

# An HPC Certification Program Proposal

## Meeting HPC Users' Varied Backgrounds

Kai Himstedt<sup>1</sup>, Nathanael Hübbe<sup>1</sup>, Julian Kunkel<sup>2</sup>, and Hinnerk Stüben<sup>1</sup>

<sup>1</sup> Universität Hamburg

<sup>2</sup> Deutsches Klimarechenzentrum

<sup>3</sup> Technische Universität Hamburg

## Acknowledgements

The authors acknowledge the discussion with Hendryk Bockelmann<sup>2</sup>, Michael Kuhn<sup>1</sup>, Thomas Ludwig<sup>1,2</sup>, Stephan Olbrich<sup>1</sup>, Matthias Riebisch<sup>1</sup>, Sandra Schröder<sup>1</sup>, and Markus Stammberger<sup>3</sup>

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.

## 1 Introduction

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. A special challenge is to provide users with skills according to their scientific backgrounds and specific demands in terms of the usage of the system. Users in the role of testers, who want to simply run a parallel program for benchmark purposes, must e.g. have a solid knowledge of operating system basics and should be able to use a workload manager like SLURM [SLUR 17] or TORQUE [TORQ 17], but in general they do not need a deeper understanding of the technical refinements of the parallelization of the program. A user who wants to develop a parallel program will usually already be able to use the operating system and the workload manager but will need further skills to apply parallelization techniques like OpenMP [OpMP 17] or GPU-computing based on CUDA [NVID 17] at the intra-node level, MPI [MPI 17] at the inter-node level or even combinations of such techniques in the sense of a hybrid or multi-level approach.

In recent years, the growing demands to improve the HPC education are in the research focus of many projects. Rüde, for example, in the final report on exascale education in the context of the "European Exascale Software Initiative 2" [Cordi 18], describes strengthening of HPC education as an important subfield of computational science and engineering (CSE) [Rüde 15]. The urgent demand for an appropriate HPC education is also indicated by the efforts of the HPC advisory council [HPCA 17] to push it, e.g. by offering workshops and releasing best practice as well as case studies. Therefore, it comes as no surprise that recently initial results of a Scientific Computing World HPC readership survey have shown that "... training and support for HPC resources are the number one concern for both those that operate and manage HPC facilities and researchers using HPC resources." [SCW 17].

Establishing an HPC Certification Program is a central part<sup>1</sup> of the joint Performance Conscious HPC (PeCoH) project. In April 2017, the three Hamburg compute centers involved in PeCoH, German Climate Computing Center (DKRZ), Regional Computing Center at the Universität Hamburg (RRZ), and Computer Center at the Technische Universität Hamburg (TUHH RZ) started the Hamburg HPC Competence Center (HHCC) as a virtual institution and central contact point for their users [HHCC 17a]. HHCC will also serve as an open-for-all education platform for HPC knowledge and competences. 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. The performance engineering aspect is of particular importance because, according to our experience at DKRZ, RRZ, and TUHH RZ, support requests are currently dominated by problems at the level of getting things to work, i.e. getting a parallel job to run. Users in this situation are far from being aware of using the expensive HPC resources appropriately.

The paper is organized as follows: In Section 2 the classical approaches for HPC education are sketched. In Section 3 our innovative approach of an HPC Certification Program will be presented, which is based on defining HPC skills a user must have to perform certain tasks like testing, building or developing parallel programs in an HPC environment. The project's progression as well as the data structures and technical modeling to define the hierarchical dependencies of the skills are also handled in this section. Finally, the major insights are concluded in Section 4, and some future work is pointed out in Section 5. The Appendix contains a detailed list of all skills we have identified for the HPC Certification Program so far.

## 2 Classical HPC Education

A good user education has traditionally been important, because it leads to cost reductions in the operation at compute centers by reducing efforts for user support and by more efficiently used HPC resources by well-trained users. From our observations at DKRZ, RRZ, and TUHH RZ, new users without a proper HPC education often use only the defaults of the respective workload manager for selecting HPC resources such as main memory or CPU time and often do not explicitly select an appropriate batch queue to submit their jobs, while a user with an adequate HPC education will take meaningful estimates into consideration to avoid reserving unnecessary amounts of HPC resources, long waiting times for the job start, or reaching the runtime limit of his job. At the same time the productivity of the users will be increased, because they feel comfortable in using an HPC system. This leads to a typical win-win situation.

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. For some time now, there have also been joint efforts to support the HPC education in Europe. The education and training strategy at the Barcelona Supercomputing Centre (BSC) as outlined in [Sanc 15] may serve as an example: as part of the Partnership for Advanced Computing in Europe (PRACE), BSC is working on the development of an appropriate European HPC professional training curricula. Classical HPC education is based on lectures, tutorials, and workshops addressing

---

<sup>1</sup> Further goals of PeCoH are e.g. to develop models to estimate the costs of batch jobs in order to give HPC users feedback indicating the impacts of running non-optimized workloads, and to develop analysis tools to (automatically) identify performance issues caused by well-known configuration mistakes in job scripts.

the various HPC topics. An HPC lecture usually involves a teacher presenting topics and concepts related to a course addressing HPC topics to users enrolled in that course and has a rather static character. An HPC tutorial is typically run in smaller groups and allows discussion of the content and interaction with other users. However, the general procedure stays rather static, which also applies to HPC workshops where users typically acquire HPC skills by involving more hands-on learning activities.

Nowadays, it is very simple to publish a live or recorded lecture on an online platform, which gives users the possibility to watch the video where and when they like. Tutorials commonly make the content (additionally) available online via the internet (see [LLNL 17] for an example). The interactive aspect of a classical tutorial may suffer, but that can be more than compensated by the improved accessibility of the hyperlinked content. In addition there are hybrid approaches. Zarestky and Bangerth [ZaBa 14], for example, performed an experiment to teach HPC with a so-called flipped classroom format that requires students to watch content videos before coming to class, thus freeing time in class. Based on qualitative data Zarestky and Bangerth report positive results in terms of being able to use the time in class efficiently, and instructors and students enjoyed the new format. Reflecting the workshop idea, there exists online content with a focus on practical HPC examples showing how to get things to work (e.g. [CAC 17]). The Extreme Science and Engineering Discovery Environment (XSEDE), a virtual organization to support open research, helps their users among other things by an online system to train the usage of an HPC system, structuring the corresponding information on their website by the help of major topics like “Getting Started”, “Working with the System”, “Visualization Resources”, and “HPC System Resources”. The user can select additional information about each topic to navigate within the content [XSED 17]. There are also Websites offering (online) HPC learning material (e.g. [FuLe 18], [PRACE 18]). However, sophisticated Web-based E-learning systems which cover the users’ varied backgrounds and their individual learning progresses do not exist – to the best of our knowledge – for teaching HPC competences.

