Friedrich-Alexander-Universität
Erlangen-Nürnberg

nhr.fau.de

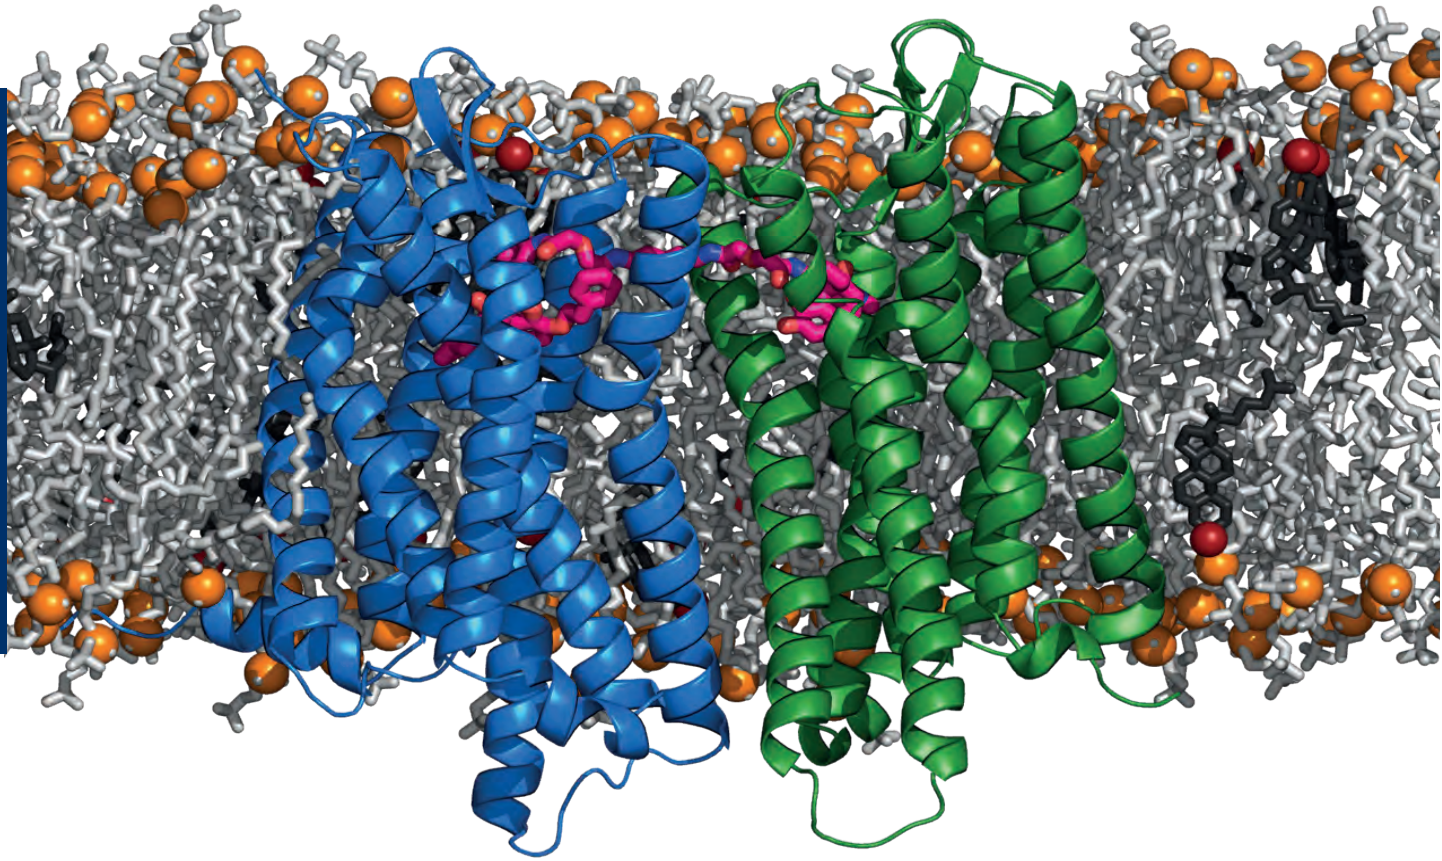

Erlangen National High Performance Computing Center (NHR@FAU)

Annual Report 2021

NHR FAU

Annual Report 2021

Erlangen National High Performance Computing Center (NHR@FAU)
Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)
c/o Erlangen Regional Computing Center (RRZE)

Martensstraße 1
91058 Erlangen
hpc-support@fau.de
`https://nhr.fau.de`

Contents

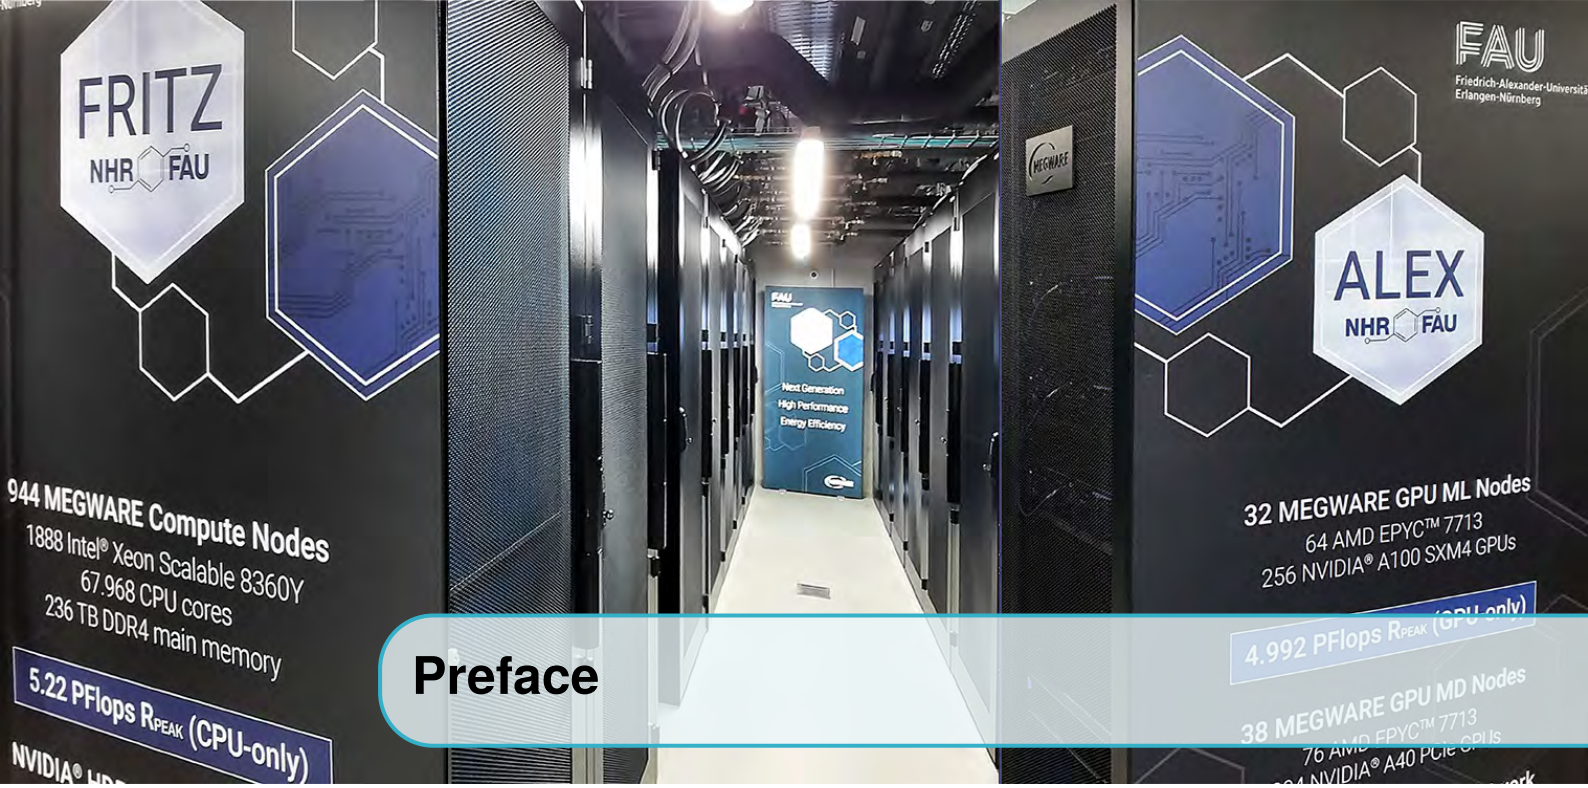

Preface

Dear Reader,

Prof. Dr. Gerhard Wellein, NHR@FAU Director

We are pleased that you take time to look at the 2021 Annual Report of the Erlangen National Center for High Performance Computing (NHR@FAU). In the face of the COVID-19 pandemic, 2021 proved to be another challenging year for all of us but it also offered some new opportunities. This was especially true for our new NHR@FAU center, which was formally established at FAU Erlangen-Nürnberg on January 1, 2021.

NHR@FAU continues the HPC services formerly provided by RRZE to FAU's researchers; it is also one out of eight national HPC centers at German universities (NHR centers), which were also established in January 2021. As an NHR center, we provide our services and expertise to scientists from all German universities and are a member of the NHR Alliance (Verein für Nationales Hochleistungsrechnen e.V.). Within the NHR Alliance, NHR@FAU leads the joint performance engineering (PE) activities and provides special application expertise in the field of atomistic simulations.

It is not surprising that 2021 was full of administrative work to implement adequate organizational infrastructures and to set the long-term foundations for our new national center. The main claim of NHR@FAU is to cover HPC across all involved fields, including HPC infrastructure design and operation, user support from level one to in-depth collaborations with developers and application scientists, HPC training, HPC research, and application expertise. To this end, we have established four divisions: *Systems & Services*, *Training & Support*, *Software & Tools*, and *Research*. The research-oriented structure of NHR@FAU is further strengthened by a group of liaison scientists. This report elaborates on the activities in these fields in detail; however, I would like to highlight some of them briefly upfront.

As the implementation of the NHR center did not hit us by surprise in January, our *Systems & Services* division had already started the procurement process for a new HPC system and initiated infrastructure upgrades in Q3–4/2020. Hence, we signed a contract with MEGWARE worth € 10 million in June 2021 to deliver a new HPC system consisting of a GPGPU module (now named *Alex*) and a CPU module (now named *Fritz*) by the end of 2021. At the same time, the extension of our cooling equipment was completed on time. Unfortunately, supply chain problems hit us during the installation of the HPC components. While *Alex* was able to go into test operation in December 2021, *Fritz* is still a little bit behind. However, we are confident that the 10 Pflop/s of aggregate double-precision peak performance of both systems will be fully at our customers' disposal in early 2022.

Although RRZE was known for its strong HPC training and support services, the NHR funding allows us to bring these activities to a new level and offer them on a national scale. In 2021, we started to extend our *Training & Support* division in to broaden our known expertise in several directions. A focus was put on GPGPU computing and the dissemination of our node-level performance engineering course by establishing a train-the-trainer concept for our partners in the NHR Alliance. Actually, our training activities benefited in 2021 from the COVID-19 pandemic forcing courses and tutorials into an online format: Regular communication channels such as the *HPC Café* or the monthly introduction for beginners have continuously received strong online attendance, and online courses saw attendees from all over the world who never would have gone through the hassle of long-distance travel for an in-person event. Furthermore, complex support cases are now regularly discussed with customers in online meetings, lowering the barrier for personal communication.

With the *Software & Tools* division we strengthened our expertise in supporting and developing tools and software for HPC. Most prominently, the team maintains our popular *LIKWID* tool suite, which is continuously adapted to new architectures and developed further with new features. Within the NHR Alliance, we also focus on the joint development of cluster-wide and job-specific monitoring solutions. In collaboration with other NHR centers, we advance our *ClusterCockpit* framework, which provides a complete monitoring stack and will be deployed on the new NHR@FAU clusters *Fritz* and *Alex*. Finally, our *Open-Source Architecture Code Analyzer* (OSACA) tool for analyzing in-core code performance is now part of the well-known Compiler Explorer at `godbolt.org` and has thus substantially increased its user base.

The *HPC research* division pools scientific activities at FAU related to the HPC focus topic of NHR@FAU. PhD students funded by third-party projects are currently pursuing research in the fields of performance modeling, hardware-efficient sparse solvers, or automatic code generation. Beyond many scientific contributions, our young researchers achieved the first and second place at the ISC 2021 Digital PhD Forum. Following our strategy of smooth transition from research to support, we could boost the GPGPU performance of a CFD solver by a factor of 80 within the EoCoE-II project.

In application support, NHR@FAU has a strong focus on atomistic simulations spanning across the fields of chemistry, life sciences, and materials science. Here, our concept of liaison scientists allows scientists all over Germany to benefit from the internationally recognized expertise of FAU's researchers. Together with NHR centers in Berlin and Paderborn, the *NHR Atomistic Simulation Center* (ASC@NHR) was established in 2021 to coordinate research, support, and training activities within the NHR Alliance.

Of course, the mission of any NHR center is to support scientists in breaking new ground in their fields. As *Fritz* and *Alex* were not in production operation in 2021, the application runs at NHR@FAU used our existing hardware and reflected the broad scientific scope at FAU with all faculties using the HPC systems. It is always hard to select the best or most important applications of the past year, so let me come back to the theme which again dominated 2021: the COVID-19 pandemic. In January 2021, GPGPUs were heavily used to investigate the dynamics of the spike protein of the corona virus, and in December 2021, massive GPGPU simulations were started on *Alex* to understand the mechanisms in mRNA vaccines. These challenging and complex simulations also reflected the need for scientific expertise, high-quality user support, and adequate hardware in order to respond to urgent scientific challenges in a timely manner.

Finally, on a national scope, NHR@FAU contributed to the establishment of the NHR Alliance and joint, regular NHR formats. Here the topic of performance engineering is especially important to us; we started and organized the public *NHR PerfLab Seminar* and the *PE Jour Fixe* together with the NHR centers in Aachen, Paderborn, and Berlin. This is only one example of how NHR provides the opportunity to actively engage in new collaborations with esteemed colleagues from other centers.

In summary, we feel that in 2021 the foundations for the next nine years of our NHR journey in Erlangen have been laid. This became only possible through our numerous enthusiastic colleagues from different disciplines who collaborated with each other, with our users, and with members of other NHR centers.

A heartfelt "thank you" goes to all who have contributed to the developments presented in this report!

Gerhard Wellein
on behalf of the NHR@FAU Executive Board

1. NHR@FAU at a glance

1.1 Organization

HPC is a key research priority at FAU. The use of HPC is central to numerous research activities across all of FAU's faculties, FAU's HPC-related research and teaching/training efforts are internationally recognized and FAU has continuously further developed its HPC infrastructure.

In 2021 these efforts have enabled FAU to establish the Erlangen National High Performance Computing Center (NHR@FAU) as one out of initially eight national centers for HPC (NHR centers) at German universities. The NHR centers receive federal and state fundings and their services are open to all German universities. NHR@FAU closely collaborates and coordinates with the other NHR centers and they jointly operate the NHR Alliance („NHR-Verein"). Besides fostering coordination and collaboration of the NHR centers, the NHR Alliance also supports the scientific computing and HPC-related research, e.g., by running the *NHR Graduate School*.

Complementing its national duties, NHR@FAU is also responsible for FAU's local HPC infrastructure and HPC services. These offerings, which have been provided by RRZE before 2021, are open to FAU researchers (and the regional RRZE-supported universities) only. Though being a separate organization within FAU, the NHR@FAU closely collaborates with RRZE in many directions, most notably in the areas of infrastructure maintenance and development, system administration, networks, and IdM-related services. Reflecting its national (Tier 2) and local (Tier 3) HPC service activities, the NHR@FAU receives annual budgets from the NHR program and the FAU.

NHR@FAU is also a pillar of the HPC activeness in the State of Bavaria. It complements the Tier-0/1 compute offerings and support services of the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and hosts one of the two KONWIHR offices. Within KONWIHR, NHR@FAU and LRZ jointly work with scientists at Bavarian universities to exploit the potential of the large computational power available in their research.

The organizational concept of NHR@FAU is to cover HPC across all involved fields including HPC infrastructure design and operation, user support from level one to in-depth collaborations with developers and application scientists, HPC training, HPC research, and application expertise. NHR@FAU aims to be the focal point of FAU's HPC activities to provide high-quality application and user support, up-to-date training and teaching offerings, and efficient and reliable compute capabilities to our users from Erlangen and all over Germany.

1.2 Executive board and divisions

The NHR@FAU director holds the professorship for HPC at the Department of Computer Science and was the lead principal investigator (PI) of the NHR@FAU proposal. Together with the other PIs and the technical director of RRZE, the center director forms the NHR@FAU *Executive Board* (see Figure 1.1), which meets regularly to discuss system operations, support and training activities, and budget planning.

NHR@FAU has four divisions: *Systems & Services*, *Training & Support*, *Software & Tools*, and *Research*. The research-oriented structure of our center is further supported by a group of *liaison scientists*, who establish a sustainable link between NHR services and ongoing research in relevant fields. The shared positions foster the idea of cross-sectional responsibilities and ensure that support requests meet the appropriate application expertise. They provide training, application consulting, and expert support to NHR@FAU; at the same time, they continue to engage in research activities within their communities.

