

Brainware for Science – the Hessian HPC Competence Center

Dörte C. Sternel and Alexandra Feith

Hessisches Kompetenzzentrum für Hochleistungsrechnen, Darmstadt

In 2013 the Hessian HPC-Competence Center (Hessisches Kompetenzzentrum für Hochleistungsrechnen, HKHLR) was founded by the five universities of Hesse (TU Darmstadt, Goethe-Universität Frankfurt, Justus-Liebig Universität Gießen, Universität Kassel, Philipps Universität Marburg), to close the so called “software gap”.

This article gives an overview about the Hessian HPC infrastructure as well as the tasks and methods of the Hessian Competence Center for High Performance Computing.

1 Motivation for founding HKHLR

High performance computing (HPC) is a key technology for scientific advancement of more and more research fields. But in the context of engineering sciences, many TOP500¹ systems achieve only a fraction of their peak computation capability [1]. In other sciences one observes a similar behaviour.

Additionally, the lifespan of a typical research code is four to five times longer than the lifespan of a computing system, with generations of researchers working on a code within their research field [2]. Thus, software development often lags behind hardware development with the result that expensive hardware is often not utilised efficiently. This is called the “software gap” [3]. Only if researchers are able to realise their ideas with a justifiable amount of effort in parallel codes that can also routinely be ported to new hardware architectures, scientific computing on high performance computers can unlock its real potential.

So, one key issue for efficient computing on codes with scientific relevance and not only on benchmarks, software competence and software specialists need to be a part of the HPC ecosystem. In [4], this human expertise was called “brainware”.

Another key issue is to make researchers and their supervisors aware of the advantages of a good performance and support them to achieve knowledge about their code performance. A third issue is to enable researchers to use HPC hardware in the most efficient way according to their programming experience.

In the end the “brainware” of full time employees in HPC support saves operational costs of high performance computers, while the hard- and software runs under best possible conditions.

¹<http://www.top500.org>

The idea of the Hessian Competence Center for High Performance Computing (Hessisches Kompetenzzentrum für Hochleistungsrechnen, HKHLR) is to make HPC expertise available for all Hessian scientists.

2 HPC Ecosystem in Hesse

2.1 HPC Hardware access in Hesse

The network of Hessian High Performance Computing consists of two tier 2 high performance computers in Darmstadt (Lichtenberg High Performance Computer) and Frankfurt (LOEWE-Cluster) with a complementary hardware architecture² as well as local tier 3 high performance computers in Giessen, Kassel, and Marburg. The clusters in Darmstadt and Frankfurt are open for all research projects in Hesse; for both clusters a proposal is necessary. In Darmstadt the access underlies a scientific review procedure, which is controlled by a steering committee.

2.2 The Hessian Competence Center for High Performance Computing, HKHLR

The funding of the first period (2013–2016) of the Hessian Competence Center for High Performance Computing was fully financed by the Hessian State Ministry of Higher Education, Research and the Arts (HMWK). After the successful evaluation by the DFG (Deutsche Forschungsgemeinschaft) in 2015, the second period has been started in 2016, which is funded again by HMWK and with a growing share by the five Hessian universities. The funding allows the employment of seven full time researchers and a management position and includes also a budget for the organisation of tutorials, workshops, trainings, and public relations. The staff is distributed over the five locations; this offers a local HPC support for the researchers. A steering committee, in which all universities are represented, decides about the focus of the HKHLR.

The coordination of the Hesse-wide work is organized by the management located in Darmstadt. In regular meetings, for example weekly web conferences as well as bimonthly meetings at one of the locations, the management and staff members report about their ongoing work, discuss local events, and future tasks.

2.2.1 The HKHLR online survey

With a yearly online survey, the HKHLR monitors the user demands for support, training, and hardware infrastructure. In 2016 239 researchers participated in that survey. In the following, selected results are presented. Figures 1(a) and (b) show an overview about affiliation and disciplines of the participants.

