

Performance Conscious HPC (PeCoH)

Julian Kunkel¹, Michael Kuhn², Thomas Ludwig^{1,2}, Matthias Riebisch², Stephan Olbrich²,
Hinnerk Stüben², Kai Himstedt², Hendryk Bockelmann¹, Markus Stammberger³

¹ Deutsches Klimarechenzentrum (DKRZ)

² Universität Hamburg

³ Technische Universität Hamburg

Contact: kunkel@dkrz.de

INTRODUCTION

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.

Existing codes and workflows must often be **adjusted in non-trivial ways** to explore the benefits of novel concepts and emerging technologies, which often **causes users to hesitate**. It is important to better understand these cost drivers in order to ultimately **increase the scientific output of data centers**.

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 will establish the **Hamburg Regional HPC Competence Center**.

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.



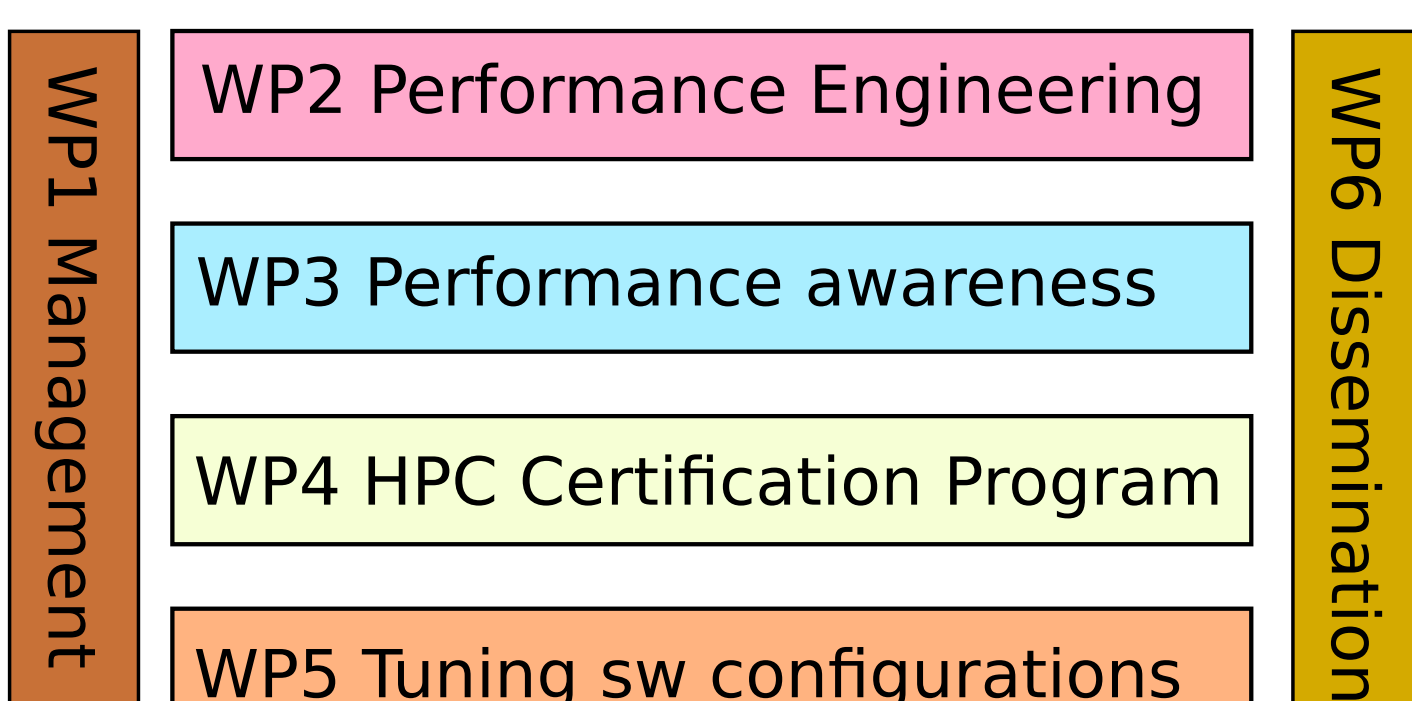
Regional Computing Center at the Universität Hamburg (RRZ) RRZ provides many central IT services to Universität Hamburg. Amongst these, it operates an HPC cluster with 396 nodes and 400 Terabyte BeeGFS file system.



