

Performance Conscious HPC (PeCoH) – 2019

Kai Himstedt¹, Nathanael Hübbe¹, Sandra Schröder¹, Hendryk Bockelmann², Michael Kuhn¹, Julian Kunkel³, Thomas Ludwig^{1,2}, Stephan Olbrich¹, Matthias Riebisch¹, Markus Stammler⁴, Hinnerk Stüben¹



¹ Universität Hamburg ² Deutsches Klimarechenzentrum (DKRZ)
³ University of Reading ⁴ Technische Universität Hamburg

Contact: kai.himstedt@uni-hamburg.de

INTRODUCTION

This poster presents project outputs and ongoing activities. For the sake of completeness, it also gives a brief overview of the major project goals and partners involved in the project.

Data centers often face **similar challenges when supporting analysis and optimization** of relevant applications. From the user perspective, the **benefit of performance engineering is difficult to quantify** and so are the costs involved in this tuning/optimization.

The German state of Hamburg has **three data centers, each independently providing compute power and support** to their users. Two general purpose data centers are responsible for the needs of Universität Hamburg and Technische Universität Hamburg, respectively. The third data center is a Tier-2 center supporting climate research. Although near to each other, the collaboration of support staff between the data centers has been limited. As part of the project, this **collaboration will be strengthened**.

GOALS

The objectives of PeCoH are to

1. raise awareness and knowledge for performance engineering and to
2. coordinate performance engineering within Hamburg's institutions.

To reach these goals, we have established the **Hamburg Regional HPC Competence Center** conduct various practical activities on established techniques but also explore some supportive and alternative approaches as research.

PARTNERS

German Climate Computing Center (DKRZ) DKRZ is a partner for Climate Research, providing tools and the associated services to investigate the processes in the climate system. The HLRE-3 supercomputer Mistral consists of more than 3,000 compute nodes, providing a peak compute performance of 3.6 PFLOPs. The system is backed by a 60 Petabyte Lustre file system.



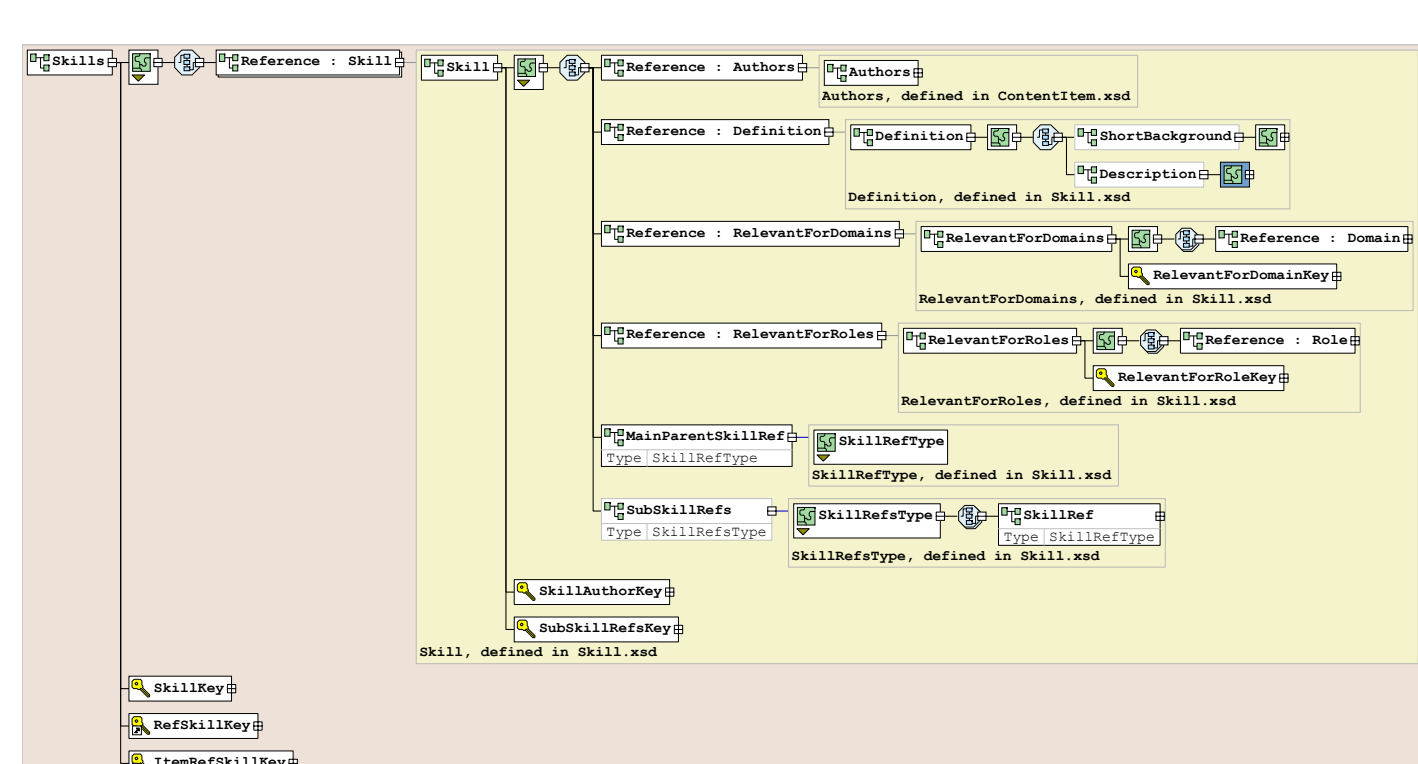
Regional Computing Center at Universität Hamburg (RRZ) RRZ provides many central IT services to Universität Hamburg. Amongst these, it operates an HPC cluster with 396 nodes and 400 Terabyte BeeGFS file system.