In addition to the benefits of using a modern Web-based approach to present the HPC content in a more dynamic and, if needed, multimedia based way, there are ideas to use computing resources more generally for additional HPC education purposes. Holmes and Kureshi [HoKu 15], for example, reported – against the background of a shortage of HPC skills and available HPC training in the UK – experiences using recycled laboratory PCs to build cluster systems for educational purposes. Not only can the students use the clusters for experiments, but the challenge to build these laboratory clusters had a positive impact in that it encourages them to search for information from a variety of sources in order to complete the building tasks, and that developed their skills and confidence in the process. Czarnel [Czar 14] presents a successful middleware approach including a Web-based interface to support easy access to HPC systems for HPC novices by hiding the queuing systems. Suh et al. [Suh<sup>+</sup> 16] adopt an approach which rather focuses on encapsulating simulation systems behind a user-friendly graphical user interface (GUI) supporting scientific workflows. This system is also made to support the education of students, but rather in the field of computational science and engineering (CSE) than in the field of HPC competences.

Summing up, online platforms for HPC education are successfully used in practice and provide great potential. However, in contrast to other areas of information technology (IT), where certificates are often used to prove IT skills<sup>2</sup> of the users, in the field of HPC neither

---

<sup>2</sup> There is a certification program for various levels of Linux system administration skills from the Linux Professional Institute [LPI 17] and a certification program for general (personal) computer skills from the European Computer Driving Licence [ECDL 17a] organization, which could serve as representative examples here.

commonly accepted standards exist, nor a certification program for the education. If a scientific institution provides learning material it will be determined by the special demands of the respective institution and its specific HPC environment. Therefore, this content will only cover a very small part even of basic HPC skills and a user with a lack of basic skills will presumably have difficulties to readily use other HPC systems. These are the issues addressed by our proposal for an HPC Certification Program.

### 3 New Approach

Living in the age of so-called digital natives, one might suppose that computer skills are picked up intuitively. The ECDL organization notes, however, that this is not the case for basic computer skills and the idea of digital natives is a dangerous fallacy that risks leaving young people without the competences they need for the workplace, and risks leaving businesses without the skilled employees they need [ECDL 17b]. It can be assumed that this is all the more true for the complex field of HPC.

The ambitious EuroLab-4-HPC project, funded in the context of the Horizon 2020 research and innovation programme, focuses on developing a structured HPC systems curriculum and training practices based on (online) courses [EURO 18a]. As a project result it is shown how the courses can also be mapped to other degree programs (e.g. physics, mathematics) at the master's level or how they can be used for a single year's program that is Bologna-aligned ([EURO 18b] p. 12). Certificates are clearly of less significance compared for example to a master's or Phd degree, but on the other hand university degrees, with their rather great scope and possibly more national character, do not attest knowledge or skills of specific and topical technologies. This training gap can be filled ideally by the help of certification programs.

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 use an HPC environment for their research. Another analogy is the transferability of skills: Anyone who is able to drive a certain type of passenger car is able to drive any other passenger car, and an HPC user who has gained the skill to use a workload manager like TORQUE will be able to use SLURM after short period of additional training, and vice versa.

Before the new approach of the HPC Certification Program is presented in more detail, we will introduce a set of terms:

*Skill*: The abilities and the knowledge specified in the skill description

*Certified Skill*: Skill of a user validated by an exam

*Content*: Learning material enabling the user to gain certified skills

*Content provider*: Institution that provides content

*Exam*: Process to validate a user's skill based on multiple-choice tests

*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 based on certified skills according to the corresponding certificate definition

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. This is comparable to the concept of a central high school graduation exam (Zentralabitur in German), where the examination is created by a central organization while the pupils are prepared for the exam by their schools. Since the start of the PeCoH project, when we had the role as certification provider as well as the role as content provider for basic HPC skills<sup>3</sup>, we welcomed the collaboration with other scientific institutions to establish generally recognized certificate definitions. Essentially, it is at the discretion of a content provider to decide which learning material is most appropriate to teach a skill. That offers freedom and flexibility in creating the learning content. We assume that collaborating scientific institutions will complement each other in producing content.

### 3.1 Previous Work

We started our development of the HPC Certification Program with the classification of HPC topics which were relevant to the three compute centers (DKRZ, RRZ, and TUHH RZ) involved. We initially identified four top level competences: “HPC Knowledge”, “Use of the HPC Environment”, “Performance Engineering”, and “Software Development” as shown in Figure 1.



**Fig. 1: Top Level Competences**

We presented a poster of the current state and goals in the PeCoH project at the ISC 2017 [Kunk<sup>+</sup> 17] and distributed a handout containing the initial classification of HPC competences and the work in progress of our HPC Certification Program [HHCC 17b], which was one of the major topics of the poster. We also presented the idea of the HPC Certification Program at the Flexible Framework for Energy and Performance Analysis in HPC Centers (FEPA) workshop [FEPA 17]. At both events we received positive feedback in several meetings and discussions, which underlines again the urgent demand for an appropriate HPC education at other compute and data centers. Additionally, we are hosting a mailing list for the HPC Certification Program [HHCC 17d].

<sup>3</sup> Within the PeCoH project we will establish all significant certification definitions. To produce content for all HPC skills listed in the Appendix we depend on the collaboration of others.

## 3.2 HPC Skill Tree

It is in the nature of the subject that HPC skills are generally built upon one another, which results in a tree structure for representing skills depending on sub-skills. The tree of HPC skills is a key component of our approach and has a role of a database for the HPC Certification Program. First of all, skills have unique names and contain a description of the HPC competences and knowledge that are associated to them. Furthermore, each skill is assigned to one of the four top level competences as described in the previous section and has additional attributes to describe its properties in more detail, like its special significance to a scientific domain (e.g. social sciences, natural sciences, earth sciences), the suitability for a user's role (e.g. tester, developer), or its educational level (e.g. basic, intermediate, or expert). This information allows to easily create different views of the skill tree in order to consider the users' varied backgrounds, e.g. for navigating within the skill tree by the help of a Web-based GUI using the attributes to filter the relevant information for them.