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.



WORK PACKAGES



CONCEPTS

The project implements traditional **services** across the virtual competence center and researches **methods** to improve the current state of the art in performance engineering. More details about selected activities are provided in the boxes on the right.

Services

- **Dissemination** of knowledge, concepts and chances of performance engineering
 - Web page: Hamburg HPC competence center
 - Publication of success stories
 - Cross references to regional HPC sites, consultants and material
- **Education** of performance engineering skills
 - Workshop and online courses
- Pro-active **feedback** of job performance
- **Support** for performance analysis and optimization
- **Co-development** to evaluate promising concepts

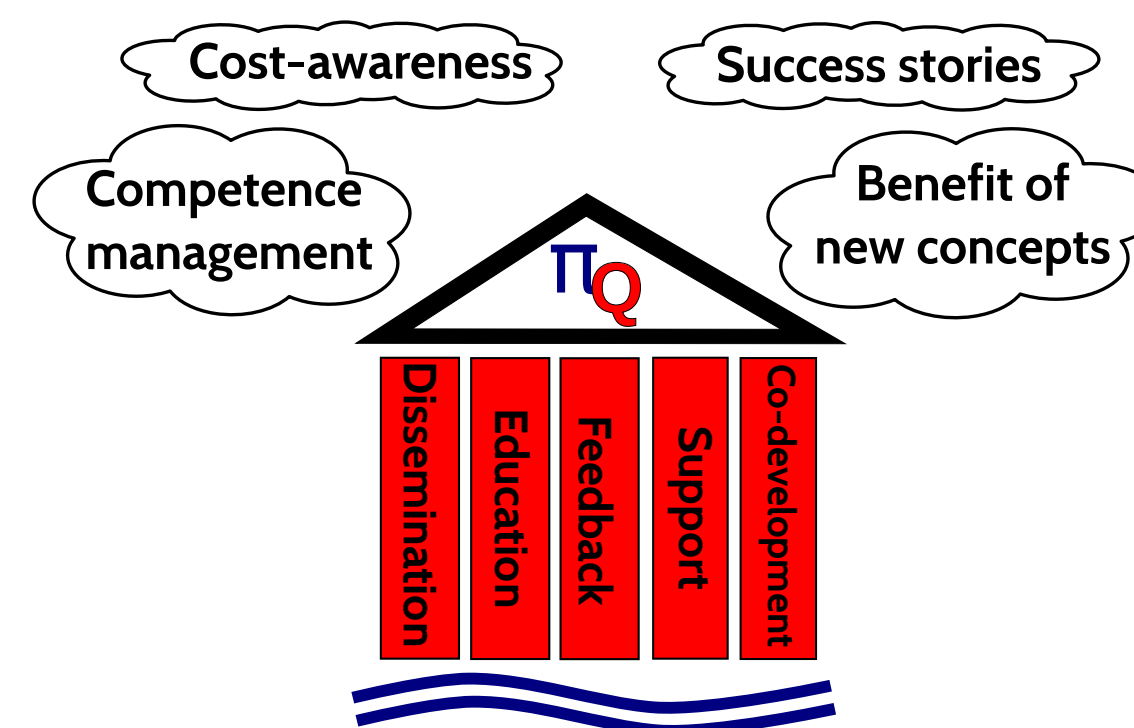


Fig. 1: PeCoH implements services (at the bottom) and researches methods (clouds).

