Performance Conscious HPC (PeCoH) – 2018

Kai Himstedt[1], Nathanael Hübbe[1], Sandra Schröder[1], Hendryk Bockelmann[2], Michael Kuhn[1], Julian Kunkel[2], Thomas Ludwig[1,2], Stephan Olbrich[1], Matthias Riebisch[1], Markus Stammberger[3], Hinnerk Stüben[1]

[1] *Universität Hamburg* [2] *Deutsches Klimarechenzentrum (DKRZ)* [3] *Technische Universität Hamburg*

Contact: kai.himstedt@uni-hamburg.de

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

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 have established the **Hamburg Regional HPC Competence Center** and are developing an **HPC certification program**.

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.

DKRZ
DEUTSCHES KLIMARECHENZENTRUM

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.

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

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.

TUHH
Technische Universität Hamburg

Skill Tree Modeling

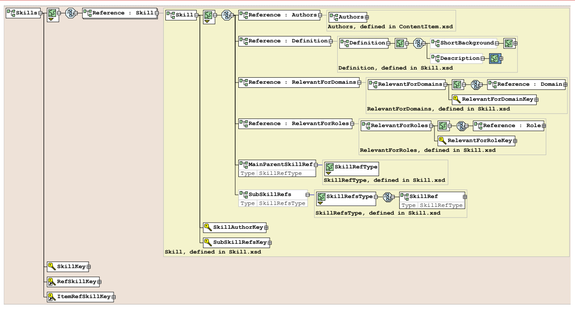

XML Schema Definition for Showing the Essential Skill Tree Structure

Introduction

Computing Centers provide great services to scientists, allowing them to run their simulations on tax payer money. The usual approach 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. The common unit for compute time accounting is the node hour. In addition to that, scientists are given a quota of data that they are allowed to store on the file system.

This approach has several problems: (1) It does not include energy consumption; (2) it cannot account for different hardware, like fat nodes with more memory, or visualization nodes with graphic cards; (3) the actual file system usage is irrelevant as long as the quota is not exceeded; (4) and once compute time and storage space are granted, there is no incentive not to use the full grant. In short, the questions scientists typically ask is "am I still within the limits?", while in this project, our goal is to raise awareness to a second question: "do I use resources efficiently?" Asking the second question should lead to better economic choices when using HPC systems.

To this end, we have implemented a SLURM job epilogue script that analyses three parameters of the finished job: (1) runtime (2) allocated nodes and (3) energy usage. These serve as inputs to a simple cost model that is fully configurable by the system administrators, and which produces a report that lists the different costs that were caused by the job. On the right is an example of what the report that is provided to the user might look like, the names of the items, their rates, and the currency used are entirely up to the system administrators. In this case, it should be clear to scientists reading the report that they can significantly reduce costs by requesting less memory.

```
==================================
job cost estimate
==================================

(10 nodes total)
BasicNode (10 nodes):
    Hardware:              2.25 Euro
    Power:                 1.25 Euro
FatNode (10 nodes):
    ExtraRam:              4.00 Euro
GpuNode (2 nodes):
    ExtraHardware:         3.50 Euro
----------------------------------
total:                    11.00 Euro

==================================
```

The same cost model configuration file can also be used by a companion script to analyse jobs in the SLURM job database post-mortem, providing a statistical overview to the system administrators.

HPC Certification Program

Computing power and complexity of HPC systems are steadily increasing. This leads to an increasing demand for a good education of their users so that they can use such systems adequately. Our HPC Certification Program approach takes the users' varied backgrounds (e.g. research area and prior knowledge) into account and focuses on performance engineering to enable them to achieve further speedups for parallel applications with efficient utilization of the HPC resources.

Classical HPC Education

Institutions which operate HPC systems usually offer regularly recurring teaching events about general aspects of Supercomputer hard- and software architectures and parallel programming at beginners' level as well as higher levels. Classical HPC education is based on

- Lectures, tutorials, and workshops, having a rather static character
- Websites offering (online) HPC learning material in a more dynamic way

However, in contrast to other areas of information technology (IT), where certificates are often used to prove IT skills of the users, in the field of HPC neither commonly accepted standards exist, nor a certification program for the education.

New Approach

We named our HPC Certification Program "HPC-Führerschein" (HPC driving licence in English) to point out that users should have a set of validated skills before they start using an HPC system.

In our approach the certificate definition is separated from content providing. While the certification board has the role of a (virtual) central authority, the learning material can be provided by different content providers, e.g. by different scientific institutions.

Set of terms for separating the definition of certificates from content providing:

Skill: The abilities and the knowledge specified in the skill description
Gained Skill: Skill of a user validated by a test (multiple-choice based, in our case)
Content: Learning material enabling the user to gain skills
Content provider: Institution that provides content
Exam: Process to validate the user's skills
Certificate definition: Set of skills as specified in the description of the certificate
Certification provider: Institution that suggests certificate definitions and corresponding exams
Certification board: Institution that establishes accepted certificate definitions and corresponding exams
Certificate: Document certifying the validation according to the corresponding certificate definition by a successful exam

By its role as a (virtual) central authority the certification board has the power to establish generally accepted certificate definitions and corresponding exams without the burden of being responsible for the content.