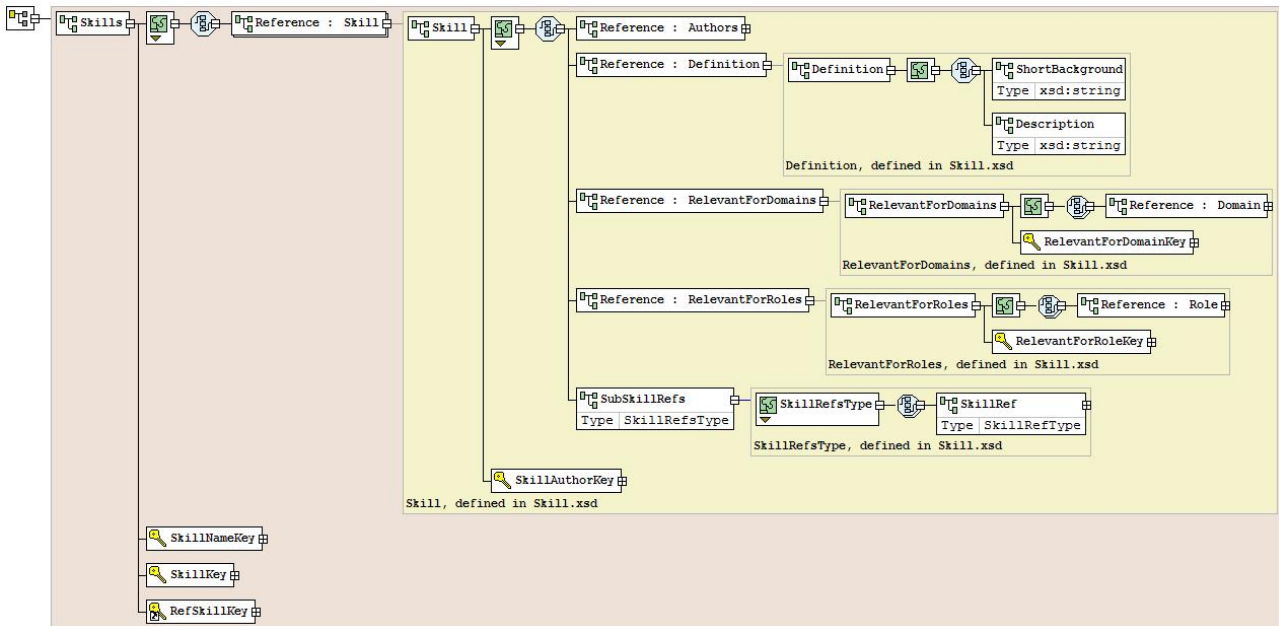
The implementation of the skill tree is based on the Extensible Markup Language (XML) [W3 17a] and a corresponding XML Schema Definition (XSD) [W3 17b]. XML is an open de-facto standard to process and exchange information in heterogeneous environments. XML data is human- as well as machine-readable, which supports the shared working on the skill tree implementation: XML files can e.g. be opened and inspected by project participants with their favoured (simple) text editor. With the machine-readable property of XML it is possible to check the syntax of an XML file having been changed, with respect to the so-called well-formedness, and validate it with the corresponding XML schema definition. A further potential is the ability to process the data with sophisticated tools, e.g. parsers, in a variety of ways. Another reason we decided to implement the skill tree on the basis of XML is the variety of powerful tools and integrated development environments (IDEs) available to support such development (e.g. MissionKit [Alto 17], Stylus Studio X16 [StSt 17], or Eclipse XML Editors and Tools [Ecli 17]). Since the skill tree is of manageable size, there is no need to use a more complex database design for its representation. JSON [JSON 17] is another popular human- and machine-readable data-interchange format, which is rated a little bit more lightweight than XML, and was also worth considering to be used to implement the skill tree. While JSON focuses on the temporary exchange of data, the XML world provides a rich family of languages, which seems to offer more potential for the modelling process. If necessary, however, XML data can easily be converted to other formats like JSON (and vice versa), in particular by the help of XSLT [W3 17c].<sup>4</sup>

The essential data structure of the skill tree is presented in Figure 2 based on the relevant part of its XML Schema Definition.

As is typical for popular naming conventions of XML data structures, the *Skills* definition in the Figure shows that the XML data of the skill tree contains, first of all, a list of *Skill* items, i.e. the XML data contains all the nodes of the skill tree in a flat data structure. In order to describe the tree, each *Skill* that depends on other sub-skills has – besides its unique name, description, and further attributes – a list of references to these sub-skills. For example, in our design, the skill to build a parallel program, e.g. via an open source package, will at least require the skill to run a parallel program in an HPC environment and that in turn will require skills to use the command line interface of the operating system and a workload manager like SLURM or TORQUE. Unique skill names are used for this referencing to other skills in the *Skills* list.

---

<sup>4</sup> At the implementation level, we plan to use JavaScript [JaSc 17], which has a native support for JSON, to make the skill tree browsable in a Web-based GUI.



**Fig. 2:** XML Schema Definition for Showing the Essential Skill Tree Structure

Obviously, this data structure allows the definition of different skills depending on the same sub-skill, so, strictly speaking, the skill tree becomes a directed acyclic graph (DAG).<sup>5</sup> This is similar to using a `Makefile` for the well-known `make` build automation tool [Feld 79] to define the dependencies of compilation units: in C, for example, header files often contain declarations that are used (i.e. included) in different source files and other header files. The `Makefile` allows to rebuild libraries and the target program in the correct order after source code changes by the help of a depth-first search to resolve all transitive dependencies between the compilation units. Similarly, a user could acquire relevant skills at the leaf level of the skill tree first and then proceed to acquire skills nearer to the root. To be able to show all skills in a tree format, e.g. in a Web-based GUI, multiple references to the same skill could be resolved by presenting the more than once referenced skill several times, so, for the sake of simplicity, the DAG property shall be neglected here.

In contrast to a `Makefile` defining a single type of relationship between dependencies<sup>6</sup>, two types of relationships are supported by the skill tree structure to define dependencies for skills: In the standard case, all skills in the list of referenced sub-skills are combined implicitly by a logical *and* operation, i.e. a skill can only be gained if *all* of its sub-skills have been gained. The second relationship is based on logical *or* operations, and allows users to gain a skill when at least *one* of the referenced sub-skills has been gained. For example, the skill to use a workload manager can be awarded to users who are able to use one of the workload managers SLURM *or* TORQUE. In practice, it will follow from context when which list type should be used in the skill definition to reference sub-skills. Within the same skill, however, it is not possible to use both types of reference lists at the same time. It would be possible to support this directly by using the composite pattern [Gamm<sup>+</sup> 95]: The basic list of implicitly *and*-combined referenced sub-skills could additionally contain lists of *or*-combined referenced sub-skills. This could be expressed in the XML Schema Definition without greater effort. But since it is easily possible to create an additional skill containing

<sup>5</sup> In the Appendix containing the detailed list of all skills, for the sake of simplicity as plain text format, such cross references to other skills begin with "see also ...".