Methods

- **Competence management** aims to identify key skills necessary for HPC
 - Develop an **HPC certification program**
 - Research methods to automatically create a knowledge base (link to existing knowledge)
- **Cost-awareness** allows predicting job/workflow execution costs with the help of models
- **Benefit of new concepts** assesses the potential of alternative architectures / software
 - Develop models to predict costs for porting / benefit for users
 - Part of general performance engineering support
- **Success stories** illustrate approaches, benefits and costs of selected optimizations

STATUS AND FUTURE WORK

The PeCoH project started on March 1, 2017 and has a duration of three years (one PhD candidate and a postdoctoral full time position). For processing the work packages three main topics can be emphasized for scientific studies:

Establishing Software Engineering for HPC

Our observations show that parallel programs are often developed in the same way as decades ago. One obvious reason is that in the scientific computing community the focus lies on publishing scientific results and not on the software that was used to obtain these results. In the past the software engineering community did not work effectively enough in the field of HPC. Therefore in the PeCoH project we want to measure the positive impact of using software engineering on scientific productivity, e.g.

- Efficient Algorithms and Data Structures
- Object Oriented Development
- Agile Software Development, Automated Testing / Test-Driven Development
- Coding Guidelines, Refactoring
- Version and Configuration Management

Development of a Cost Model

Our cost model will be based mainly on

- resource usage of the batch jobs
- time spent in different tasks to rewrite existing code on pilot studies applying performance engineering concepts

With the subsequent cost analysis we can estimate the cost-benefit ratio of novel approaches. Thus we can raise the awareness and knowledge of users for performance engineering, i.e., assist in identification and quantification of potential efficiency improvements in scientific parallel codes and parallel code usage.

Transfer of HPC Know-how and Providing Tuning Principles

According to our experience problems of the scientists are often dominated at the level of getting things to work, i.e. the user focuses on getting a parallel job to run, rather than being aware of how to use the expensive HPC resources appropriately. Hamburg HPC Competence Center (HHCC), as a virtual institution, takes this into account and provides basic and advanced HPC knowledge to improve the situation via online content e.g. for

- Basic Competences: Linux command line, shell scripts, robust job scripts, job chaining
- System Building: compilers, optimization flags, linker, libraries, debugger, code optimization
- Hardware Architectures: shared memory systems, distributed systems, hybrid systems
- I/O Operations: storage systems, choice of file sizes vs. number of files, compression
- Detecting Performance Issues: measuring speedups (compared to *best known* sequential algorithm) and efficiencies, monitoring resource utilization on the application level

HAMBURG HPC COMPETENCE CENTER (HHCC)

