software installation and hardware setup are still in flux, but early users are currently working together with NHR@FAU to bring the system to a usable state.

Minisymposium on Performance Modeling at SIAM Parallel Processing 2022

Together with Dr. Alexandru Calotoiu from ETH Zurich, Georg Hager from NHR@FAU is organizing a two-part minisymposium at the virtual SIAM *Parallel Processing for Scientific Computing 2022* conference. The two-part minisymposium, to take place on February 25, will provide an overview of the current state of the art in performance, or more generally, resource modeling of parallel code. Resource modeling is not only an active field of research but also a critical component in monitoring and operations of HPC centers. The hardware focus will be very broad, from the node to the massively parallel level, including standard multicore systems, GPUs, and reconfigurable hardware. Contributions will cover fundamental research as well as tools development and case studies. An impressive lineup of international speakers has been brought together for this minisymposium. See blogs.fau.de/hager/archives/9085 for details.

HPC-Europa: Dr. Stefano Markidis to Visit FAU

Stefano Markidis from KTH Royal Institute of Technology in Stockholm will be visiting NHR@FAU from April 4–15. The visit is funded via a grant through the HPC-Europa3 Transnational Access program. HPC-Europa3 gives European researchers the opportunity to access top-notch HPC systems and collaborate on HPC projects. Dr. Markidis, who is an associate professor and researcher in High-Performance Computing at KTH, has initiated the research on idle waves in high-performance computing in 2015. PhD student Ayesha Afzal (see the April 2021 newsletter highlight) is pushing forward in this interesting new field and will be working with him during his stay.

Annual Beginner's Parallel Programming Course

From March 8–10, our annual HPC beginner's course takes place online. Together with Leibniz Supercomputing Center (LRZ), we introduce the basic architecture of supercomputers and parallel programming with OpenMP and MPI. We also briefly touch on profiling and tracing tools. Attendees will get access to an NHR@FAU cluster for the hands-on exercises. Details and registration at: indico.scc.kit.edu/event/2608/.

GROMACS Performance on Different GPU Types

With the upcoming *Alex* GPU cluster (see the previous NHR@FAU Newsletter), more than 500 NVIDIA A40 and A100 GPUs will boost the available compute power for MD and AI workloads. This does not mean, however, that the GPUs in the TinyGPU cluster become useless. Indeed, depending on the application case at hand, even consumer GPUs can be competitive. The Gromacs Molecular Dynamics package shows some particularly interesting performance characteristics. Anna Kahler has provided some benchmark data for current GPU generations in use at NHR@FAU with relevant cases at: hpc.fau.de/?p=10609.

Spotlight: Prof. Dr. Harald Köstler

Harald Köstler holds a PhD and a Habilitation degree in Computer Science from FAU. He is a member of the Chair for System Simulation (Computer Science 10) and leads the research division at NHR@FAU. Recent research includes efficient parallel numerical algorithms and code generation technology for HPC. At NHR@FAU, he complements the research on performance engineering from the application perspective, e.g., in the areas of computational fluid dynamics and image processing. In his spare time, he plays real-time strategy games like StarCraft II. Harald is married and has a young daughter.

February Highlight: Deep Learning for Fast Aerodynamic Evaluation of Vehicle Shapes

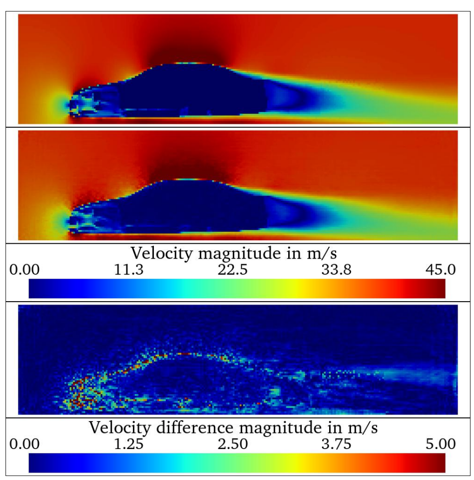

Comparing velocity field from CFD prediction (top) with DL model trained on DrivAer (bottom).