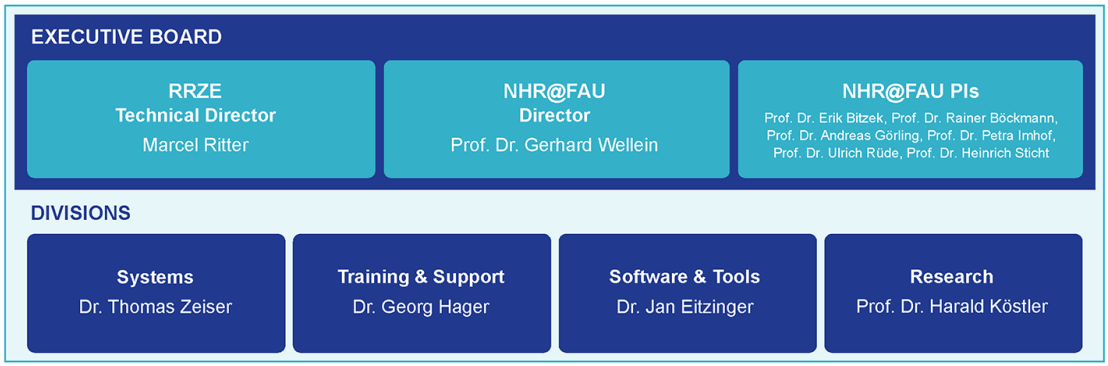

Figure 1.1: Executive board and divisions of NHR@FAU

1.3 Services

The main application topics covered by NHR@FAU are in the areas of atomistic simulations of chemical processes, biomolecular structures, and materials properties with classical and quantum-mechanical methods, which have applications in chemistry, life sciences, materials science, and beyond. FAU's unique computational chemistry, biology, and materials science community can provide comprehensive, application-specific competence and support for the full spectrum of classical and quantum-mechanical atomistic simulation methods and the most popular application codes. User consulting and expert support within NHR@FAU is provided for the full range of electronic structure techniques both for static calculations in chemistry, catalysis, and materials science.

The main HPC methodology focus is on node-level performance engineering for CPUs and GPUs. The internationally recognized model-driven performance engineering (PE) approach is the foundation of performance-related research, user consulting, training, performance monitoring, and all aspects of code optimization and parallelization. Research topics comprise novel performance models, PE concepts, efficient CPU/GPU implementation strategies, automatic generation of hardware-efficient codes, and micro-architectural analysis and benchmarking. The PE approach is applicable to a broad range of application fields. These activities are complemented by the development and maintenance of performance tools that support node-level PE.

As a core competence, NHR@FAU provides and further develops its internationally recognized expertise in node-level performance engineering for CPUs and GPUs, including performance analysis, analytic performance modeling, and resource-based performance monitoring. The pivotal guideline for support actions, services, and training activities is the PE approach. Based on hardware specifications and a performance analysis (runtime profile and inspection of hotspot code), a hypothesis is normally formulated about a potential hardware bottleneck at the main hotspot; this hypothesis should be connected with a performance expectation or even a full-fledged performance model.

Within the regular support process, PE activities are initiated in different ways: (1) detection of performance issues via job-specific performance monitoring, (2) a user's own performance analysis, (3) a user request for more effective use of a resource allotment, or (4) a user request for performance models of their own code in order to get better insight into performance issues.

The continuous job performance monitoring of relevant resource metrics (bandwidths, flop rates, network traffic, I/O utilization, etc.) that was already in place at RRZE has been adapted to NHR@FAU. In addition, the widely used LIKWID tool suite is instrumental in PE activities and a pivotal component in monitoring solutions. NHR@FAU continues to develop and adapt LIKWID to future CPU and GPU architectures. LIKWID support is provided as a service to the NHR Alliance and the entire HPC community.

Numerical methods for simulations focus on scalable iterative solvers, where FAU already had internationally recognized expertise on scalable hardware-efficient sparse solvers, modern matrix-free finite element multigrid methods, and lattice-Boltzmann solvers, including work on adaptive parallel mesh refinement with advanced mesh generation, load balancing, and visualization strategies. The work at NHR@FAU stands out because it integrates the whole method stack ranging from modeling to the hardware-aware and highly optimized implementations for extreme-scale computing on heterogeneous architectures. NHR@FAU has a large experience with building sustainable HPC software and is leading in automatic code generation techniques, where efficient parallel codes are automatically generated from abstract specifications.

Algorithmic user support is offered for these numerical methods, and expert support and training is provided in two directions: (1) manual implementation or automatic generation of hardware-efficient kernels and solvers for CPUs and GPUs and (2) defining and implementing maintainable and flexible numerical algorithms and libraries for full-scale applications. An important goal is to foster a software ecosystem of portable and inter-operable applications and libraries for the German HPC community.

In-depth courses on user code analysis and optimization are offered to staff from other centers ("train the trainer"). Advanced domain-specific HPC training complements the curriculum. As a first new contribution, a performance follow-on ("NLPE for Iterative Solvers") of the well-established lecture on "Iterative Solvers and Parallelization" is being developed. Monthly webinars with crash courses on using application codes and domain-specific tools at NHR@FAU are offered.

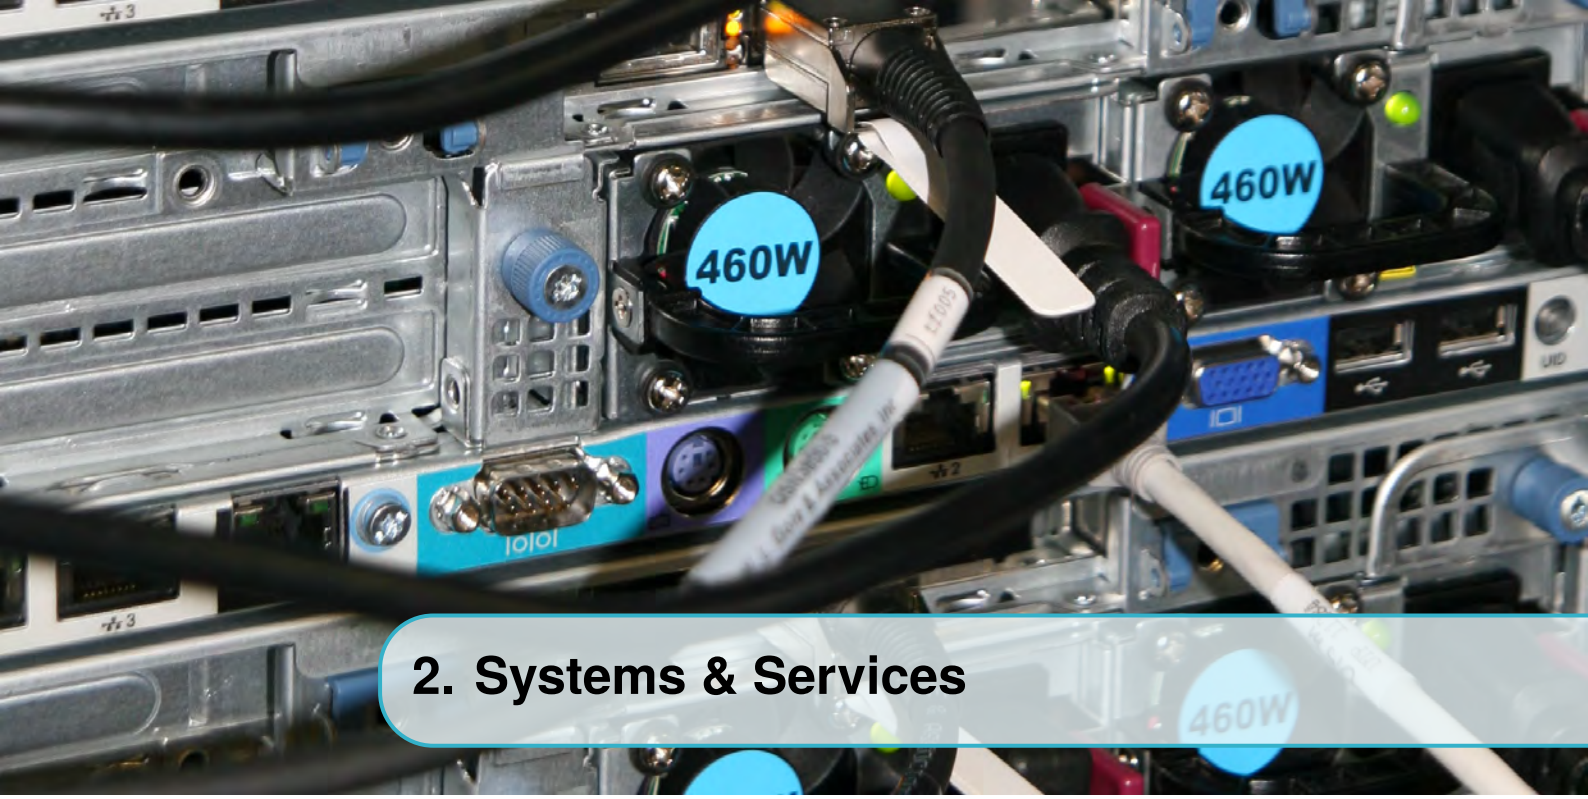

2. Systems & Services

Division head: *Dr. Thomas Zeiser*

The *Systems & Services* division at NHR@FAU is responsible for operation of the HPC systems (including procurement and software installation) and the required infrastructure (power and cooling). Many support activities are also carried out or supported by the *Systems & Services* division. The facility team within the *Systems & Services* division not only takes care of the current infrastructure but also drives the planning of a new data center building to be completed before the end of the decade.

RRZE has been operating Linux-based HPC clusters as Tier-3 compute resources for users of FAU and regional customers of RRZE since 2003. Centralisation of HPC systems and services has been well accepted and the HPC team grew over time. With the foundation of NHR@FAU in January 2021, the next level has been reached and all HPC activities and existing systems were technically taken over by NHR@FAU while still relying on the basic services of RRZE. For the Tier-3 operation of the HPC systems and services, dedicated funding is available. However, synergies of NHR and Tier 3 will be exploited wherever possible—resulting, among others, in the joint procurement and operation of systems and services.

The year 2021 was strongly characterized by the procurement of the high-performance computers *Alex* (GPGPU cluster) and *Fritz* (parallel computer). In Erlangen, HPC customers are traditionally involved closely in the design of the HPC hardware strategy and in all steps of the procurement process. Furthermore, the regular and intensive exchange of information with all leading HPC companies allows new technologies to be integrated at an early stage. In particular, the final selection of HPC systems based on a representative collection of benchmarks is labor-intensive but pays off in the form of high utilization and performance numbers for the systems. This time, the thematic orientation of NHR@FAU and, thus, potential application codes were already considered prominently in the benchmark selection.

Already in December 2020—and, thus, even before the official start of NHR@FAU—the EU-wide invitation to tender in the form of a "competitive dialogue" for the new NHR and Tier-3 flagships (parallel computers and GPGPU clusters) was launched; accordingly, only FAU people were involved here but not yet any (external) NHR users. In the dialog round, NHR@FAU talked intensively with five companies for a day each

about implementation possibilities to find an optimal solution. In April, four companies submitted a final bid. The bid submitted by the medium-sized company *MEGWARE Computer Vertrieb und Service GmbH* from Chemnitz was accepted.

Figure 2.1: Addition of a re-cooling system on the roof of the technology center of the Faculty of Natural Sciences.

With the award of the contract, the next phase of the procurement began: Various infrastructure tasks had to be addressed in the existing computer room of RRZE (renewal of the raised floor, cold water supply, power supply) and a new computer room still had to be created in the technology center of the Faculty of Natural Sciences. While the up to 200 kW of *Alex* in the existing computer room continue to be cooled with cold water via rear cooling doors on the racks, direct hot-water cooling on the processors is used for for 80% of the up to 700 kW of *Fritz*. Only the remaining 20% of network chips, power supplies, switches, etc., must be cooled via air.

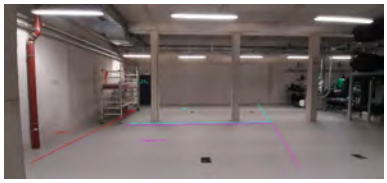

Figure 2.2: Empty space in the technology center of the Faculty of Natural Sciences, Spring 2021.

Unfortunately, the warm water of nearly 40°C cannot be further used in the existing buildings of FAU, so it is cooled down again on the roof of the building via hybrid dry coolers by free or adiabatic cooling. The "computer room" in the technology center is only a temporary solution until a new data center building is available. Two existing Jaeggi hybrid dry coolers were therefore extended by an identical model, so that a further use is possible after the HPC systems have been moved out in 10 years at the latest. With the construction of the new data center and further new buildings in the area of the Nikolaus-Fiebiger-Straße it should be possible to reuse the waste heat from the HPC systems by the end of the decade.

Figure 2.3: Preparations for the cooling system of the new parallel compute cluster, summer 2021.

In total, FAU has invested over eleven million euros in HPC systems and the associated infrastructure in 2021. The required funds came from large-scale equipment application to the German Research Foundation (DFG) under Article 91b of the German Basic Law (DFG INST 90/1171-1), NHR funding from the federal and local state governments, but also reinforcement funds from the Bavarian Ministry of Science, Research and Art (StMWK) and FAU itself.

Alex went into test operation with selected users right before the end of 2021. In the case of *Fritz*, the shortage of IT components and also other installation material (e.g., 63 A power distribution units) prevented the start of a test operation in 2021. As a result, as in the previous year, the *Meggie*, *Emmy*, and *Woody* clusters were the most important HPC resources at FAU in 2021—with users almost exclusively from FAU. *TinyGPU* is no longer tiny after various expansions in recent years and contributes significantly to HPC supply. In 2021, a significant increase in demand from machine learning could be perceived. Thanks to renewals of the network infrastructure in the HPC area, quite a few systems could already be connected with 100-Gbit Ethernet in 2021.

2.1 Compute resources

2.1.1 *Fritz* parallel computer (NHR & Tier 3)—to become operational in Q2/2022

`https://hpc.fau.de/systems-services/documentation-instructions/clusters/fritz-cluster/`

Fritz (system integrator: MEGWARE) is a high-performance compute resource with high-speed interconnect, i.e., a parallel computer. It is intended for moderately-sized, multi-node parallel workloads. *Fritz* is funded by NHR and DFG INST 90/1171-1. Thus, it will not only be the main resource for NHR projects at NHR@FAU but to a certain extent also as FAU's basic Tier-3 resource for high-end demand.

Figure 2.4: New parallel compute cluster *Fritz*, end of 2021.

The compute nodes of *Fritz* have been delivered in November/December 2021. However, owing to worldwide supply shortage, in particular the InfiniBand HCAs will only arrive in the first half of 2022. First early adopters will have access to *Fritz* from January 2022. The final configuration of *Fritz* will be:

- Four front end nodes with the same CPUs as the compute nodes but 512 GB of RAM, and 100 GbE connection to RRZE's network backbone.
- One visualization node with the same CPUs as the compute nodes but 1024 GB of RAM, one NVIDIA A16 GPU, 30 TB of local NVMe SSD storage, and 100 GbE connection to RRZE's network backbone.
- 944 compute nodes with direct liquid cooling (DLC), each with two Intel Xeon Platinum 8360Y *Ice Lake* CPUs (36 cores per chip) running at a base frequency of 2.4 GHz and 54 MB Shared L3 cache per chip, 256 GB of DDR4-RAM.
- Lustre-based parallel file system with a capacity of about 3.5 PB and an aggregated parallel I/O bandwidth of > 20 GB/s.
- Blocking fat-tree HDR100 InfiniBand with up to 100 GBit/s bandwidth per link and direction; there are islands with 64 nodes (i.e. 4,608 cores); the blocking factor between islands is 1:4.
- Operating system: AlmaLinux 8 (RHEL clone).
- Batch system: Slurm.

The LINPACK performance could not be measured yet due to lack of the InfiniBand HCAs. The direct liquid cooling of the processors and memory of the compute nodes ensures an efficient operation of *Fritz*, significantly lowering the operating costs.

2.1.2 *Alex* GPGPU cluster (NHR & Tier 3)—to become operational in Q1/2022

`https://hpc.fau.de/systems-services/documentation-instructions/clusters/alex-cluster/`

Alex (system integrator: MEGWARE) is a high-performance compute resource with NVIDIA GPGPU accelerators and partially high-speed interconnect. It is intended for single- and multi-GPGPU workloads, e.g., from molecular dynamics or machine learning. *Alex* is financed by NHR funds and DFG INST 90/1171-1. Thus, *Alex* serves both as NHR's project resource and as FAU's basic Tier-3 resource.