One issue of the survey is the self-assessment about the importance of code development and parallel performance of the used software. Fig. 2 shows some results to this aspect. 72% answered that code development is part of their scientific work (Fig. 2(a)), 83% answered that the parallel performance is part of their research or important for the success of their research (Fig. 2(b)) and more than 80% gave an assessment of their code behaviour (Fig. 2(c)). Looking at Fig. 3, which visualizes the answers to the questions about further education, a gap to the previous answers becomes evident: More than 50% answered, that they are interested in workshops on the topic “What is serial/parallel performance?” (>20% “high interest”, >30% “nice to have”), and more than 60% on the topic “How to measure code performance?” (>20% “high interest,

²hardware details: <http://www.hhlr.tu-darmstadt.de> and <https://csc.uni-frankfurt.de/>

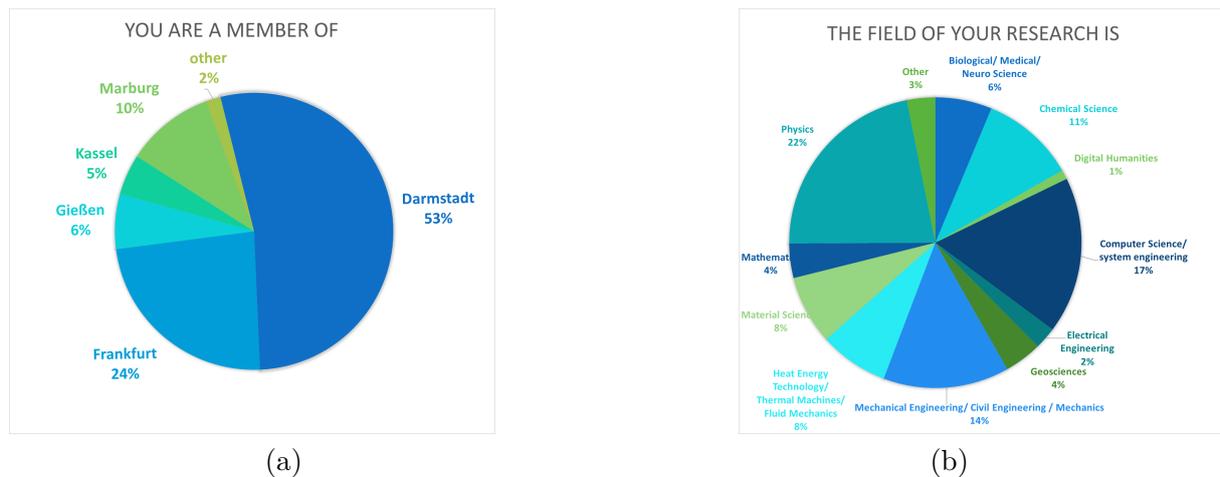


Figure 1: Participants of the online survey 2016. 239 answers.

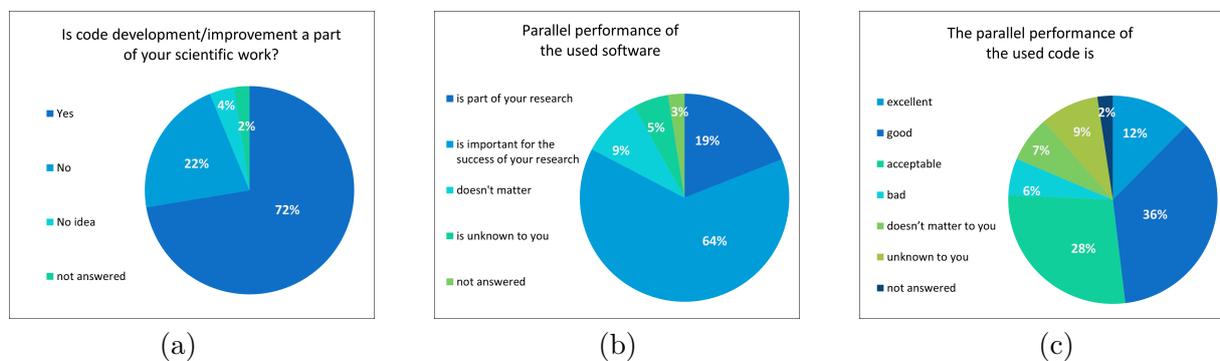


Figure 2: Code development and parallel performance. HKHLR-Online Survey 2016, 239 participants.

40% “nice to have”). This implies that many participants more or less estimated the parallel performance of their code.

About 20% of the users declare that they have a high interest in courses at beginner level, and additionally 40% of the users think that these courses are “nice to have” (Fig. 3). This shows, that for a high number of HPC users, there is a big need of courses at a relatively low level.