<sup>6</sup> The time stamp of a file can be out of date in relation to the time stamps of the files it depends on. This way `make` can for example rebuild an object file if it is out of date in relation to more recently changed source files it depends on.

the list of *or*-combined referenced sub-skills and referencing to this additionally created skill in the list of *and*-combined referenced sub-skills, we preferred to keep the data structures as simple as possible.

The attributes of the skills allow to present the skill tree in a highly dynamic manner. This way users can first of all get an overview of those custom-tailored skills which they need for the HPC environment they would like to use or the parallel program they would like to speed up. However, the skill tree itself is content-free and solely describes which HPC competences have to be taught and learned. This reflects the separation of the certificate definition, which is based in our approach on skills, from the learning material that allows the user to gain the associated skills. In the sense of an E-learning environment it is possible to present a specific content in a Web-based system, which in turn maps it to the skill tree. In further stages of the project, the skill tree can be extended to support links to learning material. In this way, a single Web-based system can be used for browsing the skill tree as well as the content.

A special challenge is to determine a reasonable granularity of skills as defined by their descriptions. One can easily imagine that an increasingly finer granularity results if one attempts to dissolve the leaves of the skill tree more and more, with a skill at the leaf level finally predefining its content. At the beginning, we actually dissolved basic skills how to use the Linux command line interface to verify the practicability of our XML implementation of the skill tree. This was indeed possible without any problems, but from the fine granularity almost a one-to-one relationship results between the skill description and the related content, so that simply put each skill would have been imparted to the extent of a single presentation slide. A representative skill definition from this ad hoc example can illustrate this: the skill "Navigate the file system" was dependent on the sub-skills "Understand the file system tree", "Print name of current working directory", "Change directory", and "List directory contents". For the entire ad hoc example the definition of 59 skills was required just for describing some frequently used Linux commands (`cd`, `ls`, `less`, `cat`, `cp`, `mv`, `mkdir`, `rm`, `help`, `info`, `chmod`, `chown`, and `chgrp`).

It is obvious that a very fine granularity not only restricts the freedom in providing the content, but also makes it more difficult to define certificates because the number of skills will strongly increase accordingly. A granularity that is too coarse, such as a limitation to the top level competences shown in Figure 1, is also not useful as it would give the content providers essentially no assistance in structuring the learning material. The Appendix contains the list of all 46 skills we have identified for the HPC Certification Program so far, 35 of which are at the leaf level. We think with this granularity we found a good compromise between both extremes in order to separate skills and content. We will use up to three levels of education (basic, intermediate, and expert) to further subdivide each skill and to define the HPC competence level a user has acquired with regard to a skill. For the sake of simplicity the educational levels are not shown in the Appendix. The process of subdividing requires experience and expert knowledge. The information about skill attributes and educational levels is contained in the XML description of the skill tree, which will soon be available on our HHCC website [HHCC 17c].

### 3.3 Certification Modeling

The basic idea of defining certificates is to bundle a set of skills corresponding to the certificate description in order to certify – by successful exams – a user's HPC qualification. Like skills, certificate definitions have unique names and contain a description of the HPC com-



petences and knowledge that are associated to them. While a skill is a more self-contained unit, a certificate definition describes on a conceptual level a further view on the skills. The skill tree represents a middle-layer between the certificate definitions and the actual content.

Initially, we intended to implement the certificate definitions by a separate data structure, which was based, like the definition of the skill tree, on XML and a corresponding XML Schema Definition. But since both XML structures were so similar that their distinction served rather a conceptual purpose than a technical one, it was natural to extend the skill tree – functioning as a central database – to be able to incorporate the properties additionally required for the definition of certificates as well, instead of using a separate structure for defining certificates. At first a skill can be easily tagged for its additional suitability as an autonomous certificate definition. A user who has gained the skill to run parallel programs in an HPC environment may thus be granted a corresponding certificate. For the skill tree it was described that two types of relationships are supported (based on *and* and *or* operations) to define dependencies on sub-skills, so that a skill is gained if all of its sub-skills are gained (*and* operation) or if one of its sub-skills is gained (*or* operation). For the definition of certificates these two types of operations were supplemented by an *n out of m* relationship, so a skill is considered to be gained if at least *n*, or a corresponding percentage value, of its sub-skills have been gained. With this type of relationship, users can be certified, for example, for their experienced ability to use version control systems, if they gain sub-skills to use two systems from the set consisting of the Revision Control System (RCS) [Tich 85], Subversion (SVN) [TASF 17], and Git [Git 17].

## 4 Conclusions

There is an urgent demand to improve the users' HPC education to enable them to use the HPC resources appropriately. This will increase their productivity and at the same time reduce the costs in the operation at compute centers. While certificates are widely applied in the IT industry to testify that users have certain skills, e.g. to administer Linux systems, this is not the case for the field of HPC. With the proposal of an HPC Certification Program we try to establish a standard for the education of HPC users. In our approach we separated the certificate definitions from the providing of the learning material. 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. The content can be provided in a variety of ways by various collaborating providers.

Sophisticated Web-based E-learning systems which cover the users' varied backgrounds (concerning for example research area and prior knowledge) do not exist for the HPC education. For our approach we implemented an HPC skill tree based on XML and a corresponding XML Schema Definition (XSD), which plays the role of a central database (see also Appendix). 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. Not only can well-trained certified users with a good knowledge of performance engineering concepts speed up their parallel programs to get their scientific results faster, but also can compute centers reduce their costs because the HPC resources will be used more efficiently.

One major challenge was to find a good compromise for the scope of the skill descriptions, in particular at the leaf level, so that a too fine granularity will not predefine the content of a skill or an all too coarse granularity will be of no help at all for the content providers for structuring the learning material. The skill tree for our HPC Certification Program contains 46 skills (see also Appendix), which we consider to be a suitable granularity.