Most compute nodes of *Alex* were delivered in October/November 2021. First early-adopters have access to *Alex* since December 2021. The final configuration of *Alex* will be as follows:

- Two front end nodes, each with two AMD EPYC 7713 *Milan* CPUs (64 cores per chip) running at 2.0 GHz with 256 MB Shared L3 cache per chip, 512 GB of RAM, and 100 GbE connection to RRZE's network backbone but no GPGPUs.
- 20 GPGPU nodes, each with two AMD EPYC 7713 *Milan* CPUs (64 cores per chip) running at 2.0 GHz with 256 MB Shared L3 cache per chip, 1,024 GB of DDR4-RAM, eight NVIDIA A100 (each 40 GB HBM2 @ 1,555 GB/s; HGX board with NVLink; 9.7 TFlop/s in FP64 or 19.5 TFlop/s in FP32), two HDR200 InfiniBand HCAs, 25 GbE, and 14 TB on local NVMe SSDs.
- 12 GPGPU nodes, each with two AMD EPYC 7713 *Milan* CPUs (64 cores per chip) running at 2.0 GHz with 256 MB Shared L3 cache per chip, 2,048 GB of DDR4-RAM, eight NVIDIA A100 (each 80 GB HBM2 @ 1,555 GB/s; HGX board with NVLink; 9.7 TFlop/s in FP64 or 19.5 TFlop/s in FP32), two HDR200 InfiniBand HCAs, 25 GbE, and 14 TB on local NVMe SSDs; These systems are to be delivered in Q2/2022; five of these nodes belong to HS Coburg.
- 38 GPGPU nodes, each with two AMD EPYC 7713 *Milan* CPUs (64 cores per chip) running at 2.0 GHz with 256 MB Shared L3Cache per chip, 512 GB of DDR4-RAM, eight NVIDIA A40 (each with 48 GB DDR6 @ 696 GB/s; 37.42 TFlop/s in FP32), 25 GbE, and 7 TB on local NVMe SSDs.
- Native access to the Lustre-based parallel file system of *Fritz* through 25/100 GbE.
- A dedicated central NVMe storage will be added in 2022.
- Operating system: AlmaLinux 8 (RHEL clone).
- Batch system: Slurm.

On 160 NVIDIA A100/40GB GPGPUs, a LINPACK performance of 1.7 PFlop/s has been measured in January 2022. A new value (including the 12 additional nodes from 2022) will be measured for the June 2022 Top500 list. For molecular dynamics codes like GROMACS, an NVIDIA A40 GPGPU delivers a very similar performance than a much more expensive NVIDIA A100 GPGPU. Even for many machine learning workloads, the price/performance ratio of the NVIDIA A40 GPGPUs is more than competitive while, of course, applications requiring double precision calculations rely on the NVIDIA A100. The mixture of NVIDIA A40 and A100 allowed to maximize the overall cluster performance for a broad range of applications while still serving very different needs.

2.1.3 *Meggie* parallel cluster (Tier 3, from 2016)

Meggie (system integrator: MEGWARE) is a system that is designed for running parallel programs using significantly more than one node. It is intended for distributed-memory (MPI) or hybrid parallel programs with medium to high communication requirements and consists of 728 compute nodes (each with two Intel Xeon E5-2630v4 *Broadwell* CPUs with ten cores each running at 2.2 GHz, 64 GB of RAM, and OmniPath interconnect).

2.1.4 *Emmy* parallel cluster (Tier 3, from 2013)

Emmy (manufacturer: NEC) is a high-performance compute resource with high-speed interconnect. It has been intended for distributed-memory (MPI) or hybrid-parallel programs with medium to high communication requirements and now serves also for high-capacity throughput workload. Emmy consists of 560 compute nodes (each with two Xeon 2660v2 *Ivy Bridge* CPUs with 10 cores each running at 2.2 GHz, 64 GB of RAM, and QDR InfiniBand). Emmy also consists of a small accelerator partition with NVIDIA Tesla K20 and Tesla V100 GPUs. (The Intel Xeon Phi accelerators were shut down some years ago when Intel abandoned their software support.) Emmy is close to end of life.

2.1.5 *Woody* throughput cluster (Tier 3, multiple phases 2013–2019)

Woody is the preferred cluster for serial/single-node throughput jobs and has a long history. The nodes changed over time while the name has been kept. The currently more than 200 nodes all consist of single-socket nodes with rather high-frequency quad-core Intel Xeon E3-12xx processors (E3-1240 v3 *Haswell*; E3-1240 v5 *Skylake*; E3-1240 v6 *Kaby Lake*). All current nodes have 32 GB of RAM, 1 Gbit-Ethernet, and a local HDD/SDD.

2.1.6 *TinyGPU* cluster (Tier 3, multiple phases since 2009)

TinyGPU started in 2009 as a small special-purpose research test bed. Over the years, *TinyGPU* grew and no longer is that tiny. It now consists of more than 40 nodes with in total more than 200 NVIDIA GPGPUs of different generations (NVIDIA GeForce GTX1080, GeForce GTX1080Ti, RTX2080Ti, RTX3080, Tesla V100, A100). Almost all nodes have been financed by individual research groups of the whole university; RRZE or now NHR@FAU only takes care of the proper housing and operation.

2.1.7 *TinyFAT* cluster (Tier 3, 2016/2020)

TinyFAT is an other small special-purpose cluster. It consists of almost 50 dual-socket Intel *Broadwell* or AMD *Rome* nodes with 256 or 512 GB of main memory, and local SSD storage. By now 512 GB are no longer huge, but compared to the 32/64 GB of most other clusters, it is still large.

2.2 Storage resources (Tier 3 with dedicated extensions for NHR)

NHR@FAU operates several storage systems of different quality but also price/performance serving different needs:

- All throughput and GPGPU/large-memory nodes have local HDDs/SSDs as fast job-local storage.
- Each parallel computer has a parallel file system which is dedicated to that cluster; there are no capacity quotas enforced but high-watermark deletion is applied. The main usage is for checkpoint-restart files.
- `$WORK` consists of several capacity-optimized NFS file servers (more than 2 PB in total); there is no or only very limited backup. These Linux servers can cope rather well with many small files.
- `$VAULT` is a high-quality file system where especially larger files can be kept for longer times; `$VAULT` has a total capacity of 6 PB in a single file system and is served by an IBM Spectrum Scale cluster consisting of six servers (two NDS and

served by an IBM Spectrum Scale cluster consisting of six servers (two NDS and four CES nodes; about 700 NL-SAS HDDs plus some SSDs in hardware RAIDs provide the capacity).

- `$HOME` is served by the same IBM Spectrum Scale cluster as `$VAULT`. The main difference is the frequency of snapshots (every 30 minutes vs. once per day).
- For long-term offline storage, an IBM TS4500 tape library with currently eight LTO8 tape drives and two expansion frames for up to 3,370 tapes is available. Archiving data is a manual process.

By the end of 2021, most file servers have 100 Gbit-Ethernet connections.

2.3 System usage 2021

As the NHR systems *Alex* and *Fritz* only arrived late in 2021, there were mainly Tier-3 users on the pre-NHR systems of NHR@FAU during the year 2021.

Table 2.1: Usage shares of the different scientific domains on the different Tier-3 compute segments of NHR@FAU not including *Alex* and *Fritz*; *TinyFAT* is included in the throughput clusters.

scientific domain	parallel	throughput	GPUs
chemistry + molecular dynamics	49%	20%	74%
physics, earth science + climate	11%	32%	1%
engineering (fluids, mechanical, materials)	34%	8%	3%
mathematics	1%	38%	–
machine learning	3%	2%	21%
others	2%	0%	1%

There were no major outages in the year 2021. In total, more than 180 million core-hours have been delivered on the Tier-3 CPU systems and more than 0.8 million GPU-hours have been consumed on the Tier-3 GPU nodes of *TinyGPU*. Early adopters used in addition more than 70k GPU-hours on *Alex* during the last few weeks of 2021.

3. Training & Support

Division head: *Dr. Georg Hager*

The *Training & Support* division at NHR@FAU is responsible for HPC training efforts and all support activities beyond level two. This entails training and event coordination, conducting local courses and user engagement events, organizing invited talks, participation in and organization of external training, application-specific support, performance engineering of user code, expert support for algorithms and libraries, and participation in third-party funded research projects. Due to the cross-cutting nature of many of these activities, members of the *Training & Support* division are typically operative in other divisions as well.

The RRZE HPC group, out of which NHR@FAU has emerged, has a long history of teaching, training, and support in HPC. As early as 1998, scientists from FAU were educated in programming and performance optimization for the then-current vector supercomputers VPP700 at LRZ Munich and its smaller satellite system, the VPP300 at RRZE. In 2000, the installation of the federal "HLRB" system at LRZ and the concomitant KONWIHR funding fostered the intensification of training and user support activities. Over the years, the RRZE HPC group could build up considerable expertise in parallel programming, performance modeling, and optimization, which led to the development of signature lectures and tutorials. This process was facilitated by a constant and lively mutual exchange among research, user consulting, teaching, systems administration, and training staff. Since 2019, the monthly *HPC Café* and *HPC in a Nutshell* events have been strengthening the communication with new and experienced HPC customers. Since 2021, with the advent of the NHR program, the spectrum of training offers could be broadened even further.

The *Training & Support* division is further responsible for all HPC user support that requires intense interaction with the customer, such as software configuration, code porting, parallelization, performance analysis and optimization, and the selection of appropriate algorithms and libraries. These services are not only offered to NHR users but also to third party-funded projects conducted by the NHR@FAU Research division and its project partners.

3.1 Training activities

3.1.1 Courses and tutorials

In the following we list courses and tutorials that were conducted in 2021 and were either organized by NHR@FAU, conducted by others with substantial contributions from NHR@FAU, or performed on invitation by NHR@FAU (see Table 3.1). All these events were conducted online via Zoom.

- *Node-Level Performance Engineering* held by Georg Hager, Gerhard Wellein, and Jan Eitzinger. This course covered performance engineering approaches on the compute node level. Even application developers who are fluent in OpenMP and MPI often lack a good grasp of how much performance could at best be achieved by their code. This is because parallelism takes us only half the way to good performance. Even worse, slow serial code tends to scale very well, hiding the fact that resources are wasted. This course conveyed the required knowledge to develop a thorough understanding of the interactions between software and hardware. This process must start at the core, socket, and node level, where the code gets executed that does the actual computational work. We introduced the basic architectural features and bottlenecks of modern processors and compute nodes. Pipelining, SIMD, superscalarity, caches, memory interfaces, ccNUMA, etc., are covered. A cornerstone of node-level performance analysis is the Roofline model, which was introduced in due detail and applied to various examples from computational science. We also showed how simple software tools can be used to acquire knowledge about the system, run code in a reproducible way, and validate hypotheses about resource consumption. Finally, once the architectural requirements of a code were understood and correlated with performance measurements, the potential benefit of code changes can often be predicted, replacing hope-for-the-best optimizations by a scientific process.

- *Hybrid Programming tutorial* given and organized, among others, by Georg Hager on June 15–17, 2021. Most HPC systems are clusters of shared memory nodes. To use such systems efficiently both memory consumption and communication time has to be optimized. Therefore, hybrid programming may combine the distributed memory parallelization on the node interconnect (e.g., with MPI) with the shared memory parallelization inside of each node (e.g., with OpenMP or MPI-3.0 shared memory). This course analyzed the strengths and weaknesses of several parallel programming models on clusters of SMP nodes. Multi-socket-multi-core systems in highly parallel environments were given special consideration. MPI-3.0 has introduced a new shared memory programming interface, which can be combined with inter-node MPI communication. It can be used for direct neighbor accesses similar to OpenMP or for direct halo copies, and enables new hybrid programming models. These models were compared with various hybrid MPI+OpenMP approaches and pure MPI. Numerous case studies and micro-benchmarks demonstrated the performance-related aspects of hybrid programming. Hands-on sessions were included on all days. Tools for hybrid programming such as thread/process placement support and performance analysis were presented in a "how-to" section. This course provided scientific training in Computational Science, and in addition, the scientific exchange of the participants among each other.

Table 3.1: Courses and tutorials in 2021. "(e)" denotes an invited teacher from outside FAU.

Event	Date(s)	Place	Comment
Node-Level Performance Engineering	Mar 10–12	VSC, TU Wien	G. Hager, G. Wellein
	Jul 12–14	HLRS	G. Hager, J. Eitzinger
	Dec 1–3	LRZ	G. Hager, G. Wellein
Node-Level Performance Engineering, train-the-trainer edition	Oct 11–13	NHR@FAU	G. Hager, J. Eitzinger, T. Gruber
Introduction to C++ for beginners	Mar 15–19	NHR@FAU	K. Iglberger (e)
Modern C++ Software Design	Sep 29 to Oct 1	NHR@FAU	K. Iglberger (e)
Hybrid Programming tutorial	June 15–17	TU Wien	G. Hager, R. Rabenseifner [HLRS] and C. Blaas-Schenner [TU Wien]
Parallel Programming of High-Performance Systems (PPHPS21)	April 13–15	NHR@FAU	G. Hager, V. Weinberg[LRZ]
VI-HPS Online Tuning Workshop	March 1–3	NHR@FAU	G. Hager
IHPCSS	July 28	online	D. Ernst
2021 Code Performance Series: From analysis to insight	July 15	Durham University	T. Gruber, G. Hager
LRZ/Intel Code Optimization Workshop	Nov 2–4	online	T. Gruber
EXA2PRO-EoCoE joint workshop	Feb 23	Paris	T. Gruber, S. Kuckuk, M. Holzer
LIKWID, OSACA, and Sparse MVM on A64FX	July 27	IACS, Stony Brook University	C. Alappat, J. Laukemann, T. Gruber, G. Hager
Using the LIKWID and OSACA tools on A64FX	June 02	NHR@FAU	T. Gruber
LIKWID	Aug 31	HPC.NRW	T. Gruber

- The course *Parallel Programming of High Performance Systems* was given by Georg Hager in collaboration with LRZ on April 13–15, 2021. This online course was targeted at students and scientists with interest in programming modern HPC hardware, specifically the large scale parallel computing systems available in Jülich, Stuttgart and Munich, but also smaller clusters in Tier-2/-3 centers and departments.

- *VI-HPS Online Tuning Workshop* organized by Georg Hager on March 1–3, 2021. In this particular event, the tools TAU , MAQAO, Score-P, Paraver/Extrae/Dimemas, and Extra-P were covered. On completion, participants were familiar with common performance analysis and diagnosis techniques and how they can be employed in practice.

- Lecture on *GPU Performance Analysis* given by Dominik Ernst at the International HPC Summer School (IHPCSS) on July 28, 2021. The summer school familiarized the best students in computational sciences with major state-of-the-art aspects of HPC and Big Data Analytics for a variety of scientific disciplines, catalyze the formation of networks, provide advanced mentoring, facilitate international exchange and open up further career options.

- A session on *Single-Node optimization* at the *2021 Code Performance Series: From analysis to insight* was held by Thomas Gruber and Georg Hager on July 15, 2021.

- A session on LIKWID was given by Thomas Gruber at the *Code Optimization Workshop* organized by LRZ in cooperation with Intel and NHR@FAU. This workshop was held from November 2–4, 2021 and answered questions on application performance, the maximum speed-up achievable on the architecture, and whether the implementation is matching the HPC objectives. The workshop provided a unique opportunity to learn techniques, methods and solutions on how to improve code, how to enable the new hardware features and how to use the Roofline model to visualise the potential benefits of an optimisation process.

- *Introduction to C++ for beginners* held by Klaus Iglberger, March 15–19, 2021. The focus of the training was on the introduction of the essential language features and the syntax of C++. Additionally, it introduced many C++ software development principles, concepts, idioms, and best practices, which enable programmers to create professional, high-quality code from the very beginning.

- Hands-on session about *Performance Engineering and code generation techniques* at the EXA2PRO-EoCoE joint workshop on February 22–24, 2021. The workshop showcased the participants' respective, breakthrough work in the field of computer science.

- Several talks were given about LIKWID. LIKWID stands for "Like I Knew What I'm Doing." It is an easy-to-use yet powerful command-line performance tool suite for the GNU/Linux operating system. While the focus of LIKWID is on x86 processors, some of the tools are portable and not limited to any specific architecture. For the upcoming release, LIKWID has been ported to ARMv7/v8 and POWER8/9 architectures as well as for Nvidia GPU co-processors. The LIKWID tools are used worldwide for teaching purposes or in production environments like NERSC at Lawrence Berkeley National Labratory, CSCS (Swiss National Supercomputing Center) in Lugano, the National Super Computer Center in Guangzhou, the

Vienna Scientific Cluster (VSC) or the Barcelona Supercomputing Center (BSC). Moreover, LIKWID is used by the Gauss Center for Supercomputing at LRZ (Garching), HLRS (Stuttgart), and JSC (Jülich), and by the IT provider for the Max Planck Society (MPCDF). Some of the Tier-2/-3 HPC centers in Germany using LIKWID are RWTH Aachen, TU Dresden, KIT Karlsruhe, University of Paderborn, University of Konstanz, University of Gießen, and the national research center DESY in Hamburg. Dr. Hatem Ltaief from the Extreme Computing Research Center at King Abdullah University of Science and Technology (KAUST) says "LIKWID rocks"! We thank the LANL HPC ENV team to test and use LIKWID on the Darwin cluster.

