

Erlangen National High Performance Computing Center

Newsletter #3

August 2021

In This Issue...

- ISC 2021 Digital PhD Forum prizes
- LIKWID and OSACA Webinars
- GROMACS Shootout: Intel vs. NVIDIA
- Spotlight: Dr. Thomas Zeiser
- August Highlight: COVID-19 Research @ FAU

Quick News

New test cluster nodes available

In preparation for the new cluster installation later this year, the NHR@FAU test cluster now features two Intel Ice Lake nodes, one with 32-core and one with 36core CPUs, respectively. We have also added one node with AMD's new *Milan* (a.k.a. Zen3) CPUs, a new large-memory machine with 2 TiB of RAM, and a node with an AMD MI100 GPU.

New docs for R modules

Recent versions of R are available as modules on the Woody and TinyX clusters (using the Microsoft R Open distribution). For more information, please consult our new R Studio documentation page.

Docs on python and Jupyter notebooks

Our documentation has been augmented with info about using python, Conda, and Jupyter notebooks on our systems. Find the details here.

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.

Two NHR@FAU students win prizes at ISC 2021 Digital PhD Forum

ISC, the premier European conference on High Performance Computing and Data Analytics, includes a PhD Forum since 2015. PhD students get the opportunity to present their work with a poster and a "lightning talk"; during the poster session, the PhD Forum committee selects the best contribution(s) for attractive prizes. This year, the purely digital format of ISC 2021 Digital has forced many changes, but the basic rules of the PhD Forum were still the same. Two NHR@FAU students submitted a prerecorded talk and a poster, and both were accepted for the conference. Even better, they eventually won the first an second PhD Forum prizes.

The first prize went to Ayesha Afzal for her contribution "Noise-driven Clusterlevel Performance Modelling and Engineering." Ayesha conducts fundamental research on idle waves, computational waves, and desynchronization effects in massively parallel programs on multicore systems.

The second prize went to Dominik Ernst for his submission "The Best Thread Block Size and other parameters you have to tune for optimal performance on GPUs." His research revolves around performance modeling and optimal tuning parameter selection for GPUs based solely on high-level code features, which is instrumental for code optimization.

Webinars on LIKWID and OSACA for A64FX

The Research Division at NHR@FAU has recently developed considerable expertise in code performance modeling and optimization for the Fujitsu A64FX processor, the chip that powers Fugaku, the world's fastest supercomputer. The A64FX was developed specifically for HPC applications in an uncompromising co-design effort, and its performance rivals that of current accelerator devices. The LIKWID and OSACA tools, developed at NHR@FAU, were ported to the new architecture and can now be used by performance-aware developers. LIK-WID ("Like I Knew What I'm Doing") is a collection of command-line tools for topology inspection, thread/process affinity, hardware performance monitoring, hardware configuration, and microbenchmarking. OSACA ("Open-Source Architecture Code Analyzer") provides runtime predictions of assembly kernels for several x86- and Arm-based architectures. Christie Alappat, Thomas Gruber, Jan Laukemann, and Georg Hager conducted two events on using these tools for code optimization on the A64FX. Video recordings are available:

- Using LIKWID and OSACA for performance analysis on A64FX (Webinar for Brookhaven National Lab and the University of Delaware)
- LIKWID, OSACA, and Sparse Matrix-Vector Multiplication (SpMV) on the A64FX Processor (Open webinar for Stony Brook University)

GROMACS Shootout: CPU vs. GPU

GROMACS is in heavy use on our clusters, and getting good performance out of it is paramount for many of our customers. However, the question arises which hardware configuration is the most suitable for the MD simulation code. We ran several benchmarks relevant for FAU scientists with systems of various sizes on both Intel and AMD CPUs and several NVIDIA GPUs. The results can be summarized as follows: Scaling is good on CPUs, but comparing the amount of resources needed to solve a given problem, the best hardware choice for GROMACS is definitely a GPU. Moreover, the procurement and maintenance costs for the GPUs are significantly less than those for a large CPU cluster. Read the full story here.

Spotlight: Dr. Thomas Zeiser



Thomas heads the Systems & Services group of NHR@FAU and HPC4FAU since the end of 2020. He is one of the main driving forces of our transition from FAU-only to a national computing center and he spends a lot of his time with the procurement of new systems, adaptions of the infrastructure for our new building, and administration of software packages available on our clusters.

Thomas' background is in chemical engineering, where he developed lattice Boltzmann methods during his PhD work at Prof. Durst's chair of fluid mechanics. He

has been with us since 2004, and he is interested in large-scale simulations in general, the evaluation of new HPC hard- and software, but also the efficient operation of HPC systems. Thomas also implemented the first job-based monitoring for RRZE's HPC systems, which is still in use today.

August Highlight: COVID-19 Research at FAU

The paper Mutations in the B.1.1.7 SARS-CoV-2 Spike Protein Reduce Receptor-Binding Affinity and Induce a Flexible Link to the Fusion Peptide by E. Socher, M. Conrad, L. Heger, F. Paulsen, H. Sticht, F. Zunke and P. Arnold was recently published Open Access in the Journal Biomedicines special issue Conformational Dynamics of Viral Proteins. Under massive use of our



HPC facilities, researchers from several FAU institutes and the University Hospital investigated the molecular consequences of mutations within the named spike protein on internal dynamics and receptor binding, compared to the wild-type system. With their theoretical study, the authors help to understand how mutations change the properties of protein-receptor complexes and serve as a starting point for deeper biochemical and cell biological analyses. Read the whole story here.

FAQ Corner

What is SMT (also known as hyperthreading)?

Simultaneous multi-threading (SMT) allows a CPU core to run more than one software thread at the same time. These "hardware threads," a.k.a. "virtual cores" share almost all resources. The purpose of this feature is to make better use of the execution units within the core. It is rather hard to predict the benefit of SMT for real applications, so the best strategy is to try it using a well-designed, realistic benchmark case.

What is thread or process affinity?

Modern multicore systems have a strong *topology*, i.e., groups of hardware threads share different resources such as cores, caches, and memory interfaces. Many performance features of parallel programs depend on where their threads and processes are running in the machine. This makes it vital to bind these threads and processes to hardware threads so that variability is reduced and resources are balanced.

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