The aerodynamic optimization process of cars requires multiple iterations between aerodynamics engineers and stylists; the improvements are slow because of the CFD computation time. Surrogate models can produce fast approximate results in a constrained design space. A primary drawback of current models is that they can work only on the parameterized geometric features they were trained with. We show that deep learning models can predict an arbitrary input geometry's drag coefficient and velocity field without explicit parameterization. We use two data sets based on the DrivAer car geometry for training a modified U-Net architecture that uses Signed Distance Fields to represent the input geometries. Our models outperform the existing models by at least 11% in prediction accuracy for the drag coefficient and brings us a step closer to a model that can be used for approximate aerodynamic evaluation of unseen, arbitrary vehicle shapes. The paper, authored by Sam Jacob *et al.* from the Chair for System Simulation, was accepted for publication in SAE International Journal of Passenger Vehicle Systems; a preprint is available at: arXiv:2108.05798.

FAQ Corner: File Systems

Why the need for several file systems?

Different file systems have different features; for example, a central NFS server has massive bytes for the buck but limited data bandwidth, while a parallel file system is much faster but smaller and usually available to one cluster only. A node-local SSD, one the other hand, has the advantage of very low latency but it cannot be accessed from outside a compute node.

What is a parallel file system?

In a parallel file system (PFS), data is distributed not only across several disks but also multiple servers in order to increase the data access bandwidth. Most PFS's are connected to the high-speed network of a cluster, and aggregated bandwidths in the TByte/s range are not uncommon. High bandwidth can, however, only be obtained with large files and streaming access.

What is file system metadata?

Metadata comprises all the bookkeeping information in a file system: file sizes, permissions, modification and access times, etc. A workload that, e.g., opens and closes files in rapid succession leads to frequent metadata accesses, putting a lot of strain on any file server infrastructure. This is why a small number of users with "toxic" workload can slow down file operations to a crawl for everyone. Note also that especially parallel file systems are ill-suited for metadata-heavy operations (see above).

Why should I care?

Not only may efficient file operations speed up your own code (if file I/O is what you must do); they will also reduce the burden on shared file servers and thus leave more performance headroom for other users of the resource. Hence, it is a matter of thoughtfulness to optimize file accesses even if *your* performance gain is marginal.

Solution to the HPC and Computer History Crossword Puzzle

Across
1. Keep it simple, make it fast!
3. This Gene limits your performance
6. Difference or Analytical Engine? He built neither
11. “It’s easier to ask for forgiveness than it is to get permission”
14. High-speed networking for the masses
18. 4.0! Finally!
19. The countess of code
20. The best computers are born in your parents’ living room
22. This Gene scales quite well