3.1.2 HPC Café

The *HPC Café* was initiated in late 2019 as a new way to foster and facilitate communication between the RRZE HPC group and HPC customers. It also complements the traditional HPC support channels by providing a deliberately informal setting. Every second Tuesday of a month, all HPC customers are invited to a freestyle Q&A session, followed by a short presentation about a "focus topic" of interest. Since April of 2020, the *HPC Café* has been conducted entirely as an online event. Whenever possible, talks are recorded and published on the FAU video portal and/or the NHR@FAU YouTube channel (youtube.com/NHRFAU). The following focus topics have been covered in 2021:

- December 14, 2021: *Python beyond the basics: Numpy, Scipy, Matplotlib* by Jan Hönig, LSS.
- November 9, 2021: *Effective Editing With Vim* by Jan Eitzinger.
- October 19, 2021: *Howto on using the Cx services based on the RRZE Gitlab instances* by Thomas Gruber.
- September 14, 2021: General Q&A, some advice on using (and not misusing) the file systems, current state of NHR system installation by NHR@FAU staff.
- July 13, 2021: *Current status NHR@FAU resources and KONWIHR projects* by Thomas Zeiser and Katrin Nusser.
- June 8, 2021: *Build systems and "Make"* by Jan Eitzinger.
- May 11, 2021: *"Continuous x" (Cx) for HPC Systems*. Guest talk by Jennifer Buchmüller and Terry Cojean, KIT.
- April 13, 2021: *Julia in HPC*. Guest talk by Valentin Churavy, MIT CSAIL Julia Lab.
- March 9, 2021: *Git part 2: advanced features and workflows* by Julian Hammer and Thomas Gruber.
- February 9, 2021: *Git: basics, common workflows, and tips and tricks* by Jan Eitzinger.
- January 12, 2021: *AI-assisted research at FAU*. Four FAU researchers give short talks about how their projects benefit from Artificial Intelligence (AI) methods and what resources they require for it.

A complete list with all events can be found at `https://hpc.fau.de/systems-services/support/hpc-cafe/`.

3.1.3 HPC in a nutshell

Each month, one day after the HPC Café, a two-hour online introduction into NHR@FAU systems, access formalities, and basic usage was provided. This well-received format aims at reducing the entry barrier for new and inexperienced users. The content is continuously updated to reflect recent changes in NHR@FAU systems and access rules.

3.2 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 talks about those topics, and more. In 2021, NHR@FAU has taken on the role of organizing the online seminar talks, speaker acquisition, and pre- and post-event dissemination. If possible, all material (e.g., slides, video recordings) were published via the NHR@FAU website and YouTube channel. The following talks have been given in 2021:

- December 15, 2021: *Memory Bandwidth and System Balance in HPC Systems–2021 Update* by John D. McCalpin, Texas Advanced Computing Center.
- November 23, 2021: *General-purpose GPU Hashing Data Structures and their Application in Accelerated Genomics* by Daniel Jünger and Bertil Schmidt, University of Mainz.
- November 2, 2021: *On-demand file systems: use cases and lessons learned* by Mehmet Soysal, Karlsruhe Institute of Technology.
- October 19, 2021: *Discontinuous Galerkin Shallow-Water Simulations on FPGAs* by Tobias Kenter, University of Paderborn.
- October 5, 2021: *Towards exascale simulations using the submatrix method and approximate computing* by Thomas Kühne, University of Paderborn .
- September 21, 2021: *Performance Measurements in HPC: Techniques, Side Effects, and Use* by Christian Iwainsky, TU Darmstadt.
- July 20, 2021: *Exclusive file systems for power users with BeeGFS and network NVMe storage* by Sebastian Oeste, ZIH Dresden.
- June 22, 2021: *Tasking in OpenMP 5.0* by Christian Terboven, RWTH Aachen University IT Center).
- June 1, 2021: *ALTO: Adaptive Linearized Storage of Sparse Tensors* by Jan Laukemann, Intel Labs.
- May 18, 2021: *A Modular Precision Ecosystem Based on a Memory Accessor* by Hartwig Anzt, Karlsruhe Institute of Technology.
- May 4, 2021: *Automatic Generation of Models of Microarchitectures* by Andreas Abel, Saarland University.

- March 23, 2021: *Parallel Selection on GPUs* by Tobias Ribizel, Karlsruhe Institute of Technology.
- March 16, 2021: *On the Interaction of Memory, Application and Operating System on a Heterogeneous Memory System* by Steffen Christgau, Zuse Institut Berlin.
- March 9, 2021: *The STX Processor—Hardware and Programming* by Franz-Josef Pfreundt and Jens Krüger, Fraunhofer ITWM.
- February 23, 2021: *A closer look at the Fujitsu A64FX processor* by Georg Hager, NHR@FAU.

3.3 Projects

3.3.1 EoCoE-II

Many pressing questions about the future global energy supply lead to highly complex scientific problems that are increasingly being researched with the help of simulations on supercomputers. The scientific topics range from photovoltaics to the use of geothermal energy or the design of wind farms to plasma physics for the possible future use of fusion energy. Simulations replace and complement expensive and lengthy experiments. Together with several partners from different European countries (France, Germany, Italy, Belgium, Great Britain, Spain, Poland), a continuation application to the expiring EU project "EoCo" (Energy oriented Centre of Excellence) was submitted in 2018, funded by the "Horizon 2020" project (`https://www.eocoe.eu`). At FAU, in addition to the NHR@FAU, the Chair of System Simulation (Prof. Rüde) is also involved. After approval of the project application and during 2019, the HPC group has started to support the application developers from the other project parts in the area of performance engineering. This includes, among others, the organization of courses and tutorials, but also "hackathons" where project collaborators can apply advanced performance analysis techniques using their own simulation codes.

Since no business trips could take place in 2021, it was again impossible to hold supervision-intensive hackathons with about 20–30 participants. Thus, these activities were split into smaller events and much of the collaboration took place in small groups. In addition to the "flagship code" Alya for simulating multi-physics problems, work now focused on GYSELA, a code for simulating plasma turbulence in fusion reactors, and EURAD-IM, a program for predicting and analyzing air quality. In the case of Alya, GPU optimization strategies developed using a simple code version ("proxy app") could be applied to many parts of the production code, resulting in a massive speedup. EURAD-IM and GYSELA benefited from detailed performance analysis and code quality improvement, especially with respect to vectorization. The EoCoE-II project expired in 2021; an extension as EoCoE-III is expected to be requested in 2022.

3.3.2 KONWIHR

The main objective of KONWIHR is foster the efficient the use of high performance computers and to increase and broaden their impact on research. KONWIHR support scientists at Bavarian colleges and universities in adapting and developing

their numerical applications or other data- or computation-intensive codes for modern parallel computer architectures. Funding is based on the duration of the project and can amount to a maximum of € 10,000 (small project of three months) or up to € 50,000 for projects with a duration of twelve months. The projects must be carried out in close cooperation with one of the two computing centers (LRZ and NHR@FAU). In particular, the funded project staff member should spend a longer period at one center.

NHR@FAU receives KONWIHR funding to coordinate KONWIHR activities in the northern part of Bavaria and to support KONWIHR projects in optimizing and adapting their codes and workflows for HPC systems. KONWIHR is currently led by Prof. Bungartz (TU Munich) and Prof. Wellein (FAU).

In 2021, KONWIHR granted the following project applications:

- *Semiempirische Molekülorbital-Theorie für sehr große Systeme II*, Prof. Clark (FAU), large project, NHR@FAU.
- *Largescale data processing and mass production of cosmic-ray background simulations for H.E.S.S.*, Prof. van Eldik (FAU), small project, NHR@FAU.
- *Greedy algorithms for fair allocations and efficient assignments within facility location optimization problems*, Dr. Singh (FAU), small project, NHR@FAU.
- *MRzero – Scaling of a differentiable end-to-end optimization for magnetic resonance imaging*, Prof. Zaiss (FAU), large project, NHR@FAU.
- *Optimisation of SeisSol for Large Scale Simulations of Induced Earthquakes*, Prof. Bader (TU Munich), large project, LRZ.
- *Efficient Remeshing in preCICE*, Dr. Uekermann (TU Munich), large project, LRZ.
- *Fast and scalable finite element algorithms for coupled multiphysics problems and non-matching grids*, Dr. Kronbichler/Prof. Wall (TU Munich), large project, LRZ.

In addition, workshops have been organized to allow the KONWIHR projects to present their work and discuss with other projects and with the expert personnel from NHR@FAU and LRZ:

- March 18, 2021: KONWIHR workshop for projects from 2020-2.
- October 11, 2021: KONWIHR workshop for projects from 2021-1.

Among the numerous consulting activities in 2021 for ongoing KONWIHR projects we would like to highlight a representative example: In a collaboration with theoretical physicists from the University of Würzburg, the performance of the Metropolis update in the MARQOV software package (`https://gitpages.physik.uni-wuerzburg.de/marqov/webmarqov/`) has been improved by 50% by optimizing the evaluation of the exponential operator.

3.4 Dissemination

Dissemination of activities and results was a vital part of the NHR@FAU activities in 2021. We concentrated predominantly on our newsletter and a limited number of social media channels; however, it is planned to heavily extend public relations and visibility measures in the years to come.

3.4.1 NHR@FAU newsletter

The bi-monthly NHR@FAU newsletter is a publication that summarizes recent and upcoming events at NHR@FAU and the NHR association. Beyond events, every newsletter highlights special or noteworthy achievements by NHR@FAU scientists and an FAQ corner. Finally, the "newsletter spotlight" introduces an NHR@FAU scientist in some detail. In 2021, five newsletters have been published.

A mailing list has been set up to alert about the publication of a new issue. It is open for everyone to subscribe at: `https://lists.fau.de/cgi-bin/listinfo/nhr-newsletter`lists.fau.de/cgi-bin/listinfo/nhr-newsletter.

3.4.2 Social media

In order to widen its reach to potential customers and audiences, NHR@FAU leverages social media channels. Currently, we concentrate mostly on dissemination via recorded lectures, seminars, and tutorials, which are published either on the NHR@FAU YouTube channel (`https://www.youtube.com/NHRFAU`) or the FAU video portal (`https://www.fau.tv/course/id/1146`). Our recordings of parallel programming lectures are particularly well received and generate considerable feedback in the HPC community and among students.

4. Software & Tools

Division head: *Dr. Jan Eitzinger*

The *Software & Tools* division at NHR@FAU is responsible for the development of Open Source software and for providing and administrating services for NHR@FAU and the NHR alliance. To this end, the division develops, maintains and supports Open-Source software that is published on our GitHub account (`https://github.com/RRZE-HPC`): LIKWID performance tools, Open Source Architecture Code Analyzer (OSACA), Loop Kernel Analysis and Performance Modeling Toolkit (Kerncraft), The Bandwidth Benchmark, MachineState, and the ClusterCockpit Monitoring Framework.

The division participates in performance and monitoring tool related third party projects, develops and maintains software used internally at NHR@FAU (User Portal), and provides and administrates services (NHR Moodle, NHR@FAU User Portal, NHR@FAU Job specific Monitoring for all cluster systems).

We are member of the Virtual Institute - High Productivity Supercomputing (VI-HPS) and the SPEC Research Group. We contribute our expertise in performance and monitoring tools and software development in the NHR central projects *Standards and interfaces for a system-wide job-specific performance monitoring* (project lead) and *Cx as a service for sustainable HPC research software engineering*.

The RRZE HPC group, out of which NHR@FAU has emerged, already started in 2009 to develop the LIKWID Performance Tool Suite. From 2013 to 2016 it had the project lead in the BMBF FEPA project, that developed a system-wide job-specific monitoring infrastructure. From 2017 to 2020 the group was a partner in the BMBF Metacca project and contributed its LIKWID and Kerncraft tools. Also from 2017 to 2020 the RRZE HPC group was the project lead in the DFG ProPE project, where among other activities the initial development for the ClusterCockpit Monitoring Framework began. ClusterCockpit is now besides LIKWID the second large Open-Source project at NHR@FAU.

In 2021, we submitted the BMBF proposal EE-HPC to further strengthen our activities in cluster-wide performance and energy monitoring. NHR@FAU is the project lead with further partners HLRS Stuttgart, RWTH Aachen, Deutsches Klimarechenzentrum, and HPE. The project proposal was accepted and the project will start in autumn 2022. In 2021, a total of four persons were working in the *Software & Tools* division.

4.1 Tool development

The tool development at NHR@FAU is well integrated with our main research topics of performance engineering and performance modeling. LIKWID and OSACA are important tools used in many of our research and support projects and enable us to adopt and research novel processor architectures at a very early stage. With Kerncraft (developed by Julian Hammer) the application of analytical performance models (the well-known Roofline model and the ECM model, which was developed at NHR@FAU) is possible also for non-expert users. The ClusterCockpit monitoring stack bundles our activities in cluster-wide monitoring solutions.

4.1.1 LIKWID performance tool suite

LIKWID (`https://github.com/RRZE-HPC/likwid`) is an easy-to-use yet powerful command line performance tool suite for the GNU/Linux operating system. LIKWID is maintained by Thomas Gruber.

Currently, LIKWID consists of seven core tools of which `likwid-perfctr` (counting hardware performance events), `likwid-topology` (display node topology), `likwid-pin` (control thread and process affinity), and `likwid-bench` (microbenchmarking framework) are most prominent. The hardware-architecture-specific tools support x86-64, ARM, and Power processors and NVIDIA GPUs.

LIKWID is by far our most popular open-source project, and its tools are used worldwide for research, teaching, or in production environments like NERSC at Lawrence Berkeley National Laboratory, CSCS (Swiss National Supercomputing Center) in Lugano, the National Super Computer Center in Guangzhou, the Vienna Scientific Cluster (VSC) or the Barcelona Supercomputing Center (BSC).

Moreover, LIKWID is used by the members of the Gauss Center for Supercomputing—LRZ Garching, HLRS Stuttgart, and JSC Jülich—and at the IT provider for the Max Planck Society (MPCDF). Some of the Tier-2/3 HPC sites in Germany using LIKWID are RWTH Aachen, TU Dresden, KIT Karlsruhe, University of Paderborn, University of Konstanz, University of Gießen, and the national research center DESY in Hamburg.

In 2021, LIKWID was endowed with support for the latest CPU and NVIDIA GPU architectures. It is now available in version 5.2.1. At various VI-HPS tuning workshops, LIKWID was introduced to a wide audience including one workshop at CINECA (Italy) on their MARCONI100 system equipped with POWER9 CPUs and NVIDIA GPUs.

4.1.2 OSACA Open Source Architecture Code Analyzer

OSACA (`https://github.com/RRZE-HPC/OSACA`) is a tool that can analyze x86 and Arm64 assembly code and produce runtime predictions assuming a steady-state execution and all data in the L1 cache. By detecting data dependencies, OSACA provides not only a throughput prediction as the best-case scenario, but also the critical path and loop-carried dependencies for kernel snippets. OSACA is maintained by Jan Laukemann.

A tool like OSACA is needed for analytic performance modeling, e.g., to formulate ECM or refined Roofline models. While there exist other tools like the Intel Architecture Code Analyzer (IACA), LLVM's Machine Code Analyzer (LLVM-MCA) or The uops.info

Code Analyzer (uiCA), they lack accuracy in prediction, are not open source, discontinued in development, or do not provide support for non-x86 architectures. However, OSACA can also handle all modern Intel and AMD x86 architectures and several Arm processors. Besides existing as a Python3 module and CLI application, it is integrated into the Compiler Explorer at `https://godbolt.org`, which allows using OSACA from a browser without any installation.

In 2021, we enhanced the supported architectures by integrating in-core performance models for several Arm-based chips like the A64FX microarchitecture or the Cortex A72 used in the Raspberry Pi computers. Furthermore, we provided support for the state-of-the-art Intel Ice Lake server microarchitecture.

4.1.3 Kerncraft loop kernel analysis and performance modeling toolkit

Kerncraft is a loop kernel analysis and performance modeling toolkit. It allows automatic analysis of loop kernels using the Execution Cache Memory (ECM) model and the Roofline model, and their validation via actual benchmarks. Kerncraft provides a framework for investigating the data reuse and cache requirements by static code analysis.

In combination with the Intel IACA tool or our OSACA tool, Kerncraft can give a good overview of both in-core and memory bottlenecks and use that data to construct predictive, white-box performance models. Kerncraft was developed and maintained by Julian Hammer, who has left FAU; we currently search for funding and a new maintainer for further development of Kerncraft.

4.1.4 ClusterCockpit monitoring framework

Figure 4.1: Screenshots of the ClusterCockpit web interface. Left: job list of running jobs; right: single-job view showing the resource utilization in a polar plot and diagnostic Roofline plot.

ClusterCockpit (`https://github.com/ClusterCockpit`) is a full-stack framework for job-specific performance monitoring on HPC clusters. It started as a side project for a web frontend during the DFG ProPE project and has grown into a cooperative effort to provide a tailor-made software stack for job-specific performance monitoring.

ClusterCockpit currently consists of cc-backend (web API backend), cc-frontend (web-based user interface), cc-metric-collector (a node agent to measure and forward node metric data), and cc-metric-store (a simple metric timeseries in-memory cache). A focus of ClusterCockpit is to define standards and interfaces for an interoperable monitoring ecosystem. We document and specify generic data structure descriptions and came up with a job archive specification that allows to archive and share job performance data in a portable manner.