2.2.2 Monitoring of HPC usage

For optimizing the usage of the Hessian HPC facilities, the HKHLR management initiated a central monitoring, which will be continuously improved.

This job monitoring, in its current form, is used to detect imbalances between required and used resources. In case of imbalances, the HKHLR staff can contact the user to clarify the reasons, and help to improve the usage.

Furthermore, to give tuning experts the information where substantial performance improvements can be achieved, monitoring helps to identify “power users”. Figure 4 shows the accumulated resource demands of the projects at Lichtenberg High Performance computer, ordered by decreasing resource demands. The first 10 users request around 50% of the complete cluster (cf. [4], with a similar situation).

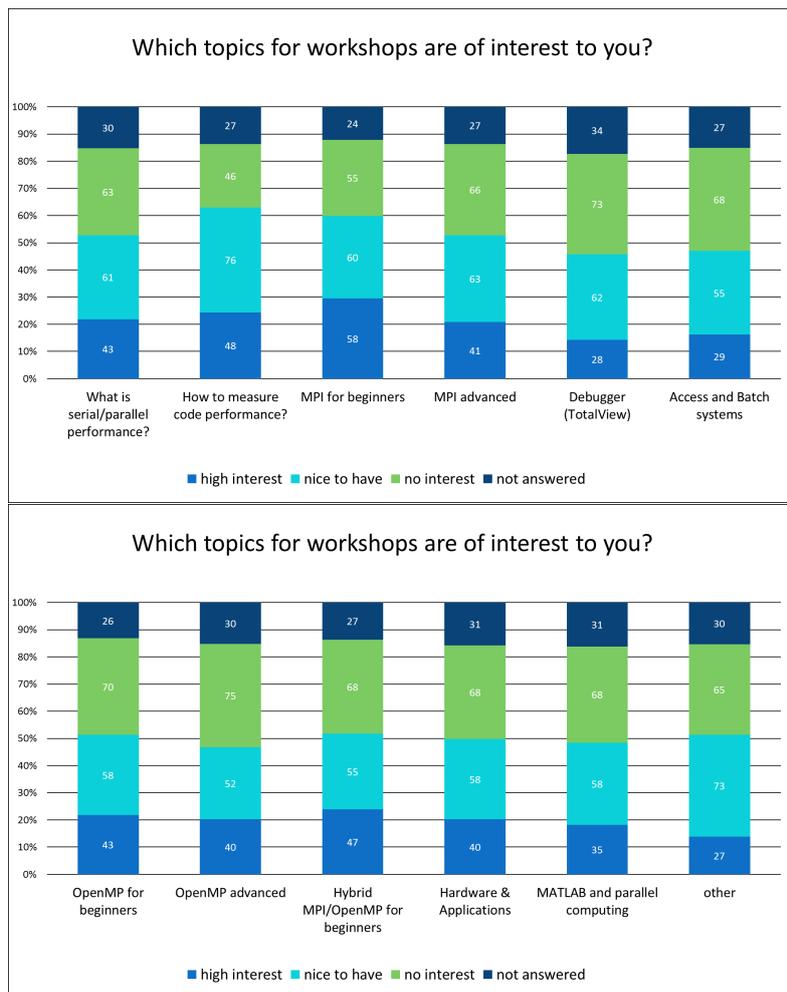


Figure 3: Demands for advanced trainings. HKHLR-Online Survey 2016, 197 answers.

Another advantage of a central monitoring in Hesse is to control free resources. When there is an overload on one location and free resources on one of the others, it is often possible to satisfy short term necessities.

2.2.3 Personal support

A user can request personal support of HKHLR via the ticket system of the local HPC cluster.

If the HKHLR staff identifies potential candidates for performance improvement by the analysis of the monitoring data (see section 2.2.2), they will contact the user. In most cases small changes in the batch script or of the compiler options can save a lot of resources.

The support of power users is a greater challenge. Here, the researchers are addressed directly. Together with HKHLR staff members, the parallel performance of the code will be analyzed, bottlenecks identified, and suggestions for improvement developed.