HPC Skill Tree

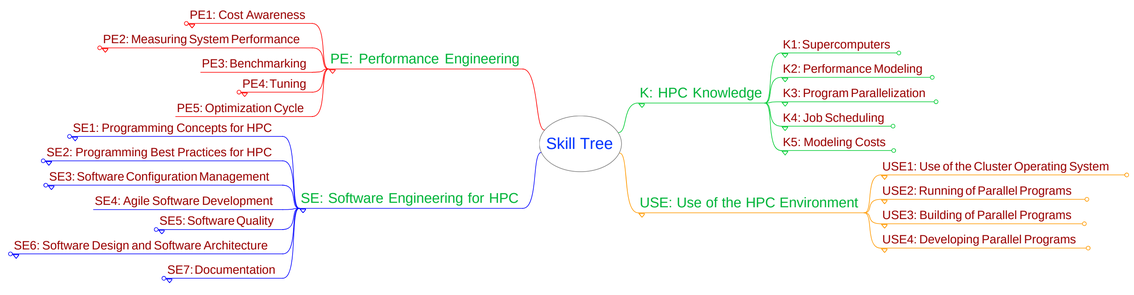

Tree of the Top Level Skills

Beside its name and description, a skill in the tree has additional attributes to describe e.g. its special significance to a scientific domain. Such information can be easily used to create different views of the skill tree in order to consider the users' varied backgrounds and to give the user an overview of those custom-tailored skills which he has to acquire to pass the exams.

HPC Software Engineering

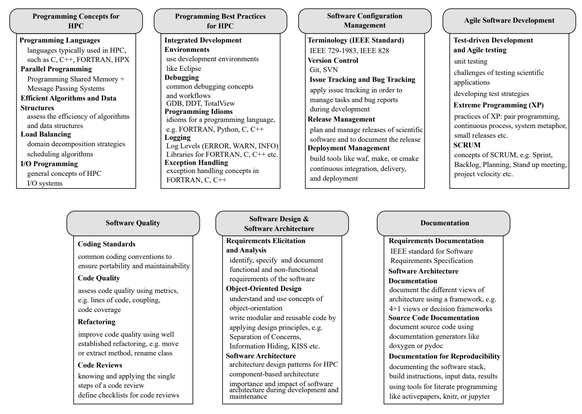

Broaden the use of software engineering techniques in the field of HPC in order to increase the performance of parallel programs.

Activities to Date

We particitpated for the presentation of the PeCoH project in the

- ISC 2017
  June 18–22, 2017, Frankfurt, Germany
  - PeCoH project poster[6]
  - Handout to the work in progress of our HPC Certification Program[3]
- FEPA workshop
  July 20–21, 2017, Erlangen, Germany[1]
- 7^{th} Gauß Allianz HPC-Status-Conference
  December 4–5, 2017, Stuttgart, Germany[2]

A concept paper for the HPC certification program is available for download.[4]

We are hosting a mailing list for the HPC certification program.[5]

We submitted a BoF (Birds of a Feather) titled "HPC Certification Program" in the field of education and training for the ISC 2018 together with Weronika Filinger (EPCC) and Jean-Thomas Acquaviva (DDN).

References

[1]FEPA. *Flexible Framework for Energy and Performance Analysis in HPC Centers – Workshop 2017*. https://blogs.fau.de/prope/fepa-workshop-2017/ Slides: https://www.hhcc.uni-hamburg.de/en/files/fepa2017-pecoh-slides.pdf

[2]Gauß. *7^{th} HPC-Status-Conference of the Gauß Allianz – 2017*. https://gauss-allianz.de/en/hpc-status-konferenz-2017?name=value Slides: https://www.hhcc.uni-hamburg.de/en/files/gauss2017-pecoh-slides.pdf

[3]HHCC. *Hamburg HPC Competence Center – Handout to the work in progress of the HPC Certification Program*. https://www.hhcc.uni-hamburg.de/en/files/isc2017-hpc-certification-program.pdf

[4]HHCC. *Hamburg HPC Competence Center – HPC Certification Program: Concept Paper - Draft Version 0.91 – June 1, 2018*. https://www.hhcc.uni-hamburg.de/en/files/hpccp-concept-paper-180601.pdf

[5]HHCC. *Hamburg HPC Competence Center – Mailing List of the HPC Certification Program*. certification.hhcc@lists.uni-hamburg.de

[6]Kunkel, Julian, Michael Kuhn, Thomas Ludwig, Matthias Riebisch, Stephan Olbrich, Hinnerk Stüben, Kai Himstedt, Hendryk Bockelmann, and Markus Stammberger. Performance Conscious HPC (PeCoH) – Project Poster. *ISC High Performance 2017 (20 June 2017)*. Frankfurt, Germany.

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.