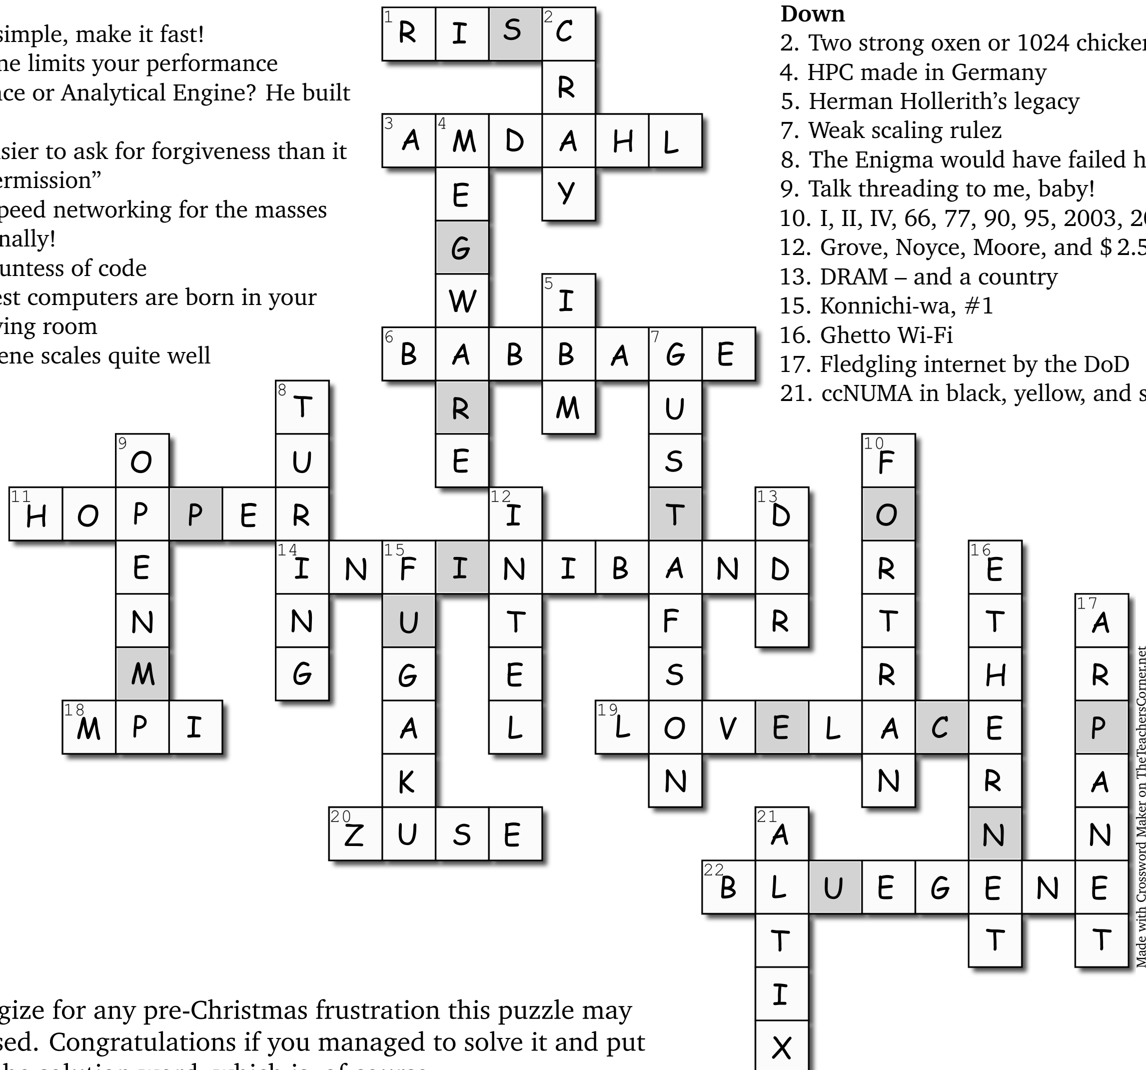

Down
2. Two strong oxen or 1024 chickens?
4. HPC made in Germany
5. Herman Hollerith’s legacy
7. Weak scaling rulez
8. The Enigma would have failed his test
9. Talk threading to me, baby!
10. I, II, IV, 66, 77, 90, 95, 2003, 2008, 2018
12. Grove, Noyce, Moore, and $ 2.5 M
13. DRAM – and a country
15. Konnichi-wa, #1
16. Ghetto Wi-Fi
17. Fledgling internet by the DoD
21. ccNUMA in black, yellow, and silver

We apologize for any pre-Christmas frustration this puzzle may have caused. Congratulations if you managed to solve it and put together the solution word, which is, of course, “SUPERCOMPUTING.”

Contact Information:
Erlangen National High Performance Computing Center
Martensstr. 1
91058 Erlangen, Germany

https://nhr.fau.de/
General contact: hpc@fau.de
Issues: hpc-support@fau.de
NHR newsletters: http://tiny.cc/NHR-NL