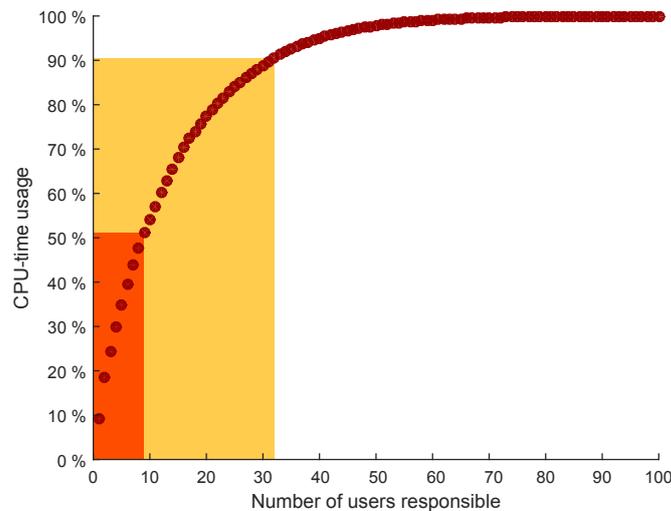


Figure 4: Accumulated resource demands of the users at Lichtenberg High Performance Computer.

2.2.4 Workshops, training, courses and networking

The workshop series HiPerCH (High Performance Computing Hesse) takes place two times a year. The experience of the last two years showed that researchers prefer local workshops. Therefore, the location of HiPerCH workshops alternates. The topics of HiPerCH are chosen by the user demands extracted from the online survey, from experience of HKHLR staff, or from direct user requests. These workshops cover basic as well as advanced topics for specific audiences.

As a result of the online surveys, HKHLR initialized introductory courses for new users of the Hessian HPC clusters, covering information about general access, available hardware, and usage. Basic introductions to Linux, shell scripting, and other basic tools are part of the introductory courses and are completed by offering more detailed knowledge by the HiPerCH workshop series.

On demand, the HKHLR organizes special trainings as well as user group meetings.

2.2.5 Unified Usage

To simplify the use of other hessian HPC clusters, the HKHLR aspires to unify the systems, whenever it is reasonable and technically feasible. One first step is to use the same scheduler at all locations as soon as possible.

To push that process, the HKHLR organizes regular meetings of the Hessian HPC-system administrators.

2.2.6 Tools

With the funding of the HMWK, the HKHLR provides the license for the performance analysing tool Vampir³ and the debugger Totalview⁴ for all Hessian Universities. Those tools and individual user support, either on request by users or by identifying power users (see section 2.2.1),

³<https://www.vampir.eu/>

⁴<http://www.roguewave.com/products-services/totalview>

foster competences in analyzing or tuning parallel tools. Local trainings empower the researchers in the use of these tools.

2.2.7 Public relations

Besides scientific publication, the awareness of Hessian research on HPC in public is one aim of HKHLR. Therefore, HKHLR offers a platform for HPC in Hesse. The Web page (<http://www.hpc-hessen.de>) summarizes information about the goals and benefits of HKHLR and gives an overview of Hessian research by HPC reports, current research projects as well as video portraits of Hessian researchers. HKHLR assists researchers by the distribution of their research for the general public. In addition, it is the first point of contact for new users and offers helpful links, flyer material, and an event calendar. Since 2015, HKHLR has an own booth at ISC High Performance in Frankfurt (<http://isc-hpc.com/>).

Acknowledgements

HKHLR is funded by the Hessian State Ministry of Higher Education, Research and the Arts.

References

- [1] D.F. Harlacher, S. Roller, F. Hindenlang, C.-D. Munz, T. Kraus, M. Fischer, K. Geurts, M. Meinke, T. Klühspies, V. Metsch, and K. Benkert. “Highly Efficient and Scalable Software for the Simulation of Turbulent Flows in Complex Geometries”, in W.E. Nagel et al. (eds): High Performance Computing in Science and Engineering, Part 5, pp. 289-307, DOI 10.1007/978-3-642-23869-7_22, Springer Verlag, 2012.
- [2] D.E. Post, R.P. Kendall, Software project management and quality engineering practices for complex, coupled multiphysics, massively parallel computational simulations: lessons learned from ASCI. *Int. J. High Perform. Comput. Appl.* 18, pp. 399-416, 2004.
- [3] Computational Science: Ensuring America’s Competitiveness, President’s Information Technology Advisory Committee. https://www.nitrd.gov/pitac/reports/20050609_computational/computational.pdf, 2005.
- [4] C. Bischof, D. an Mey, and C. Iwainsky, “Brainware for green HPC”, *Computer Science – Research and Development*, DOI 10.1007/s00450-011-0198-5, Springer, 2011.