In 2021, the previous web framework was discontinued in favor of a modular solution using the latest web technology, which is implemented in Golang (backend and system applications) and Svelte (web user interface). Three members of the team (Jan Eitzinger, Thomas Gruber, and Christoph Kluge) and one research assistant contributed in 2021 to the development of ClusterCockpit.

We are currently in the process of acquiring third-party funding and initiating collaborations to further intensify our efforts in this area. Our aim is to establish ClusterCockpit as an attractive open-source offering for job-specific performance monitoring in HPC centers.

4.1.5 MachineState

MachineState (`https://github.com/RRZE-HPC/MachineState`) is a Python3 module and CLI application to document and compare hardware and software settings known to affect performance. It is our contribution to enable deterministic and reproducible benchmark results on today's complex processor and node architectures. MachineState is maintained by Thomas Gruber.

4.1.6 The Bandwidth Benchmark

The Bandwidth Benchmark (`https://github.com/RRZE-HPC/TheBandwidthBenchmark`) is inspired by the famous STREAM main memory bandwidth micro-benchmark by John McCalpin. It can be seen as "STREAM on steroids" and contains eight streaming kernels with varying data access patterns.

Because it is a simple, modular C99 implementation with a simple yet flexible makefile-based build system, it can also be used as a blueprint for other microbenchmarking projects. For teaching purposes, there is a single-file version of this benchmark. The projects Wiki pages (`https://github.com/RRZE-HPC/TheBandwidthBenchmark/wiki`) feature results for many processor architectures. The Bandwidth Benchmark is maintained by Jan Eitzinger.

4.2 Services

4.2.1 HPC user portal

Modern digital workflows are key for efficient project management, handling of user accounts, and resource access. To this end, we are developing the *HPC Portal* web interface, which will provide principal investigators with the means for self-administration of approved projects.

Access to computational NHR resources will be streamlined by utilization of user-supplied public SSH keys, which allows secure, password-less authentication for these accounts. In its current state, the *HPC Portal* provides the aforementioned functionality and an administrative overview for authorized users.

In addition, login via Single-Sign-On (SSO) was recently implemented, allowing users to use the credentials of their respective home organization. We are currently testing this new service with a selected number of trial users in order to find bugs, but also to collect feedback and feature requests. Broader testing and, eventually, productivity of the service are expected to be achieved the course of 2022. The HPC User Portal is maintained by Christoph Kluge.

4.2.2 NHR Moodle learning platform

Since 2021 we host and administrate a Moodle server for general use within the NHR Alliance (`https://moodle.nhr.fau.de/`). Login is enabled through DFN AAI single sign-on authentication. The service is deployed as a docker container on a high-availability VM cluster at RRZE and uses the central RRZE SQL database and tape backup services. The NHR Moodle service is maintained by Christoph Kluge.

4.2.3 NHR@FAU ClusterCockpit service

For usage within the FAU, we administer a central ClusterCockpit monitoring server that can be reached at `https://monitoring.nhr.fau.de`. For this purpose, we procured a dedicated server with large main memory capacity and fast NVMe file storage.

The server is used for all cluster systems and provides job-specific performance monitoring for support personnel and HPC users. Currently only FAU HPC users can authenticate via LDAP. In the future we plan to grant access also to external NHR users. This service is jointly maintained by the *Systems & Services* (Katrin Nusser) and *Software & Tools* (Jan Eitzinger, Thomas Gruber, Christoph Kluge) divisions.

4.3 NHR Alliance central project

Standards and interfaces for a system-wide job-specific performance monitoring (PI: Jan Eitzinger)
There already exist a number of solutions for the system-wide, continuous performance monitoring at the NHR centers. These solutions have been developed partly in cooperation and partly as isolated efforts. This project aims to create interoperability and modularity among these solutions. The goal is, on the one hand, a common solution in which some components remain adaptable or interchangeable and, on the other hand, common standards on which data is collected in which form and granularity.

4.4 Publications and talks

- Jan Eitzinger, Talk at ASHPC 2021 (First Austrian–Slovenian HPC Meeting): *ClusterCockpit—A web application for job-specific performance monitoring*, online, May 31, 2021.
- Thomas Gruber, Christie Louis Alappat, Jan Laukemann and Georg Hager, Webinar about LIKWID, OSACA and SpMV on A64FX processor, Institute for Advanced Computational Science at Stony Brook University, online, July 27, 2021, Recording: `https://youtu.be/0LIvlTULdz0`, Slides: `http://tiny.cc/OOKAMI-Hackathon`
- Thomas Gruber, LIKWID introduction, NRW.HPC Tools Workshops, August 31, 2021, Webpage: `https://events.uni-paderborn.de/event/130/`
- Jan Eitzinger, Talk at NHR Monitoring project Meeting: *ClusterCockpit v1.0*, online, September 30, 2021.
- Jan Eitzinger, Talk at ZKI AK Super Computing: *Was kann ein Job Monitoring System leisten? Diskussion und Beispiele anhand der ClusterCockpit Monitoring-Anwendung*, online, October 07, 2021.
- Thomas Gruber, PRACE Course HPC Code Optimisation Workshop, Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften, online, 2–4 November, 2021, `https://doku.lrz.de/display/PUBLIC/PRACE+Course%3A+HPC+Code+Optimisation+Workshop+2021`.

- Jan Eitzinger, Talk at NHR PE Jour fixe: *Zentrales NHR Projekt: Job-spezifisches Performance Monitoring*, online, December 14, 2021.

- D. Ernst, G. Hager, M. Knorr, G. Wellein, and M. Holzer. *Opening the Black Box: Performance Estimation during Code Generation for GPUs*. Accepted for SBAC-PAD 2021, in 2021 IEEE 33rd International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), Belo Horizonte, Brazil, Oct 26–29, 2021, pp. 22–32. DOI: https://ieeexplore.ieee.org/document/9651599, Preprint: https://arxiv.org/abs/2107.01143.

- D. Pasadakis, C. L. Alappat, O. Schenk, and G. Wellein. *Multiway p-spectral graph cuts on Grassmann manifolds*. Machine Learning (2021). DOI: https://link.springer.com/article/10.1007/s10994-021-06108-1. Preprint: https://arxiv.org/abs/2008.13210.

- C. L. Alappat, N. Meyer, J. Laukemann, T. Gruber, G. Hager, G. Wellein, and T. Wettig. *ECM modeling and performance tuning of SpMV and Lattice QCD on A64FX*. Concurrency and Computation: Practice and Experience, e6512 (2021). Available with Open Access. DOI: https://onlinelibrary.wiley.com/doi/10.1002/cpe.6512, Preprint: https://arxiv.org/abs/2103.03013.

- A. Afzal, G. Hager, and G. Wellein. *Analytic Modeling of Idle Waves in Parallel Programs: Communication, Cluster Topology, and Noise Impact*. Proc. ISC High Performance 2021 Digital, June 24 to July 2, 2021, Frankfurt, Germany. DOI: https://link.springer.com/chapter/10.1007/978-3-030-78713-4_19 Preprint: https://arxiv.org/abs/2103.03175.

- C. L. Alappat, J. Seiferth, G. Hager, M. Korch, Thomas Rauber, and G. Wellein. *YaskSite—Stencil Optimization Techniques Applied to Explicit ODE Methods on Modern Architectures*. 2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), Seoul, Korea (South), 2021 pp. 174–186. DOPI: https://ieeexplore.ieee.org/document/9370316, Preprint: https://blogs.fau.de/hager/files/2021/02/cgo21main-p18-p-aeebf45-49058-preprint.pdf.

- R. Ravedutti L. M., J. Eitzinger, A. M. Maidl, D. Weingaertner. *An instrumentation framework for performance analysis of halide schedules*. Journal of Computer Languages, 101065 (2021). DOI: https://www.sciencedirect.com/science/article/abs/pii/S2590118421000447?via%3Dihub.

- R. Ravedutti L. M., J. Schmitt, S. Eibl, J. Eitzinger, R. Leißa, S. Hack, A. Pérard-Gayot, R. Membarth, H. Köstler. *tinyMD: Mapping molecular dynamics simulations to heterogeneous hardware using partial evaluation*. Journal of Computational Science 54, 101425 (2021). DOI: https://www.sciencedirect.com/science/article/abs/pii/S1877750321001095?via%3Dihub.

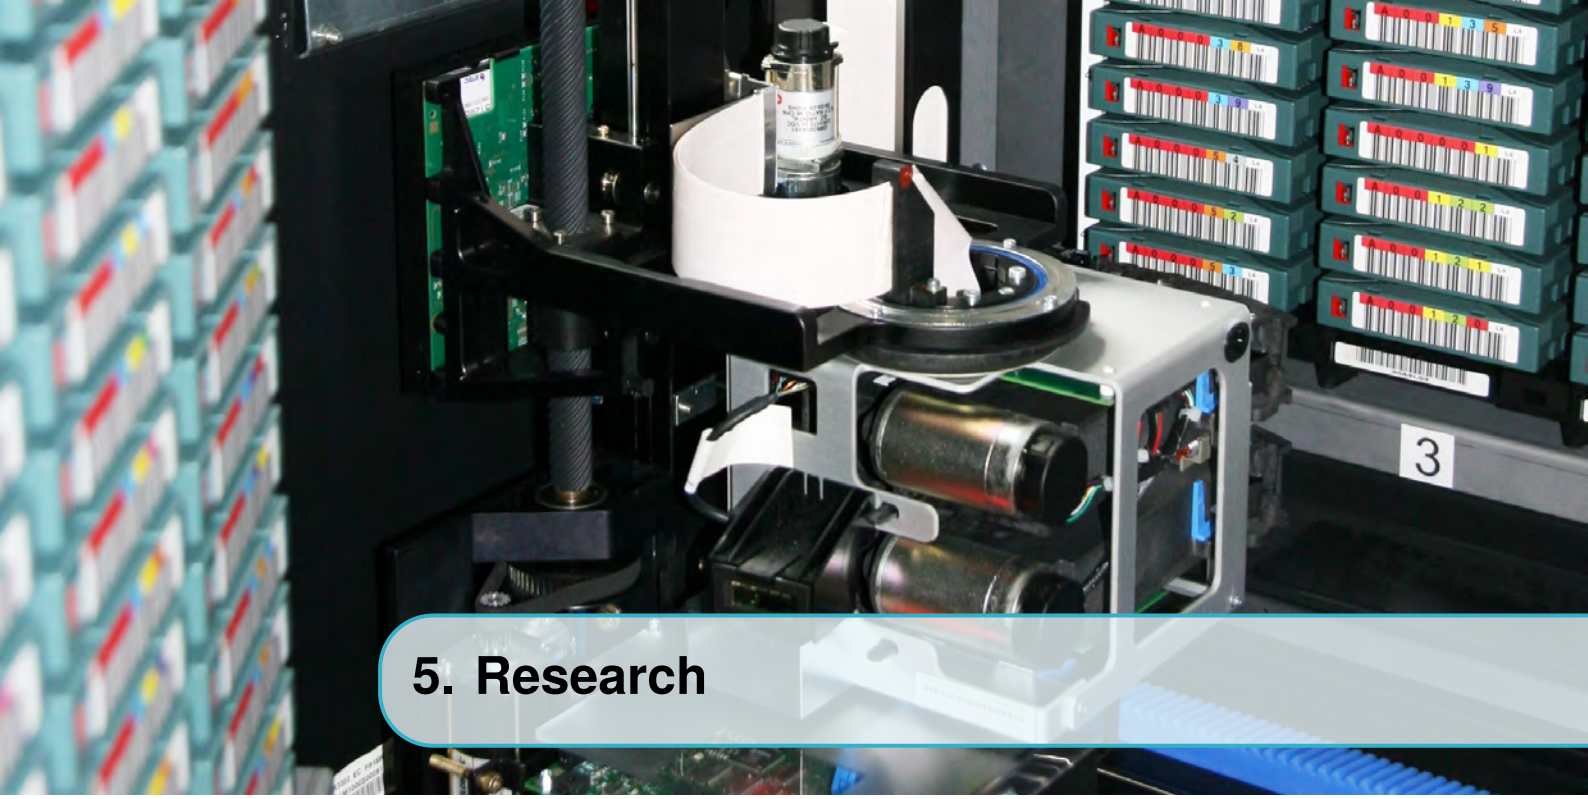

5. Research

Division head: *Prof. Dr. Harald Köstler*

The overarching goal of the research activities at NHR@FAU is to introduce a systematic and model-guided performance engineering (PE) process into all performance analysis and optimization activities within the field of computing. This process helps developers and performance analysts on all levels of expertise to understand observed performance and find optimization opportunities.

Our research is focused on developing and applying analytic performance models, tools, and libraries that support this mission. We participate in projects where we can leverage our expertise to help our partners develop deeper insights into performance issues and how to mitigate them.

Additionally, we combine our HPC expertise with software engineering and code generation technology to provide software frameworks for a variety of applications in the field of computational science and engineering.

5.1 Performance modeling and performance engineering

People

Dominik Ernst, Ayesha Afzal

GPU performance modeling

Our paper "Opening the Black Box: Performance Prediction During Code Generation on GPUs" was accepted for the SBAC-PAD conference in Belo Horizonte, Brazil in October 2021. The paper describes a new method of choosing optimal code generation and run configuration settings by predicting the performance based solely on high level information, which code generators can easily supply. Of special importance is the estimation of data transfer volumes at all stages of a GPU's memory hierarchy, which had not been modeled before with this level of accuracy without requiring execution traces or simulation.

In the context of the European EoCoE-II project (see Sect. 3.3.1), we assisted our partners from the Barcelona Supercomputing Center in analyzing and optimizing the performance of the Navier-Stokes matrix assembly routine in the finite element code

Alya. Through detailed measurements, we assessed the current state of the routine and identified a number of key performance problems and approaches in how to solve them. Our partners were able to implement the solutions according to our analysis, which resulted in speedups ranging from a factor of 3× on the CPU to over 100× on the GPU. Our project partners were aware that the performance of the OpenACC offloaded GPU version had not been in a good state, but had no confidence that a better solution was attainable. A paper about the work is in preparation for publication in 2022.

Idle waves modeling

In 2019 we had published our work on incorporating noise effects into analytic performance models at the IEEE International Conference on Cluster Computing (CLUSTER 2019) in a paper titled *Propagation and Decay of Injected One-Off Delays on Clusters: A Case Study* [1]. We have developed and validated an analytic model of the propagation speed of idle waves, which emerge from delays in execution or communication on specific processes and propagate through the parallel program, much like a train delay that causes other trains to wait and thus *ripples* through the schedule.

In 2021, we extended this work using a variety of HPC platforms and diverse application scenarios to further explore how these idle waves interact nonlinearly within a parallel code on a cluster and how they decay due to system noise, system topology, and application load imbalance. The work was published in the ISC High Performance conference paper *Analytic Modeling of Idle Waves in Parallel Programs: Communication, Cluster Topology, and Noise Impact* [3].

Parallel overlapping memory-bound kernels

In 2020, we analyzed the role of memory-bandwidth bottlenecks in the propagation of idle waves and desynchronization in the ISC High Performance conference paper *Desynchronization and Wave Pattern Formation in MPI-Parallel and Hybrid Memory-Bound Programs* [2]. Desynchronization can cause code to run faster via overlapping communication overhead with useful work.

In 2021, we extended this work towards the study of memory bandwidth sharing within a ccNUMA contention domain by different compute kernels. The results were published in the paper *Analytic performance model for parallel overlapping memory-bound kernels* [5] in the journal *Concurrency and Computation: Practice and Experience*.

Furthermore, we focused on performance aspects of desynchronization and studied how our previous findings, which were mainly obtained using microbenchmarks and toy codes, can be useful in selecting parameter changes for optimizing program performance. This work, titled *The Role of Idle Waves, Desynchronization, and Bottleneck Evasion in the Performance of Parallel Programs* [6], has been submitted to the journal *Transactions on Parallel and Distributed Systems*.

Threshold for load balancing techniques

In November 2021, in collaboration with the High Performance Computing (HPC) research group at the Department of Mathematics and Computer Science at the University of Basel, we started investigating the interplay between advanced load balancing techniques and desynchronization. The main question is whether load balancing can be counterproductive in situations where slight imbalances facilitate desynchronization of MPI processes and, consequently, improve resource utilization. Results are expected to be published in 2022.

Posters and prizes

In 2021, we presented multiple posters: *White-box Modelling of Parallel Computing Dynamics* [7] at the International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia) and *Physical Oscillator Model for Parallel Distributed Computing* [4] in the ISC High Performance 2021 research poster track. We also presented the poster *Noise-driven Cluster-level Performance Modeling and Engineering* at ISC High Performance 2021 PhD Forum. The latter was awarded the *PhD Forum Award – 1st place* by the ISC PhD Forum committee to recognize the most outstanding PhD work.

The NHR@FAU also contributed to the "International HPC Summer School" (IHPCSS) with a tutorial by Dominik Ernst about *GPU Performance Analysis* (listed in Table 3.1). The IHPCSS is jointly organized by HPC institutions from the USA, Canada, Europe and Japan.