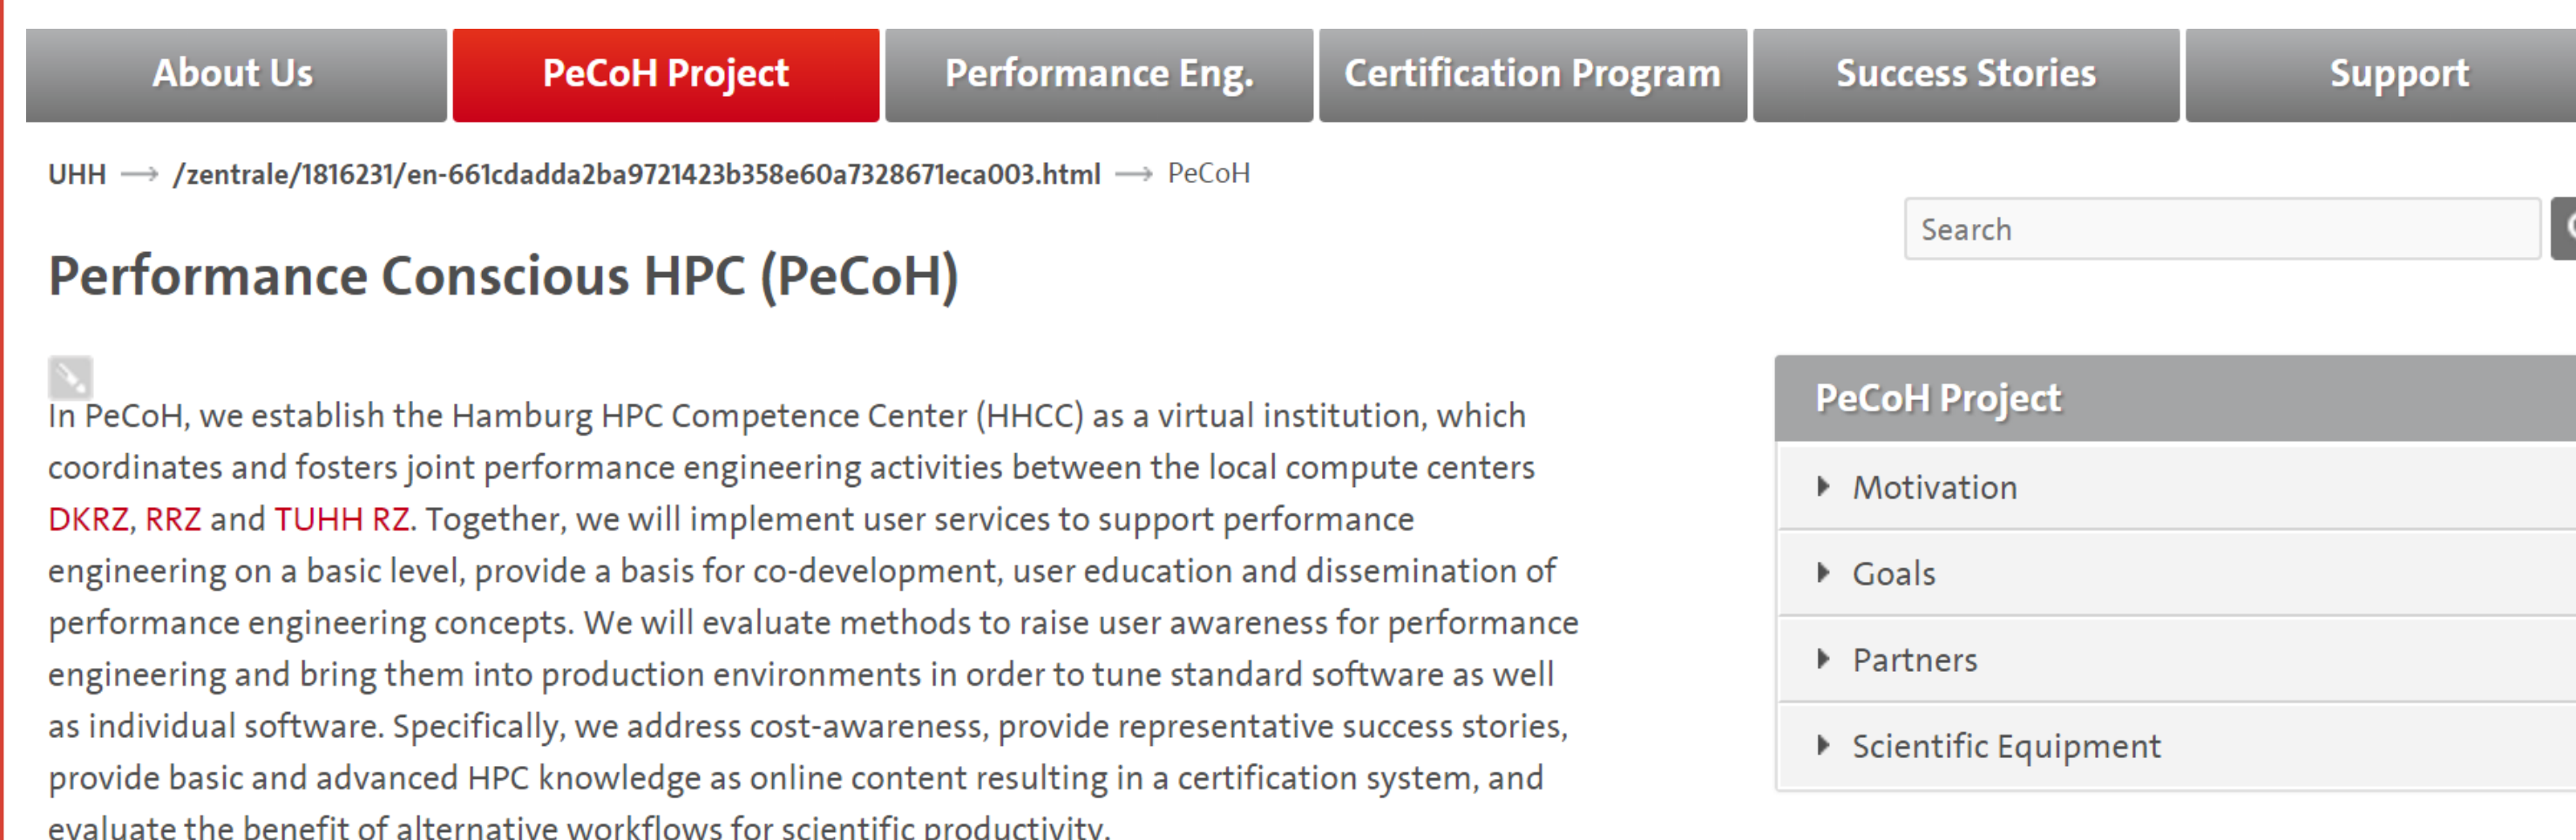


Fig. 2: The screenshot gives an impression of the HHCC Website which is currently being constructed.

PERFORMANCE ENGINEERING

This work package provides the new and adapted concepts of **performance engineering** for the target domain and, thus implements the **co-development service** and the method benefit of novel concepts. We will

1. identify suitable **performance engineering frameworks**,
2. relate typical **analysis patterns to performance deficiencies**,
3. **evaluate state-of-the-art approaches** of partitioning patterns/procedures,
4. provide means to **monitor and visualize performance** to identify relevant issues,
5. identify **best practices**,
6. **evolve coding guidelines** according to typical performance issues and
7. **forecast future performance** by applying performance models and prototypes.

Based on the results of models and prototypes:

- we extract knowledge about the effectiveness of concepts for different domains
- we will periodically capture the fraction of
 - work time spent in different tasks
 - estimated or measured performance

This will be used to quantify the cost-benefit better and to identify cost-effective strategies.

PERFORMANCE AWARENESS

This work package implements our method for **cost-awareness** and combines it with the service for providing feedback. We will

1. develop models to **approximate costs** for scientists,
2. **embed these models into the resource manager SLURM**,
3. track and report the **resource utilization and efficiency of applications and data centers**,
4. work on **tools to detect common performance issues** (e.g., by misconfigurations),
5. provide **feedback to the users**.

HPC CERTIFICATION PROGRAM

This work package covers the method **competence management** and the education service. Our approach covers

1. an analysis and classification of **competences and their value for scientists**,
2. the development of a **program to check and improve on these competences**,
3. the collection of **material to improve base knowledge** on these topics,
4. an **online tutorial** packaging lists of relevant material and online content,
5. **online examinations** to obtain the entry-level of the certificates.

Examples for **relevant competences** are command line use, scripting, batch systems, parallel execution, scaling, understanding resource usage, compilers and performance engineering.

This will not only help users but also the support staff to better **understand the experience and needs of the users**.

TUNING OF SOFTWARE

This work package provides a help desk and determines high-level tuning possibilities of frequently used software. For example: Amber, Gromacs, Abaqus, OpenFOAM, Matlab, R.

Resulting in the documentation of best practices to avoid common (performance) mistakes.

COLLABORATION

We are looking for **collaboration partners** to leverage available concepts for services and to evaluate research concepts such as the cost-awareness modifications and the HPC certification program. Contact: kunkel@dkrz.de

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.