## 5 Status, Collaboration, and Future Work

The development of the HPC skill tree is nearly completed. At DKRZ, RRZ, and TUHH RZ we already have some content to teach HPC topics, which can be used to fill the content-free structure level formed by the HPC skill tree for our online education platform. We welcome suggestions from interested readers on the tree structure and the actual classification of HPC skills. Furthermore, we encourage readers to provide us with content for HPC skills and will express our gratitude by a corresponding entry in the acknowledgement area on our website. (See also Appendix for the list of skills.)

A user will have to participate in online examinations based on multiple-choice tests to gain an HPC certificate. For each HPC skill a pool of questions is developed, of which a subset is selected for each individual examination. Once the test is completed, the system will automatically assess the results and create a PDF with the certificate. At the beginning, we will manually approve the test results. Later on in the development, the individual learning progress could be stored as a part of the user account, allowing users to interrupt their exam preparation at any time and continue later to navigate seamlessly in the learning content.

It will be particularly interesting to measure the success of the certificate-based approach for the HPC education. One idea is to see if there will be less support requests of new users with simple demands for running parallel programs on the clusters at DKRZ, RRZ, and TUHH RZ. With additional surveys, the users' satisfaction with the certification program can be determined. It will also be possible to check if the performance awareness of certified users is raised, i.e. if they use the HPC resources more appropriately.

## 6 References

- [Alto 17] Altova. *Altova MissionKit – Award-winning Suite of XML, SQL, & UML Tools*. <https://www.altova.com/missionkit>
- [CAC 17] CAC. *Cornell University Center for Advanced Computing – Cornell Virtual Workshop*. <https://cvw.cac.cornell.edu/default>
- [CORDI 18] CORDIS. *Community Research and Development Information Service: European Exascale Software Initiative 2 – Towards exascale roadmap implementation*. <http://cordis.europa.eu/project/rcn/105840.en.html>
- [Czar 14] Czarnul, Pawel. Teaching High Performance Computing Using Beesy-Cluster and Relevant Usage Statistics. *International Conference On Computational Science (ICCS 2014), Procedia Computer Science*. Vol. 29 (2015):1458-1467.
- [Ecli 17] Eclipse. *Eclipse XML Editors and Tools*. <https://marketplace.eclipse.org/content/eclipse-xml-editors-and-tools-0>
- [ECDL 17a] ECDL. *European Computer Driving License – Home Page*. <http://ecdل.org/>
- [ECDL 17b] ECDL. *European Computer Driving License – The Fallacy of the Digital Native*. <http://ecdل.org/policy-publications/digital-native-fallacy>
- [EURO 18a] EUROLAB-4-HPC. *EuroLab-4-HPC – Home Page*. <https://www.eurolab4hpc.eu/>
- [EURO 18b] EUROLAB-4-HPC. *D3.2 Best Practices in HPC Training*. <https://www.eurolab4hpc.eu/static/deliverables/D3-2-best-practices-HPC-training.610d055cf370.pdf>
- [Feld 79] Feldman, Stuart I. Make – A Program for Maintaining Computer Programs. *Software Practice & Experience*. Vol. 9 (1979):255-265.
- [FEPA 17] FEPA. *Flexible Framework for Energy and Performance Analysis in HPC Centers – Workshop 2017*. <https://blogs.fau.de/prope/fepa-workshop-2017/>
- [FuLe 18] FutureLearn. *Online Course Supercomputing*. <https://www.futurelearn.com/courses/supercomputing#section-topics>
- [Gamm<sup>+</sup> 95] Gamma, Erich, Richard Helm, Ralph Johnson, John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley. Boston, San Francisco, New York 1995.
- [Git 17] Git. *git –fast-version-control*. <https://git-scm.com/>
- [HPCA 17] HPCA. *HPC Advisory Council – Home Page*. <http://www.hpcadvisorycouncil.com/>
- [HHCC 17a] HHCC. *Hamburg HPC Competence Center – Home Page*. <https://www.hhcc.uni-hamburg.de>
- [HHCC 17b] 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>

- [HHCC 17c] HHCC. *Hamburg HPC Competence Center – Download Area*. <https://www.hhcc.uni-hamburg.de/en/support/downloads.html>
- [HHCC 17d] HHCC. *Hamburg HPC Competence Center – Mailing List of the HPC Certification Program*. [certification.hhcc@lists.uni-hamburg.de](mailto:certification.hhcc@lists.uni-hamburg.de)
- [HoKu 15] Holmes, Violeta, and Ibad Kureshi. Developing High Performance Computing Resources for Teaching Cluster and Grid Computing courses. *International Conference On Computational Science (ICCS 2015), Procedia Computer Science*. Vol. 51 (2015):1714-1723.
- [JaSc 17] JavaScript. *JavaScript – Reference*. <https://developer.mozilla.org/en-US/docs/Web/JavaScript>
- [JSON 17] JSON. *JavaScript Object Notation – Introducing JSON*. <http://www.json.org/>
- [Kunk<sup>+</sup> 17] 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. Download via [http://isc-hpc.com/isc17\\_ap/presentationdetails.htm?t=presentation&o=1196&a=select&ra=personendetails](http://isc-hpc.com/isc17_ap/presentationdetails.htm?t=presentation&o=1196&a=select&ra=personendetails)
- [LLNL 17] Lawrence Livermore National Laboratory. *Livermore Computing Center – High Performance Computing: Tutorials*. <https://hpc.llnl.gov/training/tutorials>
- [LPI 17] LPI. *Linux Professional Institute – Home Page*. <http://www.lpi.org/>
- [MPI 17] MPI. *The Message Passing Interface (MPI) standard*. [www.mcs.anl.gov/research/projects/mpi/](http://www.mcs.anl.gov/research/projects/mpi/)
- [NVID 17] NVIDIA. *CUDA Zone*. <https://developer.nvidia.com/cuda-zone>
- [OpMP 17] OpenMP. *The OpenMP API Specification for Parallel Programming*. [www.openmp.org](http://www.openmp.org)
- [PRACE 18] Partnership for Advanced Computing in Europe. *Training Portal – Training Courses*. [http://www.training.prace-ri.eu/nc/training\\_courses/index.html](http://www.training.prace-ri.eu/nc/training_courses/index.html)
- [Rüde 15] Rüde, Ulrich. *European Exascale Software Initiative 2: Deliverable D2.3 WP2 Final Report on Exascale Education*. [www.eesi-project.eu/wp-content/uploads/2015/05/EESI2\\_D2.3\\_Final-report-on-exascale-education.pdf](http://www.eesi-project.eu/wp-content/uploads/2015/05/EESI2_D2.3_Final-report-on-exascale-education.pdf)
- [Sanc 15] Sancho, Maria-Ribera. BSC Best Practices in Professional Training and Teaching for the HPC Ecosystem. *Journal of Computational Science*. Vol. 14 (2015):74-77.
- [SLUR 17] SLURM. *SLURM Workload Manager Overview*. <https://slurm.schedmd.com/overview.html>
- [SCW 17] Scientific Computing World. *Training and Support Number One Concern for the HPC Community*. <https://www.scientific-computing.com/news/training-and-support-number-one-concern-hpc-community>