Technische Universität Hamburg (TUHH-RZ) RZT provides central IT services to the Technische Universität Hamburg. It operates a 244 node cluster with a 250 Terabyte BeeGFS file system.



SKILL TREE MODELING




XML Schema Definition for Showing the Essential Skill Tree Structure

PERFORMANCE AWARENESS: PROVIDING COST FEEDBACK

The usual approach to accounting on a scientific supercomputer is that the scientists apply for compute time with their projects, get their projects granted, and then use their compute time to reach their project goals. Grants usually include an amount of node hours and a quota for disk storage.

This approach has several problems, which essentially boil down to the fact that scientists only works against limits. The true costs associated with their actions are not visible to them. Our cost modeling efforts were aimed at providing more information about the true costs to HPC users, enabling them to make more informed decisions. On the other hand, our cost models can also be used to derive statistics of the use of an HPC system. The figure on the right shows the results of performing this analysis on a day's worth of jobs.



Statistical analysis of costs per job that were executing during a single day on a supercomputer, each line shows the result of a different cost model

Most interestingly, the vast majority of jobs does not consume any significant amount of resources, the costs are driven by the few very large and long running jobs, which are usually prepared by much smaller jobs. As such, it seems very important to educate users to efficient HPC usage during the time span in which they are still running many small jobs, before they start their really big runs. This can have a positive effect on the scientists ability to make the best scientific use from an HPC system.

The used scripts and a docker image to test the SLURM extension are available online at <https://github.com/pecoh/cost-modelling>

DEVELOPMENT OF THE HPC CERTIFICATION PROGRAM

Users should have a set of validated skills before they start using an HPC system.

For the classification of HPC competences and the definition of certificates we implemented an HPC skill tree (based on XML).

Beside its name and description, a skill in the tree has additional attributes to describe e.g. its special significance to a scientific domain, in order to

- enable the creation of different views of the skill tree to consider the users' varied backgrounds
- give the user an overview of custom-tailored skills which he has to acquire to pass the exams

A certificate definition

- bundles an appropriate set of HPC skills
- is used to certify a users' HPC qualification by successful exams

For the development of the certification program, we separated the certificate definition from the providing of content, similar to the concept of a high school graduation exam (Zentralabitur in German).

We helped to setup the HPC Certification Forum, an international board that will sustain this activity and curate the skills, certificate definitions, and exam questions.

The certification board has the role of a (virtual) central authority, which results in major advantages:

- Generally accepted certificates and corresponding exams can be established by the HPC Certification Forum without the burden of being responsible for the learning material
- The learning material can be provided by different scientific institutions

The recently kick-started HPC Certification Forum (HPC-CF) welcomes international collaborators.

ONGOING ACTIVITIES

Tuning: We have successfully started our examination with performance improvements for examples using the language R, along with efficient libraries like OpenBLAS or MKL and parallelizing loops with the `foreach()` paradigm in OpenMP and MPI environments.