5.2 Performance tools

Our developments and research in the field of performance related tools is bundled in the *Software & Tools* division (for details see Section 4.1), which collaborates closely with the *Research* division.

5.3 Building blocks for sparse linear algebra and stencil solvers

People

Christie Alappat, Jan Laukemann

Adaptive linearized storage of sparse tensors

In collaboration with Intel and the University of Oregon, Jan Laukemann developed a novel adaptive mode-agnostic tensor format. The analysis of high-dimensional sparse data is becoming increasingly popular in many important domains. However, real-world sparse tensors are challenging to process due to their irregular shapes and data distributions. The *Adaptive Linearized Tensor Order* (ALTO) format keeps neighboring nonzero elements in the multi-dimensional space close to each other in memory and uses an adaptive bit-encoding scheme that trades off index computations for lower memory usage and more effective use of memory bandwidth.

Moreover, by decoupling its sparse representation from the irregular spatial distribution of nonzero elements, ALTO eliminates the workload imbalance and greatly reduces the synchronization overhead of tensor computations. As a result, the parallel performance of ALTO-based tensor operations becomes a function of their inherent data reuse and shows a speedup of 8x over state-of-the-art approaches [12].

Multiway *p*-spectral clustering

In collaboration with the Advanced Computing Laboratory (Prof. Olaf Schenk) at the Università della Svizzera italiana (USI), we developed an algorithm that improved the accuracy of spectral clustering. In this research we used the information from multiple eigenvectors of the graph Laplacian in the p-norm to improve the quality of the clustering. The joint work resulted in a paper titled *Multiway p-spectral graph cuts on Grassmann manifolds* [15], which was published in the Springer joournal *Machine Learning*.

A64FX processor

In November 2020, our paper *Performance Modeling of Streaming Kernels and Sparse Matrix-Vector Multiplication on A64FX* [9] won the PMBS20 Best Short Paper Award. The paper investigated the then newly developed Arm A64FX processor and developed an ECM performance model for the chip.

In 2021, we were invited to make a more detailed study of our work for a contribution to a special issue the Wiley journal *Concurrency and Computation: Practice and Experience* (CCPE). We carried out extensive benchmarking and revised our ECM model for the newly deployed Fugaku supercomputer (No. 1 in the Top500 list since June 2020), which was powered by the A64FX processor. We further used the developed performance model to optimize two important computational routines: sparse matrix-vector multiplications (SpMV) and the domain-wall kernel from quantum chromodynamics (QCD). The work was done in collaboration with the Department of Physics at the University of Regensburg (Prof. Wettig) and resulted in a journal paper titled *ECM modeling and performance tuning of SpMV and Lattice QCD on A64FX* [10].

Recursive Algebraic Coloring Engine (RACE)

In 2019, we had published our work on the RACE coloring library with the *ACM Transactions on Parallel Computing* in a paper titled *A Recursive Algebraic Coloring Technique for Hardware-efficient Symmetric Sparse Matrix-vector Multiplication* [8]. The library used a level-based approach to generate parallelism for efficiently running sparse linear algebra kernels with dependencies.

In 2021, we extended our RACE framework to optimize performance of sparse matrix-power-vector multiplication (MPK) kernels. The research shows that the same level-based technique used for coloring can be used in the case of MPK kernels to block the matrices for better cache reuse. The idea allows for a generalization of temporal blocking on stencil computations to general sparse matrices. The initial results are very promising and we attain a speedup of 2–3× on modern CPUs compared to the baseline MPK kernel, which uses a sequence of SpMV computations.

5.4 Software engineering for HPC and data analytics

The research presented in this area is a joint work of NHR@FAU and the Lehrstuhl für Systemsimulation (LSS), where mainly the group of Prof. Köstler is involved.

People

Rafael Ravedutti, Sebastian Kuckuk, Sam Jacob

Generation of Higher-Order Discretizations Deployed as ExaSlang Specifications (GHODDESS)

In 2020, we had published our first work on GHODDESS in the *Advances in Water Resources* journal paper *Quadrature-free discontinuous Galerkin method with code generation features for shallow water equations on automatically generated block-structured meshes* [11]. The extension to ExaStencils implements an efficient discontinuous Galerkin discretization of the shallow water equations.

In 2021, we extended our GHODDESS framework with respect to different aspects: First, we implemented additional functionality to do computations on masked block-structured grids which lead to better performance compared to unmasked ones with

the same grid quality and in case of complex domains with small islands or narrow channels make the generation of block-structured grids possible.

Furthermore, we incorporated p-adaptivity and thereby designed a new indicator which will be published soon in the International Journal on Geomathematics. Also, we started working on a new approach of distributing the p-adaptive algorithm to optimize performance. The initial results are very promising and will also be published in the future.

Performance engineering and code generation for molecular dynamics and particle simulations

Within the the NHR-funded project *Optimierung von Bibliotheken für datenparallele Prozessorarchitekturen*, Rafael Ravedutti Lucio Machado provided support in analyzing *libxc*, a library for the computation of exchange-correlation and kinetic energy functionals for density-functional theory. The analysis focused on the overall workflow, the compute kernel composition, the Maple-to-C scripts, and potential for optimization. Based on this, he provided suggestions and a proof-of-concept implementation by extracting the parts of interest from libxc and providing an easy way to replace and evaluate the correctness of these kernels with the new implementation using Python and SymPy. A part of the results was included in a project poster submitted to ISC'22.

Moreover, we worked on the development of MD-Bench (`https://github.com/RRZE-HPC/MD-Bench`), a proxy application to evaluate the performance of molecular dynamics (MD) kernels with several strategies and configurations on different machines. Recently, the implementation of the GROMACS M×N pair list scheme was introduced into MD-Bench, making it the first proxy app with such an optimization. This is particularly important for studying the performance improvements obtained with it and evaluate whether this optimization can be leveraged in other particle simulation scenarios such as Discrete Element Methods (DEM), Smoothed-Particle Hydrodynamics (SPH), and other MD simulations [14] beyond GROMACS (such as the material modeling simulation cases that currently can only be carried out in LAMMPS, for instance).

Finally, he also develops P4IRS (`https://github.com/rafaelravedutti/pairs`), a framework for the code generation of particle simulation kernels targeting multiple hardware such as CPUs and GPUs. The framework uses a simple symbolic description in Python to setup the simulation and to describe the potentials and/or force fields, and then generates optimal code for the chosen hardware using domain knowledge from such simulations achieved in MD-Bench experimentation and performance studies.

Surrogate models for computational fluid dynamics

The aerodynamic optimization process of cars requires multiple iterations between aerodynamics engineers and designers; the improvements are slow because of the required CFD computation time. Surrogate models can produce fast approximate results in a constrained design space; a primary drawback of the 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 the drag coefficient and velocity field of an arbitrary input geometry without explicit parameterization.

We used 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 eleven percent in prediction accuracy for the drag coefficient and bring us a step closer to having a model that can be used for approximate aerodynamic evaluation of unseen, arbitrary vehicle shapes. The

paper, authored by Sam Jacob *et al.*, was accepted for publication in the *SAE International Journal of Passenger Vehicle Systems*; a preprint is available at arXiv [13].

5.5 Projects

The Research division of NHR@FAU led and participated in several research projects in 2021. Funding sources were the NHR Alliance, the DFG, and the EU *Horizon 2020* program.

5.5.1 NHR Alliance central projects

- *Highly scalable numerical solution algorithms and numerical libraries (PI: Ulrich Rüde)*
  The development of massively parallel, matrix-free multigrid solvers is a core ingredient of the research that is conducted in the context of previous and ongoing projects at FAU. One bottleneck of the parallel implementation is the solution of the coarse grid problem. This can hardly be parallelized efficiently since the number of unknowns may be roughly of the order of the employed parallel processes. The communication overhead of a coarse grid solver outweighs the computational intensity by far. Especially the performance of more complex coarse grid solvers such as algebraic multigrid or even sparse direct methods suffers. However, such solvers are necessary in the presence of numerical challenges such as viscosity jumps or distorted element shapes.

- *Cx as a service for sustainable HPC research software engineering (PI: Harald Köstler)*
  A large number of active developers contributing to the same code base increases the probability of new bugs being introduced because not all possible configurations are being tested. This either decreases the quality of the code, or requires significant amounts of human resources to be spent on integration tests before a new release is made. Continuous Integration (CI) can reduce both the amount of time necessary for these tests, as well as provide near instant feedback after bugs that cause build failures have been introduced, by automatically building all possible configurations at given points in time (e.g., after every source code commit). In addition, Continuous Benchmarking (CB) frames not only unified testing but also a continuous benchmark environment. Finally, Continuous Deployment (CD) enables scientists to provide efficient, reliable and sustainable HPC software.

- *Optimization of software libraries for DFT codes for data-parallel processor architectures (PI: Bernd Meyer)*
  There is a significant potential for modernization in some widely used application programs especially in the area of atomistic simulations, both with respect to the more efficient use of existing data parallelism and the classical problem of code scaling. The area of electron structure calculation using density functional methods (DFT) and methods for calculating the dynamics based on these methods has been identified as an area in which the researchers can join their expertise and promptly implement it in efforts to increase the efficiency of application codes. Especially in this area, the development and implementation of complex program packages increasingly relies on modularization, specialization in code development and the use of widely applicable software libraries. In the present project, we address the libxc, DBCSR and FFTXlib libraries and support, among others, the strategy for modularization of electron structure codes targeted by the Materials at the Exa-Scale, European Center of Excellence (MaX).

5.5.2 DFG project

Dynamic HPC software packages: Seamless integration of existing software packages and code generation techniques

Complex phenomena in the natural and the engineering sciences are increasingly being studied with the help of simulation techniques. This is facilitated by a dramatic increase of the available computational power, and Computational Science and Engineering (CSE) is emerging as a third fundamental pillar of science. CSE aims at designing, analyzing, and implementing new simulation methods on high-performance computing (HPC) systems such that they can be employed in a robust, user-friendly, and reliable fashion to a wide variety of scientific and engineering problems. Considering the high cost of supercomputing, reaching the best possible computational efficiency becomes a primary criterion and is central to the research agenda of CSE. Furthermore, HPC software must support a range of increasingly complex applications on modern heterogeneous and volatile hardware platforms, where often many different algorithms are combined to model interacting physical processes. For that HPC software has often to be modified extensively to make full use of the additional performance of newly released architectures. Our main goal is to provide a new class of dynamic software frameworks to the HPC users that combine existing and established HPC frameworks with current code generation technology to increase the productivity when introducing new applications or porting to new platforms. We will show the benefits of this approach for three real-world multi-physics applications, optimization of wind turbines and wind farms, formation and dynamics of dunes that occur in many environmental systems such as riverbeds, and simulation of charged particles in microfluidic flows.

5.5.3 EU project

EoCoE-II—Energy oriented Centre of Excellence: towards exascale for energy (`https://www.eocoe.eu`). For detailed information please see Section 3.3.1.

Publications

[1] A. Afzal, G. Hager, and G. Wellein. "*Propagation and Decay of Injected One-Off Delays on Clusters: A Case Study*". In: *Proceedings - IEEE International Conference on Cluster Computing, ICCC* (Albuquerque, NM). Volume 2019-September. Institute of Electrical and Electronics Engineers Inc., Sept. 23–26, 2019. ISBN: 9781728147345. DOI: 10.1109/CLUSTER.2019.8890995.

[2] A. Afzal, G. Hager, and G. Wellein. "*Desynchronization and Wave Pattern Formation in MPI-Parallel and Hybrid Memory-Bound Programs*". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Frankfurt). Edited by Ponnuswamy Sadayappan et al. Volume 12151 LNCS. Cham: Springer International Publishing, June 22–25, 2020, pages 391–411. ISBN: 9783030507428. DOI: 10.1007/978-3-030-50743-5_20.

[3] A. Afzal, G. Hager, and G. Wellein. "*Analytic Modeling of Idle Waves in Parallel Programs: Communication, Cluster Topology, and Noise Impact*". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Virtual, Online). Edited by Bradford L. Chamberlain, Ana-Lucia Varbanescu, and Hatem Ltaief andPiotr Luszczek. Volume 12728 LNCS. Springer Science and Business Media Deutschland GmbH, June 24–July 2, 2021, pages 351–371. ISBN: 9783030787127. DOI: 10.1007/978-3-030-78713-4_19.

[4] A. Afzal, G. Hager, and G. Wellein. *Physical Oscillator Model for Parallel Distributed Computing*. Poster at ISC High Performance 2021. June 24–July 2, 2021. URL: https://www.researchgate.net/publication/354208484_Physical_Oscillator_Model_for_Parallel_Distributed_Computing?_sg=lVYEm3XW4E92lsf55nIWTxkXYhgpB13cA2Zi3zbsaP-YcPn2zMsmYHnIURLA_eADEZioHef0aYKgIpaCYNIPew0H4GtDTB14Q-E_avYu.W3TTE5Nzo7hjNZ45wagExeFeJB8qTBJQe764hOdV9pK0lbGaFk1muIggkEpWAveETapO_wOS8AKyPVJ1-4g0WA.

[5] A. Afzal, G. Hager, and G. Wellein. "*Analytic performance model for parallel overlapping memory-bound kernels*". In: *Concurrency and Computation-Practice & Experience* (Jan. 2022). DOI: 10.1002/cpe.6816. URL: https://onlinelibrary.wiley.com/doi/10.1002/cpe.6816.

[6] A. Afzal, G. Hager, and G. Wellein. "*The Role of Idle Waves, Desynchronization, and Bottleneck Evasion in the Performance of Parallel Programs*". In: *arXiv*. Submitted. 2022. arXiv: 2205.04190 [cs.DC].

[7] A. Afzal, G. Hager, and G. Wellein. *White-box Modelling of Parallel Computing Dynamics*. Poster at HPC Asia 2022. Jan. 12–14, 2022. URL: https://sighpc.ipsj.or.jp/HPCAsia2022/poster/115_abstract.pdf.

[8] C. Alappat et al. "*A recursive algebraic coloring technique for hardware-efficient symmetric sparse matrix-vector multiplication*". In: *ACM Transactions on Parallel Computing (TOPC)* 7.3 (2020), pages 1–37.

[9] C. Alappat et al. "*Performance Modeling of Streaming Kernels and Sparse Matrix-Vector Multiplication on A64FX*". In: *2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. IEEE. 2020, pages 1–7.

[10] C. Alappat et al. "*Execution-Cache-Memory modeling and performance tuning of sparse matrix-vector multiplication and Lattice quantum chromodynamics on A64FX*". In: *Concurrency and Computation: Practice and Experience* (2021), e6512.

[11] Sara Faghih-Naini et al. "Quadrature-free discontinuous Galerkin method with code generation features for shallow water equations on automatically generated block-structured meshes". In: *Advances in Water Resources* 138 (2020), page 103552.

[12] A. E. Helal et al. "*ALTO: adaptive linearized storage of sparse tensors*". In: *Proceedings of the ACM International Conference on Supercomputing*. 2021, pages 404–416.

[13] S. J. Jacob et al. "*Deep Learning for Real-Time Aerodynamic Evaluations of Arbitrary Vehicle Shapes*". In: *arXiv preprint arXiv:2108.05798* (2021).

[14] R. R. L. Machado et al. "*tinyMD: Mapping molecular dynamics simulations to heterogeneous hardware using partial evaluation*". In: *Journal of Computational Science* 54 (2021), page 101425.

[15] D. Pasadakis et al. "*Multiway p-spectral graph cuts on Grassmann manifolds*". In: *Machine Learning* (2021), pages 1–39.