- [Suh<sup>+</sup> 16] Suh, Young-Kyoon, Hoon Ryu, Hangi Kim, and Kum Won Cho. EDISON: A Web-based HPC Simulation Execution Framework for Large-scale Scientific Computing Software. *16th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid)*. IEEE Conference Publications. (2016):608-612.
- [StSt 17] Stylus Studio. *X16 – Powerful XML Development*. <https://www.stylusstudio.com/index.html>
- [TASF 17] The Apache Software Foundation. *Apache Subversion – Enterprise-class centralized version control for the masses*. <https://subversion.apache.org/>
- [Tich 85] Tichy, Walter F. RCS – A System for Version Control. *Software Practice & Experience* 15 (1985):637-654.
- [TORQ 17] TORQUE. *Torque Resource Manager*. [www.adaptivecomputing.com/products/open-source/torque/](http://www.adaptivecomputing.com/products/open-source/torque/)
- [W3 17a] W3. *World Wide Web Consortium – Extensible Markup Language (XML)*. <https://www.w3.org/XML/>
- [W3 17b] W3. *World Wide Web Consortium – XML Schema*. <https://www.w3.org/2001/XMLSchema>
- [W3 17c] W3. *World Wide Web Consortium – XSL Transformations (XSLT)*. <https://www.w3.org/TR/xslt/>
- [XSED 17] XSEDE. *Extreme Science and Engineering Discovery Environment – Training*. <https://portal.xsede.org/web/xup/training/overview>
- [ZaBa 14] Zarestky, Jill and Wolfgang Bangerth. Teaching High Performance Computing: Lessons from a Flipped Classroom, Project-Based Course on Finite Element Methods. *Workshop on Education for High Performance Computing (EduHPC) held in conjunction with SC14: The International Conference for High Performance Computing, Networking, Storage and Analysis*. New Orleans, Louisiana, November 16-21. IEEE Conference Publications (2014):34-41.

## 7 Appendix

At the end of the Appendix, the HPC skill tree is presented as a compact diagram.

In the following, all skills we have identified for the HPC Certification Program so far are listed in a hierarchical manner to reflect their underlying tree structure as described in Section 3. The hierarchy here is based on four top level competences: “HPC Knowledge”, “Use of the HPC Environment”, “Performance Engineering”, and “Software Development”. The *Description* section of a skill specifies the abilities and the knowledge a user will gain. When appropriate, some additional information may be presented in the *Short Background* section.

For the sake of simplicity, the attributes of the skills to indicate a special significance for users in dependence of their varied backgrounds (e.g. social scientist, natural scientist, earth scientist) or roles (e.g. tester, developer) are not included here. The same holds for the finer differentiation of a skill description regarding its educational level (e.g. basic, intermediate,

or expert). This type of additional information is contained in the full XML description of the skill tree, which will soon be available on our HHCC website [HHCC 17c].

## K HPC Knowledge

*Description:*

Knowledge of the field of High Performance Computing

### K1 Supercomputers

*Description:*

Knowledge of various system-, hardware-, and I/O-architectures used for supercomputers, i.e. computers that led the world in terms of processing capacity, and particularly in speed of calculations, at the time of their introduction, or share key architectural aspects with these computers

Knowledge of typical operation of data and computing centers

Knowledge of the differentiation between Supercomputing and Big Data

#### K1.1 System Architectures

*Description:*

Knowledge of various system-, hardware-, and I/O-architectures used for supercomputers, i.e. shared memory systems, distributed systems, and cluster systems

Knowledge of the typical architecture of cluster systems consisting of nodes with different roles (e.g. so-called head, login, compute, interactive, visualization nodes, etc.)

Knowledge of storage and compute deployments for cluster systems

#### K1.2 Hardware Architectures

*Description:*

Knowledge of elementary processing elements like CPUs, GPUs, many core architectures, and other special or application-specific hardware (e.g. TPUs)

Knowledge of parallelization techniques at the instruction level of a processing element (e.g. pipelining, SIMD processing)

Knowledge of vector systems, and FPGAs

Knowledge of hybrid approaches, e.g. combining CPUs with GPUs or FPGAs

Knowledge of the NUMA architecture used for symmetric multiprocessing systems where the memory access time depends on the memory location relative to the processor

Knowledge of network demands for HPC systems (e.g. high bandwidth and low latency)

Knowledge of typical network topologies and architectures used for HPC systems, like fat trees based on switched fabrics using e.g. fast Ethernet (1 or 10 Gbit) or InfiniBand

#### K1.3 I/O Architectures

*Description:*

Knowledge of typical I/O systems used in HPC environments

Knowledge of different types of storage media (e.g. tape, disk, and SSD)

Knowledge of the differentiation between standard file systems (e.g. Ext3, Ext4, XFS, Btrfs) and distributed file systems (e.g. Lustre, BeeGFS)

Knowledge of when to use local and global storage

Knowledge of when to use data compression

#### K1.4 Operation of an HPC System

*Description:*

Knowledge of the typical infrastructure of data and computing centers, also against the background of economic, business, and organizational aspects  
Knowledge of administration aspects of an HPC system  
Knowledge of user support aspects (typically on different levels)

#### K1.5 Supercomputing and Big Data

*Short Background:* In the recent past, Supercomputing as well as the analysis of Big Data are increasingly growing in importance for scientific research.

*Description:*

Knowledge of the differentiation between Supercomputing and Big Data

#### K2 Performance Modeling

*Short Background:* HPC systems are massively parallel and therefore sophisticated parallel programs are required to exploit their performance potential as much as possible.