For collecting further success stories we will focus on the following tuning possibilities to transfer performance engineering know-how:

- Tuning without the need to (re-)build a parallel program, e.g. by tuning a parallel program from the outside via runtime options
- Tuning without modifying the source code, e.g. by using optimized libraries and setting appropriate compiler/linker options
- Tuning of individual software based on co-development with users

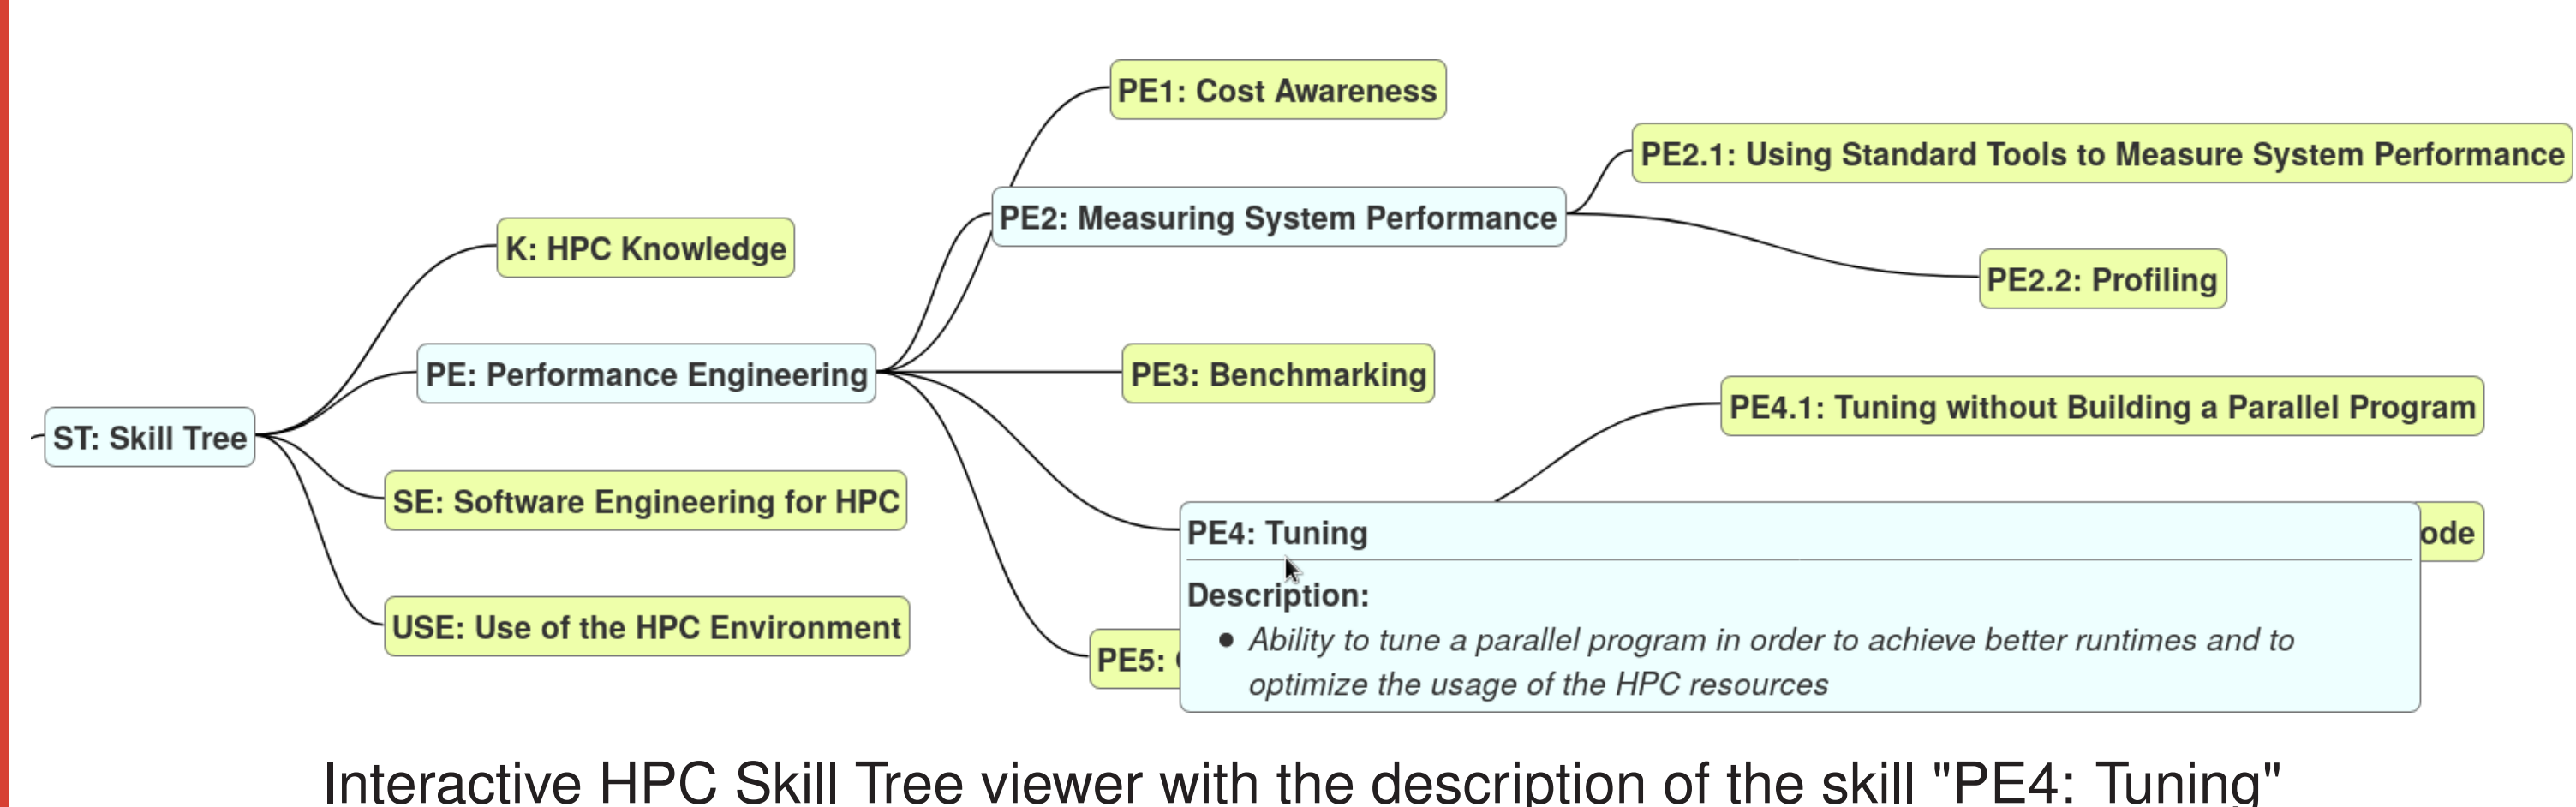
Workshop Material and Online Tutorial: In the PeCoH project we have a focus on the HPC Certification Program and the role as certification provider. As we have additionally the role as content provider for basic HPC skills in the project, we designed a workflow based on XSLT programs and the well-known Pandoc-Tool to convert the Markdown content files associated with each single skill to various target formats like HTML, PDF, and TeX.

