NHR FAU

Erlangen National High Performance Computing Center

Newsletter #1

April 2021

In This Issue...

Quick News

Missed a talk?

The NHR@FAU YouTube channel at tiny.cc/YT-NHR-FAU provides recordings of some talks, lectures, and courses, as far as this is covered by copyright restrictions.

New Resources

As of this month, the NHR@FAU resources have been augmented by several new large-memory and GPU-equipped compute nodes in our TinyFat and TinyGPU clusters. See our systems documentation for details.

New LIKWID release upcoming

The bugfix release 5.1.1 of LIKWID came out by the end of March. Shortly afterwards, AMD released the documentation for their "Zen 3" chips and Intel brought their "Ice Lake SP" chips to market. The LIKWID team is preparing the new major release 5.2.0 offering full support for both architectures.

HPC Statustagung

On April 27, 2021, the German Federal Ministry of Education and Research, together with the Gauss Alliance and the University of Göttingen, will host a conference on current developments in the field of High Performance Computing (HPC). The conference will be conducted as an online event.

Welcome to the first NHR@FAU newsletter! Here we publish a summary of latest news, past and upcoming events, and frequently asked questions with their answers. Every newsletter will also highlight one NHR@FAU employee so you gradually get to know who is behind this new organization.

Past Events

NHR@FAU is ramping up its course program in order to fulfill its mission to educate HPC beginners, users, developers, and experts at FAU and beyond. In the first half of 2021, a number of courses have been conducted or contributed to by NHR@FAU staff or invited lecturers:

The *EXA2PRO-EoCoE joint workshop* (February 22–24) provided insights into the research results of the EXA2Pro and EoCoE EU projects and presented their tangible output in terms of software and tools. Researchers from NHR@FAU and the Chair for System Simulation (LSS) showcased state-of-the-art performance engineering and code generation techniques.

In the *38th VI-HPS Virtual Tuning Workshop* (March 1–3), lecturers from the Barcelona Supercomputing Center (BSC), the University of Versailles (UVSQ), and the Jülich Supercomputing Centre (JSC) introduced the use of advanced performance analysis tools such as MAQAO, Score-P, Paraver, Tau, and Dimemas. 23 participants from all over the world (including Brazil and the USA) got their hands dirty on RRZE's *Meggie* cluster.

Guest lecturer Klaus Iglberger conducted his *beginner's C++ course* from March 15–19. More than thirty participants embarked on an intense virtual *tour de force* through this popular but complex programming language. This event was made possible through the generous support by KONWIHR, the Competence Network for Scientific High Performance Computing in Bavaria.

Continuing a long-standing collaboration with TU Wien, Georg Hager and Gerhard Wellein from NHR@FAU provided their popular tutorial *Node-Level Performance Engineering* as a three-day online event with hands-on exercises at the *Vienna Scientific Cluster*.

The *Parallel Programming for High Performance Systems* course has been conducted for more than two decades as an RRZE-LRZ collaboration. For 2021, it was compacted to three days (April 13–15) with a strong focus on MPI and OpenMP parallel programming, and converted to an online event. By the generous support of KONWIHR, participants from Bavarian universities and research institutions could participate free of charge.

Our *HPC Café* (on every second Tuesday of the month) has established itself as a popular event for discussing current events and developments. After the usual freestyle Q&A we provide short (or not so short) presentations about topics of interest. Many of these are available on the HPC Café website, including video recordings in some cases. On April 13, guest speaker Valentin Churavy from the MIT Julia Lab talked about "Julia in HPC." On May 11, Harald Köstler, Head of Research at NHR@FAU, will give a practical introduction to "Continuous Integration, Deployment and Benchmarking for HPC."

NHR PerfLab Seminar Series

The NHR PerfLab is a collaborative effort by the NHR centers at FAU, RWTH Aachen, ZIB Berlin, and the University of Paderborn to combine resources and activities around performance engineering, code analysis, HPC computer architectures, and tools. The *NHR PerfLab seminar* is a series of public and internal talks. It provides a platform for NHR researchers and practitioners to share their results and ideas. Public talks are announced on the NHR website. Beyond the seminar, performance engineers and support personnel from all four centers will meet once per month to discuss interesting support cases.

Spotlight: Prof. Gerhard Wellein

Gerhard Wellein is the director and lead PI of the Erlangen National High Performance Computing Center (NHR@FAU), a Professor for High Performance Computing at the Department for Computer Science at FAU, the deputy speaker of KONWIHR, and member of the scientific steering committee of the Gauss Centre for Supercomputing (GCS).

Gerhard Wellein has more than twenty years of experience in teaching HPC techniques to students and scientists from computational science and engineering. His research interests focus on performance modeling and performance engineering, architecture-specific code optimization, novel parallelization approaches, and hardware-efficient building blocks for sparse linear algebra and stencil solvers. He has been the lead PI in numerous HPC projects, including the German-Japanese project "Equipping Sparse Solvers for Exascale" (ESSEX) within the DFG priority program SPPEXA ("Software for Exascale Computing").

April Highlight

A research paper titled "Analytic Modeling of Idle Waves in Parallel Programs: Communication, Cluster Topology, and Noise Impact" has been accepted for publication at ISC High Performance 2021, the premier European conference on high performance computing. Ayesha Afzal, PhD student at the professorship for high performance computing and lead author of the paper, has been researching analytic modeling techniques for MPI-parallel programs for more than two years. The paper presents new insights about *idle waves,* phases of idleness that "ripple" through a parallel code running on a cluster. The work was supported by KONWIHR and the BMBF projects Metacca and SeASiTe.

Read more at arXiv:2103.03175.

FAQ corner

What is "memory bandwidth"?

Memory bandwidth (measured in bytes per second) is the maximum rate at which a CPU can transfer data to and from memory. It is one of the major bottlenecks in computing – many algorithms, no matter how well they are optimized, are so hungry for data that the performance of their implementation is limited by the memory bandwidth. Current server CPUs have memory bandwidths of 100-200 Gbyte/s, GPUs are in the 1 Tbyte/s range.

How much faster is a GPU compared to a CPU?

Current top-notch NVIDIA GPUs have memory bandwidths of around 1 Tbyte/s and peak performance rates beyond 10 Tflop/s in double precision. A high-end server CPU such as an AMD Rome 64-core model has a memory bandwidth of 180 Gbyte/s and tops out at 2.5 Tflop/s. You do the math, but in reality it's more complicated than that because not every algorithm is well suited for GPUs, and not every software package is well optimized for both platforms.

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

Link to NHR newsletters