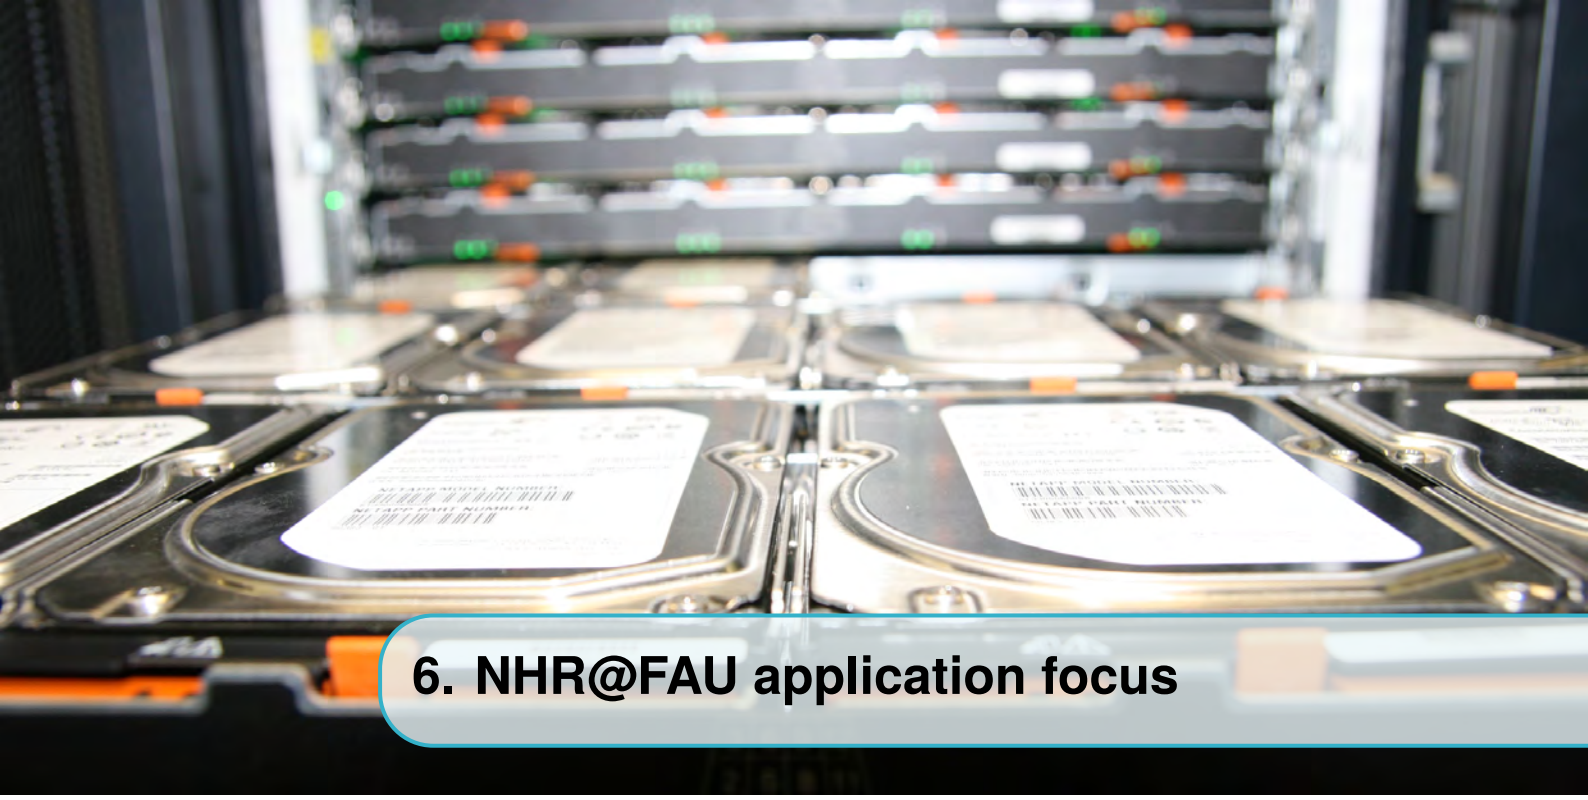

6. NHR@FAU application focus

NHR@FAU offers the whole spectrum of both atomistic simulation methods and their fields of application: Scientists at FAU use methods from classical MD to quantum chemistry, and everything in between. Atomistic simulation methods are applied in chemistry, biology, physics, medicine, materials science, and engineering.

We have established a Germany-wide unique interdisciplinary competence center that helps users to select and use atomistic simulation methods in an HPC environment and actively accompanies and coordinates the development of high-performance simulation codes. An interdisciplinary approach promises not only synergy effects through the exchange and joint development of simulation and evaluation tools but a cross fertilization of materials and life sciences, which often use the same or very similar simulation techniques.

A further core project is the education and lifelong training of scientists and engineers. The close cooperation between theory, simulation, and experiment ensures that the training is not aimed specifically at modelers but that it is made available to experimental colleagues as well. NHR@FAU makes an important contribution to the key technologies of scientific computing and software development through sustained concentration of methodological competence both in application and development of computer codes and their hardware-related optimization.

6.1 Overview

Atomistic Simuation Center (ASC)

The NHR centers in Paderborn, Berlin, and Erlangen have a strong focus on atomistic simulations. Researchers develop methods for and perform interdisciplinary simulations with various applications in the materials and the life sciences, spanning the areas of physics, biology, and chemistry. In order to join forces, share expertise, and bundle competence, the centers have formed the *Atomistic Simulation Center*. This umbrella organization serves as a platform providing guidance in tackling research questions by simulations. The *liaison scientists* at the three centers, who are local experts in software and application domains, will assist in providing individual support through advice

and training on methods, software, tricks and pitfalls, best-suited architectures, and best practices.

The Atomistic Simulation Center was made public to the community at a (fully online) inauguration symposium on October 26–27, 2021. In 16 talks, PIs from the three centers and invited speakers presented their research in the different areas of atomistic simulations. With up to 120 participants attending simultaneously, the symposium was truly a success.

CECAM node

To further foster synergies, the sites of the NHR centers constituting the ASC, ZIB (Berlin), University of Paderborn, and FAU (Erlangen) have signed an agreement to form a joint CECAM node *Mathematics and Computation in Molecular Simulation* as a consortium (until the end of 2025 and renewable). The node is represented on the board of CECAM directors by node director Felix Höfling (ZIB) or one of the deputy directors Petra Imhof (FAU) or Thomas Kühne (UPB). In the agreement, the goals of the node were states as follows:

> The general purpose of CECAM [. . .] is to enhance European cooperation in computational research in biology, chemistry, engineering, and physics, and broaden its interest beyond these research areas.

This will be realized by, e.g., organization of workshops and schools and a visitor and exchange program, which is in line with the objectives of the ASC and the NHR centers: training of researchers and knowledge exchange.

6.2 Activities of liaison scientists

6.2.1 Dr. Frank Beierlein

Computer-Chemie-Centrum (CCC) and NHR@FAU

NHR activities, projects, and support

Frank Beierlein was part of the organizing team of the NHR Atomistic Simulation Center (ASC) Inauguration Symposium 2021, which took place online from October 26–27, 2021. Here, he supported Prof. Petra Imhof with the preparations and was in close contact with the other ASC centers (Paderborn and Berlin). During the conference, Frank was in charge of the technical back office and the video recordings (together with Tobias Klöffel). These videos were later processed and published by Harald Lanig.

At the NHR board meeting on September 20, 2021, Frank had the opportunity to present his projects to the NHR community in Erlangen. The regular NHR liaison scientist seminars and the AG Imhof group seminar serve as further opportunity to present and discuss current and planned activities. Close contact with NHR@FAU is further ensured by regularly attending the NHR@FAU *HPC Café*.

As one of the liaison scientists (LS) in charge of the MD simulation package Amber, Frank benchmarked and tested Amber20 on the new RTX 3080 (*TinyGPU*) and Tesla A40 (*Alex*) GPUs, in addition to local GPUs (RTX A5000). As the new clusters use Slurm instead of torque as batch system, the submit scripts for different types of Amber jobs were adapted for Slurm. Watching the Amber list and the Amber literature ensures being up-to-date on what happens in the Amber community. As one of the two liaison scientists responsible for Amber Frank provides support for Amber20 customers of NHR@FAU.

Together with the *Systems & Services* group of NHR@FAU (Thomas Zeiser), Frank optimized the setup of Amber 20 free energy thermodynamic integration simulations for efficient use of the new GPU cluster *Alex* (and the new nodes on *TinyGPU*).

Frank is an expert in parameterizing nonstandard residues for Amber. In his current research on DNA repair mechanisms Frank develops high-quality parameters for damaged DNA bases which are part of a DNA strand that binds to a DNA repair enzyme. These systems are extensively simulated on the NHR@FAU's GPU clusters. In addition to conventional, unbiased MD, Frank is an expert in the field of alchemistic free energy methods, like thermodynamic integration, a technique that is increasingly used in pharmaceutical research. He adapted and refined the setup, simulation and analysis process of this kind of simulations, so that these simulations can now be run routinely on the NHR@FAU HPC clusters. Also, the process of analyzing DNA/protein MD simulations and free energy simulations was optimized. The results of Frank's research were published in 2 peer-reviewed papers, and a HPC user report was written for the NHR@FAU/HPC web pages. Petra Imhof and Frank Beierlein also closely collaborate with Andriy Mokhir (Organic Chemistry, Erlangen, proposal for a new DFG-funded Research Training Group).

In their teaching (see below) and research activities, Frank Beierlein and Petra Imhof take care that documentation, scripts and inputs are as user-friendly as possible to ensure that even inexperienced users or potential NHR customers can use advanced simulation techniques on HPC systems.

Teaching and training

Frank is closely involved in teaching computational chemistry courses/hands on user training for the FAU students at the Computer-Chemistry-Center (CCC). These were in the winter semester (WS 20/21 and 21/22): *Moderne Softwareapplikationen* (2 courses), *Molecular Modelling* and in the summer semester (SS 21): *Computational Chemistry* (2 courses), *Molecular Modeling*, Bio-Organic & Bio-Inorganic LAB and Drug Discovery LAB.

Publications

- F. Beierlein, S. Volkenandt, P. Imhof. *Oxidation Enhances Binding of Extrahelical 5-Methyl-Cytosines by Thymine DNA Glycosylase*, J. Phys. Chem. B 2022, 126, 1188–1201.

- S. Volkenandt, F. Beierlein, P. Imhof. *Interaction of Thymine DNA Glycosylase with Oxidised 5-Methyl-cytosines in Their Amino- and Imino-Forms*, Molecules 2021, 26, 5728.

Outreach

The NHR Atomistic Simulation Center (ASC) Inauguration Symposium 2021, online October 26–27, 2021 was not only a highly interesting and very successful scientific conference, but it also offered the opportunity to inform a broader audience about NHR. Unfortunately, our Molecular Modeling Workshop, which is annually organized in Erlangen by the Molecular Graphics and Modeling Society—Deutschsprachige Sektion e.V was cancelled due to COVID-19. Frank is an active member in this society and supports the managing board as cash auditor („Kassenprüfer“).

Collaboration with other liaison scientists and HPC staff

Together with Anselm Horn, Frank is the responsible liaison scientist for Amber. Together with Thomas Zeiser, the Amber performance was optimized on the NHR@FAU

HPC systems. Together with all other LS, Harald Lanig, and Tobias Klöffel from the Training & Support division of NHR@FAU, Frank was part of the organizing team of the NHR Atomistic Simulation Center (ASC) Inauguration Symposium 2021.

6.2.2 PD Dr. Anselm Horn

Department of Medicine (Professorship for Bioinformatics) and NHR@FAU

NHR activities, projects, and support

Anselm Horn provided technical and scientific support about the HPC systems in Erlangen for Carina Höring from the Universität Regensburg during her PhD project "Histamine H2 and H4 receptor ligands: Sub-type, functional and ortholog selectivity." The project was conducted under supervision of PD Dr. Andrea Strasser (Pharmazeutische/Medizinische Chemie II, Universität Regensburg) as part of the DFG-funded Research Training Group 1910 "Medicinal Chemistry of Selective GPCR Ligands." LS support is acknowledged in a resulting research paper [1].

Further, Anselm Horn was involved in a cooperative project with the workgroup of Prof. Dr. Dr. Lorenz Meinel (Institut für Pharmazie und Lebensmittelchemie, Julius-Maximilians-Universität Würzburg). Using his expertise on Amber force field parameterization, he developed ff14SB-compatible parameters for a PEG3 and a PEG6 spacer unit. In conjunction with earlier developed parameters for isopeptide bonds, he was able to investigate the dynamics of a covalently coupled PEG-containing peptide linker around insulin-like growth factor using the A100 GPU cards of *TinyGPU*. The results are to be published in a joint research paper [2] and are used as preparatory work for a CRC proposal, which in turn will comprise several computational projects for NHR@FAU.

Moreover, he participated in a cooperative project with Dr. Lisa Link-Paulus from the workgroup of Prof. Dr. Anja Boßerhoff (Institut für Biochemie, Friedrich-Alexander-Universität Erlangen-Nürnberg). The experimental researchers had found a new splice of the regulatory protein Argonaute2 influencing microRNA function and melanoma cell viability. Molecular dynamics simulations yielded a distinct difference in structure and flexibility between wild type and spliced protein form. Only the availability of A100 GPUs on *TinyGPU* and the expertise of the LS made the computational work feasible in the short timespan available for this competitive project. The results are to be published in a joint research paper [3] and are to be presented as conference poster [4].

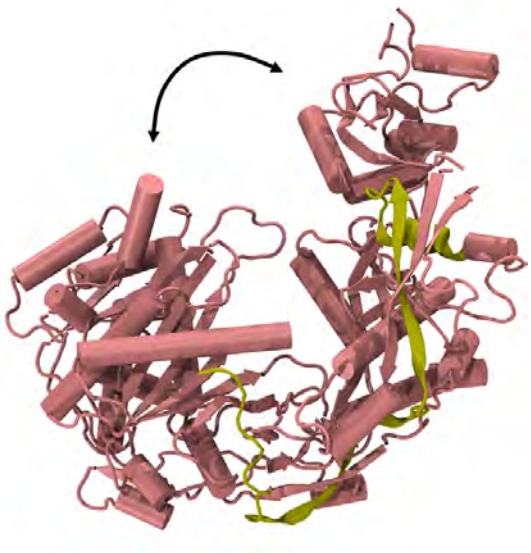

Figure 6.1: Structural model of Argonaute2 (PDB-ID: 4z4d) that served as template for the two different protein species. The main part of the protein is colored pink, the N-terminal residues not present in the variant are colored yellow. The molecular dynamics simulations yielded a significant widening of the RNA-binding cleft in the truncated species compared to the full-lenght protein and an enhanced mutual flexibility of the protein parts.

In addition, he was granted an early-user-access for the new clusters *Alex* (GPU) and *Fritz* (CPU) in order to especially test the large number of A40 cards in simulations. The respective project aims at investigating the feasibility of

modelling the binding of a ligand to a single chain antibody fragment (scFv); for that many different initial ligand orientations in the simulation setup ensure truly independent simulation runs. The project experienced some delay, due to challenges in the ligand parameter generation, but is to be continued and finished in 2022. Since this project was the first larger-scale application of AMBER on *Alex* and *Fritz*, the experience gained will be most valuable for later cluster users and NHR@FAU projects.

Teaching and training

In order to provide a first help for new users, Anselm Horn started to work on an extension of the Tips & Ticks for Amber with a focus on practical aspects with regard to the new compute clusters *Alex* and *Fritz*. Some of his experiences from the early-access-project have been shared within the Bioinformatics group.

Meetings, travel, and community outreach

Anselm Horn was part of the core organizational committee that planned and realized the NHR Atomistic Simulation Center (ASC) Inauguration Symposium 2021, which took place as online conference on October 26–27, 2021 `https://www.atomistic-simlab.hpc.fau.de/asc-nhr-2021/`.

He was active in support on the Amber Mailing List helping Amber users all over the world to overcome problems, especially in the area of parameter generation. There, he ensured that the NHR@FAU affiliation was clearly visible in his posts thus helping to increase the publicity of the abbreviation as a kind of brand and the NHR initiative in general.

Collaboration with other liaison scientists

On suggestion of Dr. Thomas Zeiser and PD Dr. Harald Lanig, he initiated a joint benchmark project between the LS involved in atomistic simulation. The idea for that is to investigate the practical performance of different molecular dynamics codes on systems, which are or have been under active investigation in the different workgroups, in order to provide potential users with some notion about actual run times on the NHR@FAU resources for their project proposals and applications. After an organizational hold at the end of 2021, the project will be pursued in 2022.

References

1 C. Höring, M. Conrad, C. A. Söldner, J. Wang, H. Sticht, A. Strasser, Y. Miao. *Specific Engineered G Protein Coupling to Histamine Receptors Revealed from Cellular Assay Experiments and Accelerated Molecular Dynamics Simulations* Int. J. Mol. Sci. 2021, 22, 10047; DOI: 10.3390/ijms221810047.

2 M. Beudert, L. Hahn, A. H. C. Horn, N. Hauptstein, H. Sticht, L. Meinel, R. Luxenhofer, M. Gutmann, T. Lühmann. *Merging bioresponsive release of insulin-like growth factor I with 3D printable thermogelling hydrogels*, submitted.

3 L. Linck-Paulus, A. H. C. Horn, A. Matthies, S. Fischer, G. Meister, H. Sticht, M. Kappelmann-Fenzl, A. Bosserhoff. *A previously unknown Argonaute 2 variant positively modulates the viability of melanoma cells*, submitted.

4 L. Linck-Paulus, A. H. C. Horn, A. Matthies, S. Fischer, G. Meister, H. Sticht, M. Kappelmann-Fenzl, A. Bosserhoff. *A new Argonaute 2 splice variant influences microRNA function and melanoma cell viability* to be presented at the 73rd Mosbacher Kolloquium *The World of RNAs—Principles & Applications*, Mosbach/Baden, Germany; March 31 to April 2, 2022.

6.2.3 Dr. Sebastian Kuckuk

Chair of Computer Science 10 (System Simulation) and NHR@FAU

NHR activities, projects, and support

Sebastian Kuckuk contributed to the NHR-funded project *Optimierung von Bibliotheken für datenparallele Prozessorarchitekturen*. This included providing support in analyzing libxc, a library for the computation of exchange-correlation and kinetic energy functionals for density-functional theory. The analysis focused on the overall workflow, the compute kernel composition, the maple-to-C scripts, and revealing potential for optimization. Based on this, he provided suggestions and a proof-of-concept implementation for employing code generation techniques using Python and SymPy. A part of the results was included in a project poster submitted to ISC'22.