To efficiently equip new users with certified HPC skills, we will

- finish the coordination with other stakeholders that may provide content
- work on the completion of the complementary open source workshop material
- develop a multiple choice questionnaire to validate the basic HPC skills

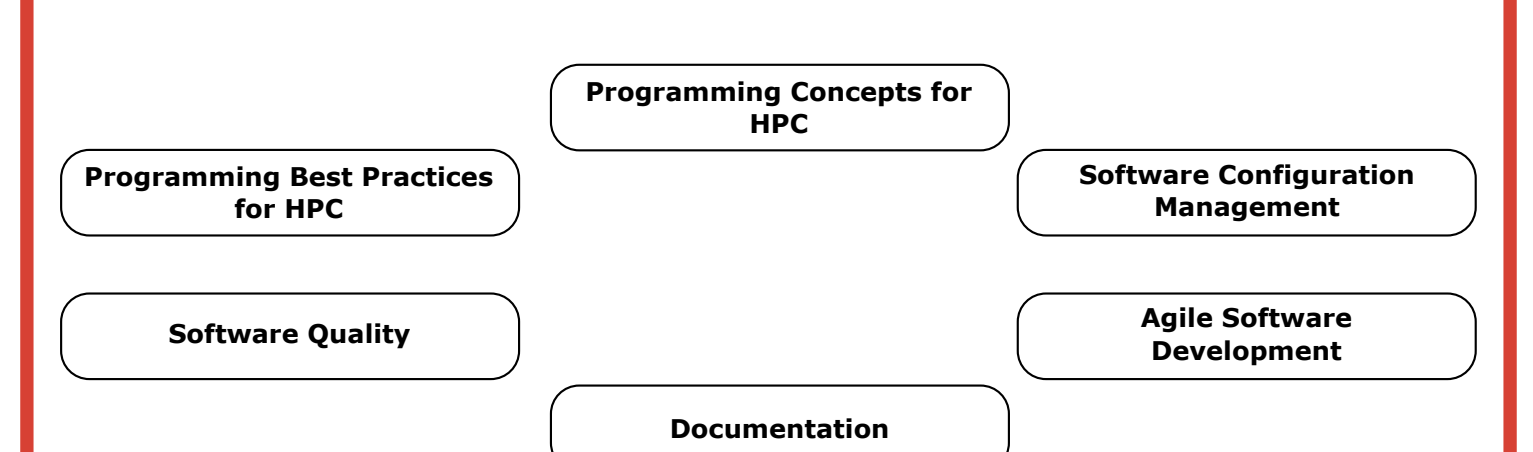
HPC SKILL TREE

The developed HPC Skill Tree contains descriptions of the various skills that are relevant in the field of HPC (we are still in the process of refining it and open to suggestions). The skills will be linked to the relevant teaching material. Our current HPC Skill Tree is available online via our interactive skill tree viewer (<https://www.hbcc.uni-hamburg.de/en/hpc-certification-program/hpc-skill-tree.html>). This viewer can easily be included into third party web-sites and can be configured to provide different views of the skill tree data.