*Description:*

Knowledge of how the performance of parallel programs may be assessed

##### K2.1 Performance Frontiers

*Description:*

Knowledge of the definitions for key terms like speedup, efficiency, and scalability

Knowledge of the key measure floating point operations per second (FLOPS) for the performance of HPC systems and its pitfalls

Knowledge of Moore's and Amdahl's laws and their significance for performance frontiers in modern HPC

Knowledge of the roofline model, used to provide performance estimates for parallel programs based on multi-core or accelerator processor architectures, by showing inherent hardware limitations

##### K2.2 Bounds for a Parallel Program

*Description:*

Knowledge of how performance bounds of the various components of the HPC system (e.g. CPU, caches, memory, network, I/O) can limit the overall performance of a parallel program

#### K3 Program Parallelization

*Description:*

Knowledge of the typical parallelization techniques used at the intra- and inter-node level of cluster systems

Knowledge of the causes of parallelization overheads, which eventually prevent efficient use of an increasing number of processing elements

Knowledge of domain decomposition strategies (i.e. splitting a problem into pieces that allow for parallel computation)

##### K3.1 Levels of Parallelization

*Description:*

Knowledge of the auto parallelization capabilities of current compilers (e.g. to automatically parallelize suitable loops), which are applicable at the intra-node level

Knowledge of parallelization techniques at the intra-node level (e.g. based on advanced OpenMP features and GPU-computing)

Knowledge of hybrid parallelization approaches, combining for instance OpenMP and GPU-Computing

Knowledge of the message passing paradigm based on environments like MPI, which is the de-facto standard at the inter-node level for parallelizing programs using more than a single node  
Knowledge of multi-level hybrid approaches (e.g. combining OpenMP and MPI)

### K3.2 Parallelization Overheads

*Description:*

Knowledge of the various overheads, i.e. overheads for communication, synchronization, and redundant computations  
Knowledge of the problems of load imbalances, execution speed noise (OS jitter, cache contention, thermal throttling, etc.), and typical trade-offs (e.g. reducing the synchronization overhead by increasing the communication overhead)

### K3.3 Domain Decomposition

*Description:*

Knowledge of typical decomposition strategies to split a domain into subdomains to make it suited for parallel processing  
Knowledge of measures like surface to volume ratio and how to map domains to machines

### K4 Job Scheduling

*Description:*

Knowledge of how workload managers control the unattended background execution of programs or jobs respectively by the help of job queues  
Knowledge of typical scheduling principles (e.g. first come first served, shortest job first, fair share, and backfilling) to achieve objectives like minimizing the averaged elapsed program runtimes, treating the users fair, and maximizing the utilization of the available HPC resources

### K5 Modeling Costs

*Short Background:* The user's awareness of the costs related to the operation of an HPC system is raised. For the resources of an HPC system, a distinction is made between costs for the computing elements of the supercomputer and costs for the storage system.

*Description:*

Knowledge of the impact of a cluster nodes type (e.g. CPU type, main memory expansion, or GPU extensions) and of the storage media type (SSD, disk, or e.g. tape for long term archiving (LTA) purposes) on its costs  
Knowledge of how to assess runtime costs for jobs  
Knowledge of how to assess the costs for the infrastructure of data and computing centers as well as their personnel costs  
Knowledge of economic and business aspects, e.g. break-even considerations, when personnel costs for tuning a parallel program and savings through speedups achieved are compared

### USE Use of the HPC Environment

*Description:*

Ability to use a cluster operating system as well as to run, build, and develop parallel programs

### USE1 Use of the Cluster Operating System

*Short Background:* HPC systems are usually accessed via a command line interface



(CLI). The user acquires skills to use a – generally Linux based – CLI to interact with the HPC system.

*Description:*

Ability to use and write shell scripts e.g. to automatically execute several commands in a row that otherwise would have to be entered manually one by one and to automate (possibly more complex) tasks

Ability to select the right environment setting to build programs with the proper compiler, linker, and libraries versions or to run programs

#### USE1.1 Use of the Command Line Interface

*Description:*

Ability to execute frequently used commands, e.g. to navigate the file system, copy, rename, and delete files, view the contents of files, and to get detailed help for the usage of a command with all its options

Ability to use regular expressions and wildcards to select or filter several items at once (e.g. files)

Ability to login remotely to cluster nodes using e.g. SSH with password or SSH key authentication

Ability to access local and remote files (e.g. via SSHFS) in remote sessions

Ability to check disk quotas commonly used to limit the amount of disk space available for the user

#### USE1.2 Using Shell Scripts

*Description:*

Ability to use and write shell scripts

Ability to use flow control, e.g. for conditional and/or repeated execution of statements in scripts

Ability to use shell functions to break large, complex tasks into a series of small, simple tasks

Ability to read keyboard input to add interactivity to scripts

Ability to write robust job scripts

Ability to use troubleshooting, e.g. to handle syntactic and logical errors in scripts

#### USE1.3 Selecting the Software Environment

*Short Background:* HPC systems have generally installed multiple versions of a number of key software tools and software environments. Package managers like Spack are sketched.

*Description:*

Ability to select the appropriate software versions for deployment to the session environment, e.g. via the so-called environment modules system

#### USE2 Running of Parallel Programs

*Description:*

Ability to run parallel programs in an HPC environment

Ability to use the command line interface is required (see also USE1.1 Use of the Command Line Interface)

Ability to select the appropriate software environment is required (see also USE1.3 Selecting the Software Environment)

Ability to use a workload manager like SLURM or TORQUE to allocate HPC resources (e.g. CPUs) and to submit a batch job

Ability to use a workload manager to allocate HPC resources for running a parallel program interactively

Ability to write robust job scripts, e.g. to simplify job submissions by the help of automated job chaining is required (see also USE1.2 Using Shell Scripts)

Ability to consider cost aspects is required (see also PE1 Cost Awareness)

Ability to measure system performance as a basis for benchmarking a parallel program is required (see also PE2 Measuring System Performance)

Ability to benchmark a parallel program is required (see also PE3 Benchmarking)

Ability to tune a parallel program from the outside via runtime options is required (see also PE4.1 Tuning Without Modifying the Source Code)

Ability to apply the workflow for tuning is required (see also PE5 Optimization Cycle)

### USE3 Building of Parallel Programs

*Description:*

Ability to build parallel programs, e.g. from open source packages

Ability to run parallel programs in an HPC environment is required (see also USE2 Running of Parallel Programs)

Ability to use a compiler and to assess the effects of optimization switches available for the relevant compilers (e.g. GNU, Intel, PGI, NAG)

Ability to use a linker and to assess the effects of linker specific options and environment variables (e.g. -L and LIBRARY\_PATH, LD\_LIBRARY\_PATH, -rpath and LD\_RUN\_PATH)