Moreover, he provided consultation for the group of Vadym Aizinger (University of Bayreuth) on the efficient use of GPUs. Noteworthy topics included porting legacy code to accelerators, profiling and performance analysis for GPU-enabled applications, employing hybrid (CPU-GPU) parallelization for application codes, and porting applications to energy-efficient SoC architectures such as NVIDIA Tegra. The target applications were two different approaches to solving the shallow water equations (SWE) using higher-order discontinuous Galerkin (DG) discretizations for the simulation of ocean currents. One is realized with a code generation and domain-specific language (DSL) approach. The target code is MPI/OpenMP hybrid C++ accelerated with CUDA and employs block-structured grids. The other is based on a "traditional" implementation and unstructured grids. It uses C/C++/FORTRAN and was later accelerated with OpenCL.

Teaching and training

In order to extend the training offered by NHR@FAU, Sebastian Kuckuk prepared for the delivery of GPU programming tutorials in 2022 via completing multiple courses offered by the NVIDIA Deep Learning Institute (DLI) on GPU programming and CUDA-accelerated applications that scale across multiple GPUs, either within one compute node or across multiple nodes. He also attained certification as DLI ambassador and was certified to teach the course *Fundamentals of Accelerated Computing with CUDA C/C++*. The first delivery to members of NHR@FAU is planned for early 2022.

Secondly, he developed a training unit on performance analysis for stencil codes on GPUs. This entailed implementing benchmark cases, profiling different configurations, data editing, and creating slides. A first version was presented as part of the Programming Techniques for Supercomputers (PTfS) lecture. An extension of the popular tutorial *Node Level Performance Engineering (NLPE)* with the newly created materials is planned for the future.

Thirdly, he conceptualized a new workshop on performance portability and programmer productivity and realized first teaching units as interactive Jupyter notebooks.

Meetings, travel, and community outreach

- Invited talk on *Introduction to Code Generation Techniques* at the EXA2PRO-EoCoE joint workshop (February 22–24, 2021, online).
- NHR@FAU liaison scientist meeting (June 16, 2021, online).
- In-person meeting with the group of Vadym Aizinger (Bayreuth, November 19, 2021).
- Weekly meetings of the *Training & Support* group.

Collaboration with other liaison scientists

Joint work with Rafael Ravedutti on *Optimierung von Bibliotheken für datenparallele Prozessorarchitekturen*.

6.2.4 Dr. Samaneh Nasiri

Institute for Materials Simulation (WW8) and NHR@FAU

NHR activities, projects, and support

Samaneh Nasiri joined NHR@FAU as a Liaison scientist (Department of Materials Science and Engineering) on October 01, 2021. She contributed to preparing a DFG-Funded proposal entitled *Molecular simulation of deformation processes in composites of light metals and covalent nanoparticles*. This project aims to find out how embedded nanoparticles change the deformation and fracture properties of the resulting nanocomposites and identify optimal parameters for mechanical property enhancement. This study involves large-scale atomistic simulation, which is planned to be carried out using the computation resources at NHR@FAU. The research will be conducted in collaboration with colleagues in the UK and China who investigate the same metal-carbon systems experimentally and with other simulation methods such as density functional theory. Samaneh Nasiri has conducted preliminary investigations and proof-of-concept simulations required for the project proposal.

Samaneh Nasiri revised and published two pre-reviewed journal papers on the atomistic simulations of metal-carbon nano-particles systems, which have been conducted using the NHR@FAU computation resources.

Meetings, travel, and community outreach

Samaneh Nasiri organized recurring online monthly meetings with the collaborators in the UK and China to discuss the preliminary investigation and progress of the above-mentioned DFG-Project.

Moreover, she contributed to the NHR Atomistic Simulation Center Symposium held on October 27, 2021 by designing and creating the symposium flyer and assisting the chairman in the "materials science at the atomistic scale" session.

Collaboration with other liaison scientists

Samaneh Nasiri collaborated with Rafael Ravedutti in providing LAMMPS input scripts for benchmarking MD-bench for different system sizes, thermodynamics ensemble, and inter-atomic potentials.

6.2.5 Rafael Ravedutti

Chair of Computer Science 10 (System Simulation) and NHR@FAU

NHR activities, projects, and support

Rafael Ravedutti LucioMachado contributed to the NHR-funded project *Optimierung von Bibliotheken für datenparallele Prozessorarchitekturen*. This included providing support in analyzing libxc—a library for the computation of exchange-correlation and kinetic energy functionals for density-functional theory. The analysis focused on the overall workflow, the compute kernel composition, the maple-to-c scripts and revealing potential for optimization. Based on this, he provided suggestions and a proof-of-concept implementation by extracting the parts of interest from libxc and providing an easy way to replace and evaluate the correctness of these kernels with the new implementation using Python and SymPy. A part of the results was included in a project poster submitted to ISC'22.

Moreover, he also worked (and still works) on the development of MD-Bench (`https://github.com/RRZE-HPC/MD-Bench`), a proxy-application to evaluate the performance of Molecular Dynamics kernels with several strategies and configurations on different machines. Recently, the implementation of the GROMACS MxN pair list scheme was introduced into MD-Bench, making it the first proxy app with such optimization. This is particularly important to study the performance improvements obtained with it and evaluate whether this optimization can be leveraged into other particle simulation scenarios such as Discrete Element Methods (DEM), Smoothed-Particle Hydrodynamics (SPH) and other Molecular Dynamics simulations beyond GROMACS (such as the material modeling simulation cases that currently can only be executed in LAMMPS, for instance).

Finally, he also developed (and still develops) P4IRS `https://github.com/rafaelravedutti/pairs`), a framework for the code generation of particle simulation kernels targeting multiple hardware such as CPUs and GPUs. The framework uses a simple symbolic description in Python to setup the simulation and to describe the potentials and/or force fields, and then generates optimal code for the chosen hardware using domain-knowledge from such simulations achieved in MD-Bench experimentation and performance studies.

Teaching and training

Since 2021, Rafael has been responsible for the *Exercises for High-end Simulation and Practice* (HESP) lecture from the Chair of Computer Science 10 (System Simulation). The exercises consist of implementing particle simulations (MD, DEM, and/or SPH) to execute efficiently on GPUs with domain-specific optimizations such as the Linked Cells and Verlet Lists. He also supervises the final projects for the lecture which can focus on implementing DEM or SPH, as well as porting the application to execute in multiple GPUs with MPI.

Rafael also supervised the work of a Master's student (Maximilian Gaul) project in the *Effiziente numerische Simulation auf multicore-Prozessoren* (MuCoSim) lecture, which consisted of porting the Lennard-Jones kernel in MD-Bench to CUDA in order to execute it efficiently on NVIDIA GPUs. The final report presented a performance analysis of the implemented kernel.

Meetings, travel, and community outreach

- NHR@FAU liaison scientist meeting (June 16, 2021, online).
- Weekly meetings of NHR Research Day.
- Talk *Introducing GROMACS optimizations into MD-Bench* (March 28, 2022, online).

Collaboration with other liaison scientists

Joint work with Sebastian Kuckuk on *Optimierung von Bibliotheken für datenparallele Prozessorarchitekturen*, collaboration from Samaneh Nasiri and Anna Kahler providing experimental cases (material modeling and bio-sciences, respectively) into MD-Bench, and Talk *Vorstellung und Abstimmung MD-Bench* to present our proxy-app and gather suggestions and feedback from other Liaison Scientists concerning next steps for our research in Molecular Dynamics performance analysis (November 23, 2021, online).

6.2.6 Marius Trollmann

Department of Biology (Professorship for Computational Biology) and NHR@FAU

NHR activities, projects, and support

Marius Trollmann contributed to the organization and execution of the inauguration symposium of the NHR *Atomistic Simulation Center (ASC)*. The ASC is represented by the three NHR centers in Paderborn, Erlangen and Berlin. The center trains, supports and connects users in the application of different simulation techniques for their research in a high-performance cluster (HPC) environment. He maintained the virtual meeting during the symposium and took part as a co-chair in the session *Modelling Soft Matter at the Atomic Scale*.

He also took an active part in the recent test phase of the new established GPU cluster *A*lex at the NHR@FAU. He provided a wide-range of systems to analyze the performance of the simulation software GROMACS on the available A40 and A100 GPUs. The results of the analysis could help later NHR applicants to get a reliable estimation of their required computation time. It will further help them to exploit their granted resources more efficiently.

Moreover, he included the already available resources in his own research projects to study biomolecular systems in a pharmaceutical context. One of his projects approaches the toxicity of hundreds of antimicrobial peptides (AMPs) against red blood cells in a high-throughput method. The huge number of simulation systems requires the extensive optimized use of multiple GPUs. A publication is currently under preparation.

Another recent project considers the study of mRNA-based vaccines. Approximately 28,500 GPU-hours were spent to characterize the structure and the physico-chemical properties of the lipid nanoparticles (LNPs) used in the BioNTech&Pfizer vaccine employing classical molecular dynamics (MD) simulations. The computationally demanding atomistic simulations of the LNPs with a size between three and seven Mio atoms required the extensive use of multiple A40 and A100 GPUs. Parts of the work were presented at the virtual international *Biophysics of membranes and beyond* SFB-GEM Meeting 2021 and the 13th European Biophysics congress (EBSA) in Vienna. Some of the results will be also presented at the *Biophysics at the Dawn of Exascale Computers* thematic meeting of the Biophysical society in Hamburg. A corresponding publication is currently under review.

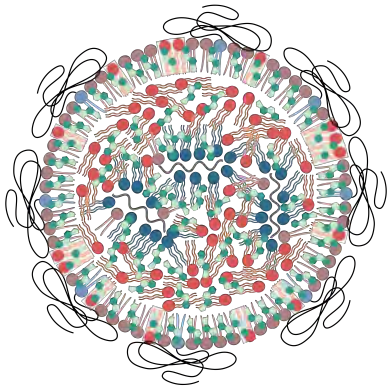

Figure 6.2: Schematic view of a solid-core lipid nanoparticle enveloping multiple mRNA-strands. The coloring of the molecules corresponds to the different types of lipids employed in the composition.

Teaching and training

To offer a practical insight into the theory and application of MD simulation of biological systems, Marius Trollmann is involved in the master module *Orientierungsmodul Strukturbiologie I* and *Orientierungsmodul Strukturbiologie II*. Participants get in touch with a realistic research hypothesis for which they then prepare and perform own simulations. This also includes the usage of NHR resources to achieve reliable results at a reasonable time scale.

Secondly, he contributed to the practical part of the module *Molecular Dynamics Simulation of Biological Membranes* which deals in particular with the programming of a simple MD-Engine employing Python and the application of free-energy calculations with GROMACS.

Meetings, travel, and community outreach

- NHR@FAU liaison scientist meeting (June 16, 2021, online).
- 13th European Biophysics conference (July 24–28, 2021, Vienna).
- SFB-GEM Meeting 2021, *Biophysics of membranes and beyond* (September 27–30, 2021, online).
- Inauguration Symposium of the Atomistic Simulation Center (October 26–27, 2021, online).
- Thematic meeting of the Biophysical society, *Biophysics at the Dawn of Exascale Computers* (May 16–20, 2022, Hamburg).
- Weekly meetings of the *Training & Support* group.

Collaboration with other liaison scientists

Joint report with other liaison scientists (Anselm Horn, Frank Beierlein, Samaneh Nasiri, Anna Kahler) summarizing benchmark results from different software packages (GROMACS, Amber, NAMD and LAMMPS). The report will be published on the open-access repository Zenodo.

6.2.7 Dr. Egor Trushin

Chair of Theoretical Chemistry and NHR@FAU

NHR Activities, projects, and support

Egor Trushin performed benchmark tests of GPU-accelerated Vienna Ab initio Simulation Package (VASP) on *TinyGPU* and *Alex* GPGPU clusters. Performance for various computational setups was analysed including various GPU models and number of used GPUs. Comparison with the CPU-executed version was also done. The report was written. Moreover, he participated in the development and implementation of new quantum chemistry and density functional methods in the Molpro quantum chemistry package.

Publications

- E. Trushin, A. Görling. *Numerically stable optimized effective potential method with standard Gaussian basis sets*, J. Chem. Phys. 155 (2021) 054109.
- S. Fauser, E. Trushin, C. Neiss, A. Görling. *Chemical accuracy with σ-functionals for the Kohn-Sham correlation energy optimized for different input orbitals and eigenvalues*, J. Chem. Phys. 155 (2021) 134111.
- J. Erhard, E. Trushin, A. Görling. *Numerically stable inversion approach to construct Kohn-Sham potentials for given electron densities within a Gaussian basis set framework*, Submitted to J. Chem. Phys.

Editor:

Prof. Dr. G. Wellein

Editorial staff:

Katja Augustin
Dr. Georg Hager
Dr. Anna Kahler
Dr. Harald Lanig

Image credits:

Prof. Dr. Rainer Böckmann (front and back cover)
MEGWARE Computer Vertrieb und Service GmbH (p. 5)
CECAM 1969 logo, CC BY-SA 4.0 (p. 44)
PD Dr. Anselm Horn (p. 46)
Marius Trollman (p. 51)

ISSN 2751-8647

Erlangen, August 2022

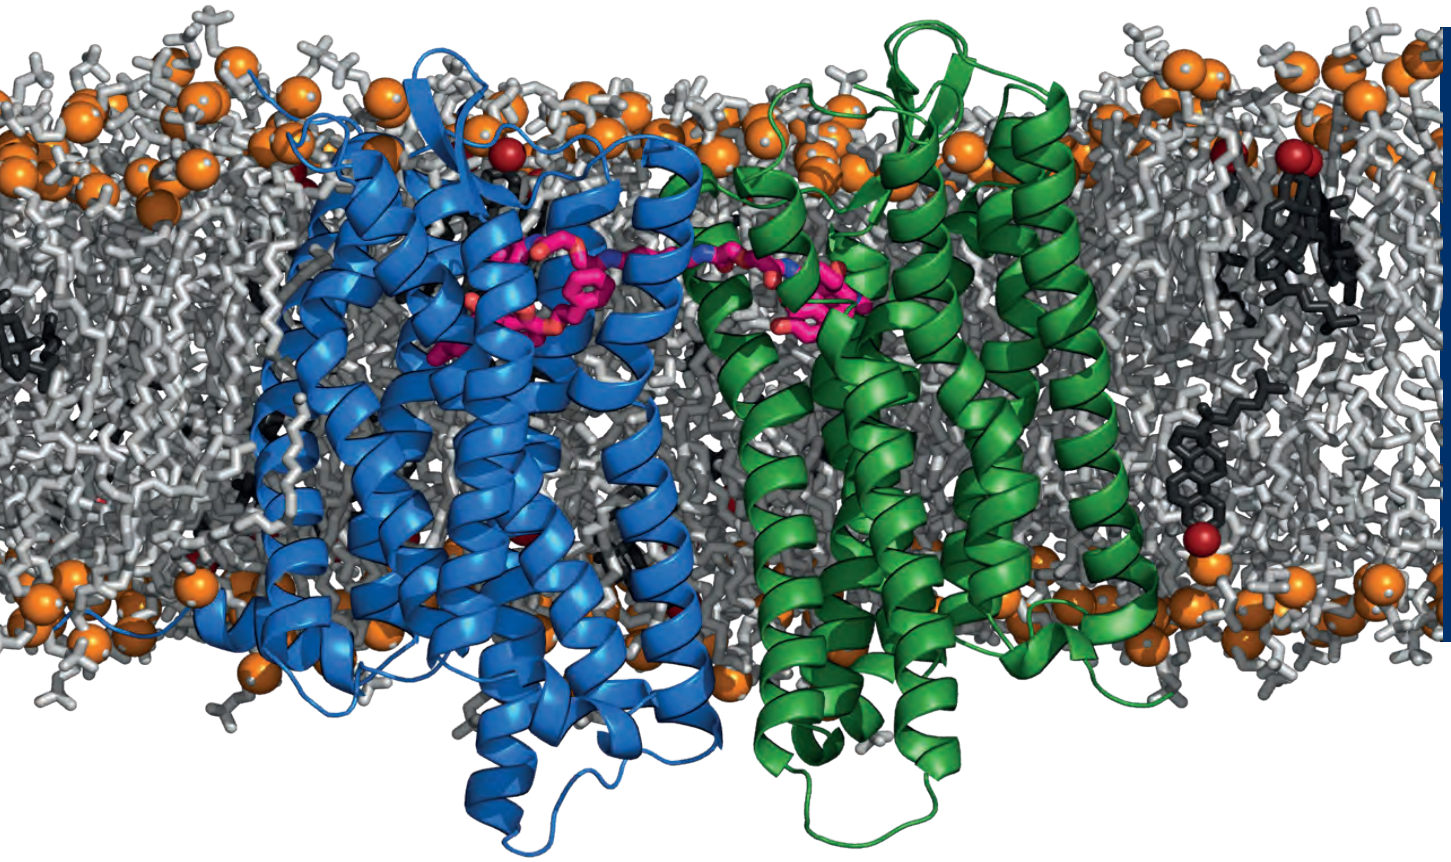

nhr.fau.de