NHR FAU

Erlangen National High Performance Computing Center

Newsletter #2

June 2021

In This Issue...

Quick News

New MD software modules

A new module for the CP2K software package has been made available on the Meggie cluster.

LAMMPS is now available as a module on the Emmy and Woody clusters. LAMMPS was built with OpenMP and Kokkos support as well as CUDA support for use on TinyGPU.

New LIKWID release

Version 5.2.0 of the LIKWID tool suite has just been released. See the news item on our website for a list of updates.

KONWIHR proposal deadline on September 1

KONWIHR project proposals can be submitted anytime, but they will only be reviewed at semi-annual intervals. The next deadline is September 1, so if you have a project that is in line with KONWIHR's list of eligible activities, now is the time to get into gear.

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.

Welcome to the NHR@FAU newsletter! This is where 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.

Breaking News: NHR@FAU and RRZE to Install New Cluster System From MEGWARE

Late in 2020, NHR@FAU put out the first NHR system for tender. At a total volume of 10 million €, this is the largest parallel computer ever to be installed at FAU. After a fierce competition, *MEGWARE Computer Vertrieb und Service GmbH* from Chemnitz, Germany came out first in the ranking of bidders. In the Q4/21–Q1/22 timeframe, a three-part system will be installed.

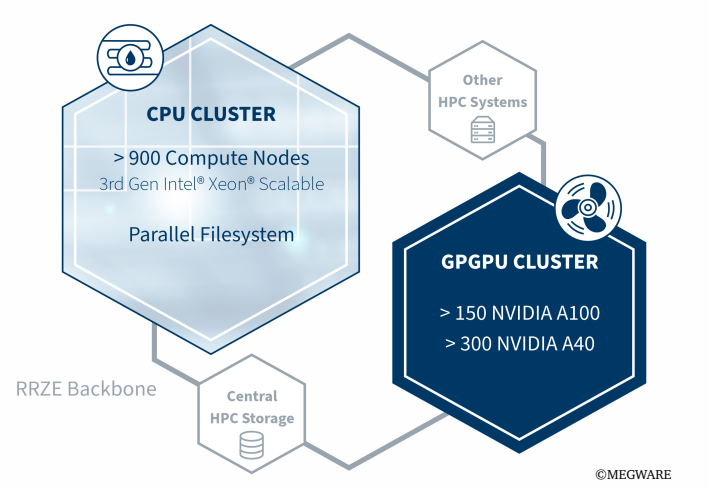

Besides a traditional multicore-based cluster with about ten times the peak performance of *Meggie*, two GPGPU partitions will be tailored to the needs of Molecular Dynamics (MD) and Artificial Intelligence (AI) researchers. A part of the computational resources will be dedicated to Tier-3 customers from FAU and Northern Bavaria, and will thus be available with a low entry threshold. The rest will be accessible after a scientific and technical review process, the structure of which is currently being discussed among the NHR centers.

NHR PerfLab Seminar Series

The NHR PerfLab seminar series has gained considerable momentum. It provides a platform for the presentation of high-profile research conducted in and around the NHR centers. From April to June, guest speakers from KIT Karlsruhe, Saarland University, Fraunhofer ITWM, ZIB Berlin, Intel, and FAU have shared their wisdom about various current research topics. If possible, presentation slides and video recordings are made available on the PerfLab Seminar website.

Upcoming Talk: Tasking in OpenMP 5.0

On June 22, at 2:00 p.m., Dr. Christian Terboven from the RWTH Aachen University IT Center will give a talk on "Tasking in OpenMP 5.0" in the NHR PerfLab seminar. The presentation will focus on performance aspects of tasking and new functionality introduced by the latest OpenMP standards. Get the full abstract and Zoom link here.

Multi-GPU Gromacs Jobs on TinyGPU?

If you ask yourself whether you should use more than one GPU in a node for a GROMACS job, the answer is that this depends crucially on the problem size. In a benchmark test running a standard MD simulation with explicit water (65,209 atoms) on an RTX2080Ti with GROMACS 2019, we obtain about 134 ns/day. On four GPUs, this goes up to 230 ns/day, which seems like a nice acceleration. However, switching to GROMACS 2021, one GPU will already yield 233 ns/day, and on four GPUs the additional performance is less than 50%. This operating point is thus a waste of resources.

NHR@FAU recommendation: Before increasing GPU count, use the newest version of GROMACS provided as a module and perform some benchmark tests to see if the resources are properly utilized. Read the full story.

Spotlight: PD Dr. Harald Lanig

Harald Lanig received his Dr. rer. nat. from the Institute of Physical Chemistry at the Julius-Maximilians University of Würzburg. After joining the research group of Prof. Dr. Tim Clark at the Computer-Chemie-Centrum at FAU, he finished his Habilitation in Pharmaceutical Chemistry in 2012. From 2013 to 2019 he was the Managing Director of the Central Institute for Scientific Computing (ZISC) at FAU and responsible for the ZISC's Research Data Management activities.

Currently, Harry is holding the position of an Executive Secretary at the NHR@FAU administrative office. He is responsible for the reviewing process of the compute time applications at NHR@FAU and for coordinating our activities with the NHR Alliance. His own research is well aligned with the application topic of NHR@FAU and focuses on the application of comparative modeling techniques, molecular dynamics simulations, and protein-ligand docking methods to investigate structure, properties, and function of biomolecules and their interactions. He coordinates the Liaison Scientists and contributes his own scientific expertise to their training, support and research activities.

June Highlight

The paper *ECM modeling and performance tuning of SpMV and Lattice QCD on A64FX* by C.L. Alappat, N. Meyer, J. Laukemann, T. Gruber, G. Hager, G. Wellein, and T. Wettig has just been accepted for publication in the Wiley journal *Concurrency and Computation: Practice and Experience*. In a collaboration with colleagues from the University of Regensburg, we modeled the performance of sparse matrix-vector multiplication and a Domain Wall kernel from Quantum Chromodynamics on the processor used in *Fugaku,* the world's fastest computer. Along the way, we could uncover interesting quirks in the processor architecture and useful insights for efficient code development. Get all the details at arXiv:2103.03013.

FAQ corner

How can I leverage node-local storage on TinyGPU to increase job performance?

Each node has at least 880 GB of local SSD capacity for temporary files under `$TMPDIR`. This directory will be deleted automatically as soon as the user has no jobs running on the node any more. Data to be kept can be copied to a cluster-wide volume at the end of the job.

How do I need to change my batch script to use the new TinyGPU nodes?

These nodes use SLURM as a batch system and run Ubuntu 20.04. SLURM automatically exports the environment of the submit host (woody) to the job. We recommend using the `sbatch` option `-export=none` to avoid this. Additionally, `unset SLURM_EXPORT_ENV` has to be done before srun to ensure that it is executed correctly.

Also, the SLURM nodes use different software versions and modules than the Torque nodes and the woody frontend. Applications should be recompiled on these nodes to ensure compatibility.

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