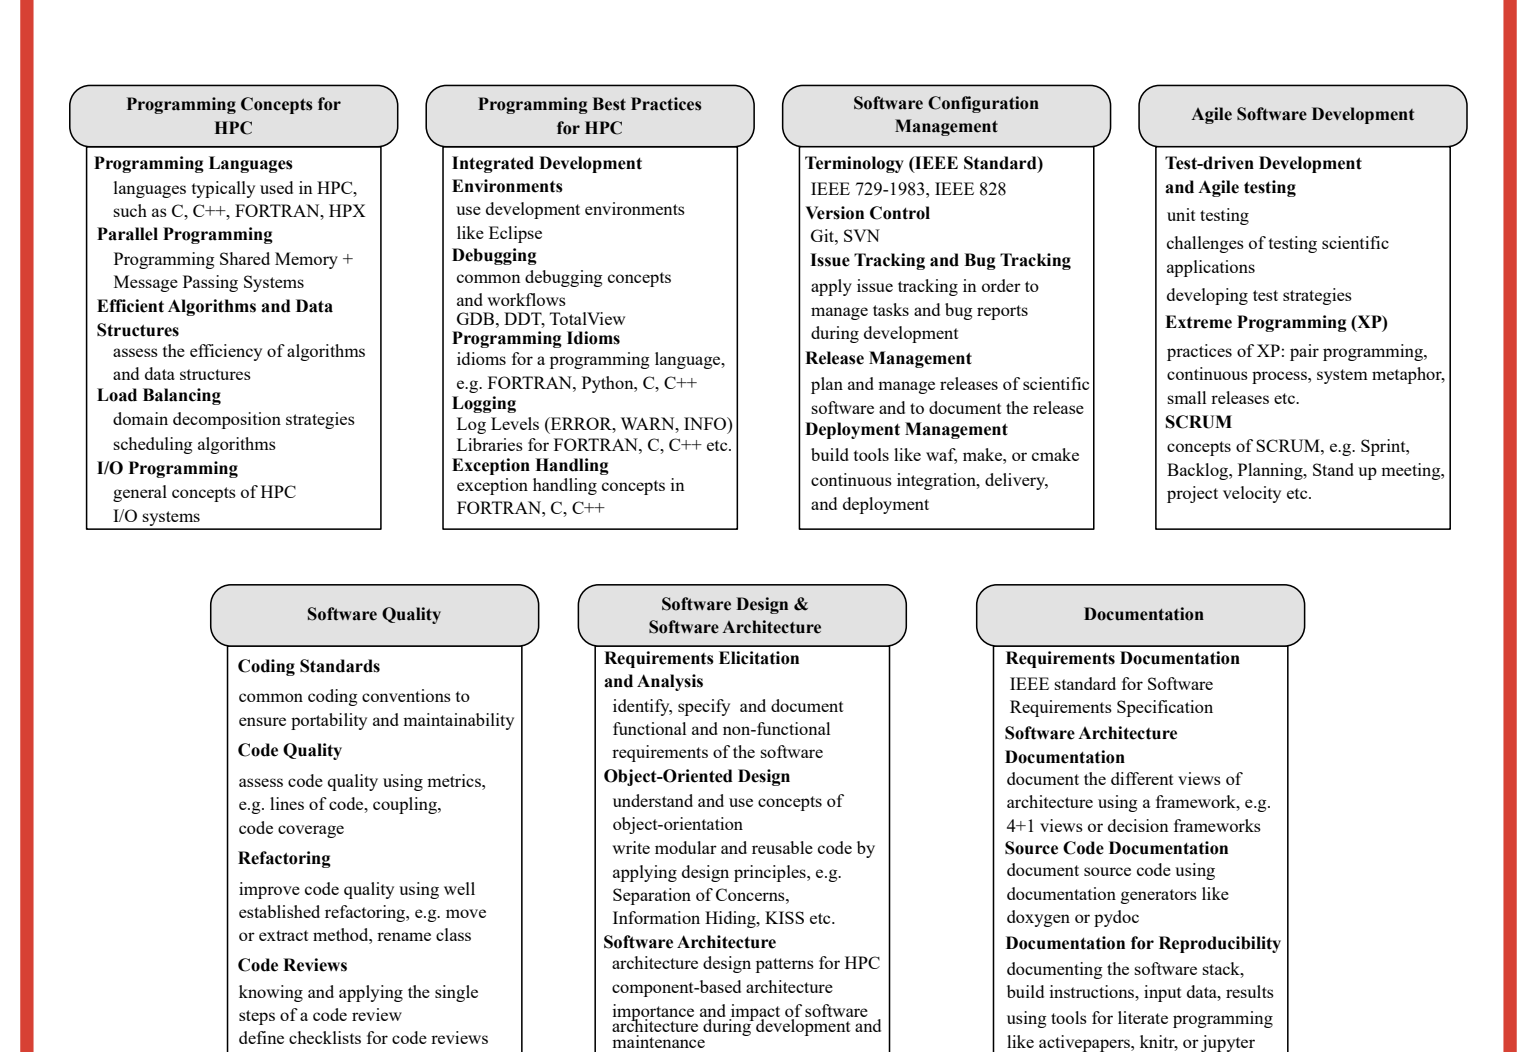


HPC SOFTWARE ENGINEERING

This activity broadens the use of software engineering techniques in the field of HPC in order to increase the performance of parallel programs.



Efficiently Create, Maintain and Reuse Code



Assess Practicability with Scientists

ACTIVITIES TO DATE

We participated for the presentation of the PeCoH project in the

- ISC 2017
June 18–22, 2017, Frankfurt, Germany
 - PeCoH project poster
- FEPA workshop
July 20–21, 2017, Erlangen, Germany
- 7th Gauß Allianz HPC-Status-Conference
December 4–5, 2017, Stuttgart, Germany
- ISC 2018
June 24–28, 2018, Frankfurt, Germany
 - PeCoH project poster
- 8th Gauß Allianz HPC-Status-Conference
October 8–9, 2018, Erlangen, Germany
- SC 2018
November 11–16, 2018, Dallas, Texas
 - Extended abstract titled "Towards a HPC Certification Program"

All contributions are available for download.¹

A concept paper for the HPC certification program is available for download.²

We are hosting a mailing list for the HPC certification program.³

In the sense of a spin-off of the PeCoH project we established the HPC Certification Forum.⁴

We submitted a BoF (Birds of a Feather) titled "HPC Certification Program" in the field of education and training for the ISC 2019 together with international partners.

REFERENCES

- ¹HHCC. *Hamburg HPC Competence Center – Downloads*. <https://www.hbcc.uni-hamburg.de/en/support/downloads.html>
- ²HHCC. *Hamburg HPC Competence Center – HPC Certification Program: Concept Paper - Draft Version 0.91 – June 1, 2018*. <https://www.hbcc.uni-hamburg.de/en/files/hbccp-concept-paper-180601.pdf>
- ³HHCC. *Hamburg HPC Competence Center – Mailing List of the HPC Certification Program*. certification.hbcc@lists.uni-hamburg.de
- ⁴HPCCF. *The HPC Certification Forum – Home Page*. <https://www.hpc-certification.org>

ACKNOWLEDGEMENTS

This work was supported by the German Research Foundation (DFG) under grants LU 1353/12-1, OL 241/2-1, and RI 1068/7-1.

Please visit also the HPC Certification Forum <https://hpc-certification.org>.