Ability to use efficient open source libraries (e.g. OpenBLAS, FFTW) and highly optimized vendor libraries (e.g. Intel-MKL, IBM-ESSL)

Ability to configure the relevant settings (e.g. by setting compiler and linker options), which determine how the application ought to be build with regard to the parallelization technique(s) used (e.g. OpenMP, CUDA, OpenACC, C++ AMP, MPI)

Ability to use the profile guided optimization (PGO) technique (see also PE4.1 Tuning Without Modifying the Source Code)

Ability to use software building environments like Autotools, CMake, Scons, and Waf

### USE4 Developing Parallel Programs

*Description:*

Ability to develop parallel programs

Ability to build parallel programs is required (see also USE3 Building of Parallel Programs)

Ability to develop software is required (see also SD Software Development)

## PE Performance Engineering

*Description:*

Ability to use systematic approaches (e.g. benchmarking and tuning, cost models) to meet performance requirements in a cost-effective way, i.e. by reducing the runtimes of parallel programs and using the resources of the HPC system appropriately for that purpose

### PE1 Cost Awareness

*Description:*

Ability to assess the costs related to the runtimes of parallel programs (see also K5 Modeling Costs)

Ability to assess the ratio of personnel costs to resource costs against the background of break-even considerations and time-to-solution constraints

## PE2 Measuring System Performance

*Description:*

Ability to measure the system performance by the help of standard tools and by profiling in order to assess the runtime behavior of parallel programs

### PE2.1 Using Standard Tools to Measure System Performance

*Short Background:* This includes information about utilization of resources like CPU as well as elapsed runtimes of a program, its unshared and shared memory usage, input and output statistics for devices and file systems, and page faults, with tools like `/usr/bin/time`, `ps`, `top`, `htop`, `vmstat`, `iostat`, and `perf` in Linux-based environments.

*Description:*

Ability to use standard tools of the operating system to get information about the behavior of parallel programs in terms of their resource utilization

### PE2.2 Profiling

*Short Background:* Profiling is explained for the CPU level, where it can be supported by hardware performance counters and by sampling techniques. Sampling is used to see, by examining the program counter, what routines and source code lines of a program are responsible for which portions of the total runtime. Automatically adding trace code to a parallel program by so-called instrumentation to record its execution in a strict chronology is explained and the difference to profiling is emphasized. Similar techniques are explained for profiling the network level (e.g. based on InfiniBand counters and I/O server states).

*Description:*

Ability to get the base data for tuning the performance of parallel programs by profiling

Ability to detect performance issues and bottlenecks caused, for example, by inefficient programming, memory accesses, I/O operations, cache-misses, page-faults, and parallelization overheads (see also K3.2 Parallelization Overheads)

Ability to assess how different views of the profiling data (e.g. timeline graphs and communication matrices to illustrate the traffic between processes) can give insights in the runtime behavior of the program

Ability to use performance analysis tools like ScoreP, Scalasca, Vampir

Ability to use the standard MPI profiling interface (PMPI) and environment variables like `$I_MPI_STATS` to control the built-in performance analysis functionality in MPI

## PE3 Benchmarking

*Description:*

Ability to assess speedups and efficiencies as the key measures for benchmarks of a parallel program (see also K2.1 Performance Frontiers)

Ability to benchmark the runtime behavior of parallel programs, performing controlled experiments by providing varying HPC resources (e.g. 1, 2, 4, 8, ... cores on shared memory systems or 1, 2, 4, 8, ... nodes on distributed systems for the benchmarks)

Ability to differentiate between strong and weak scaling

Ability to assess the performance characteristics of parallel programs with regard to CPU usage, memory accesses (e.g. latencies for random access, cache sizes,

strided access patterns, and bandwidth), I/O operations (e.g. record length, IOPs, latency, bandwidth, throughput, and multi-stream processing), and communication (message sizes, network bandwidth and latency)

#### PE4 Tuning

*Description:*

Ability to tune a parallel program in order to achieve better runtimes and to optimize the usage of the HPC resources

##### PE4.1 Tuning Without Modifying the Source Code

*Description:*

Ability to select appropriate tasks sizes (big vs. small) that may have positive performance impacts on the workflow, and to run several (smaller) tasks by the help of job chaining (see also USE2 Running of Parallel Programs)

Ability to use mapping of processes to nodes, pinning of processes/threads to CPUs or cores, and setting memory affinities to NUMA nodes in order to speed up a parallel program

Ability to speed up program execution by using optimized libraries and setting appropriate compiler/linker options (including PGO workflow)

Ability to speed up program execution by setting appropriate runtime options (e.g. for MPI and OpenMP)

Ability to speed up program execution by setting package specific options (e.g. selected by environment variables and command line arguments)

##### PE4.2 Tuning via Reprogramming

*Short Background:* The potential for tuning via reprogramming exists on the hardware as well as on the software level. At the software level, performance improvements are achievable by using more efficient algorithms. This is explained by the help of popular practice-relevant examples.

*Description:*

Ability to to reprogram appropriate parallel code for improved performance on the processing element level e.g. by using functional units (for executing fused multiply-add instructions and variants thereof), by using vectorization techniques with SIMD instructions, etc.

Ability to assess how appropriate computationally intensive functions (which have been identified earlier by profiling the parallel program) can be ported to many core architectures like GPUs to achieve further speedups

#### PE5 Optimization Cycle

*Short Background:* The workflow is represented by an optimization cycle with the steps benchmarking, gathering system performance data (e.g. via profiling), analyzing, and tuning.

*Description:*

Ability to apply the full workflow for tuning a parallel program

#### SD Software Development

*Description:*

Ability to develop parallel programs

##### SD1 Efficient Algorithms and Data Structures

*Description:*

Ability to assess the efficiency of algorithms and data structures, especially with respect to their suitability for typical (scientific) parallel programs, e.g. by the help of popular practice-relevant examples

## SD2 Programming

*Short Background:* The user learns how to complete programming tasks and gets a short overview of machine- and assembly-languages toward so-called high-level programming languages. The focus lies on the programming languages that are in widespread use within the HPC community.

*Description:*

Ability to program in languages typically used in HPC environments, such as C, C++, FORTRAN, HPX

Ability to use interoperability between languages, for example by calling C or C++ from FORTRAN and vice versa

Ability to use integrated development environments (IDEs) like Eclipse, e.g. to seamlessly perform the typical development cycle with the steps edit, build (compile and link), and test

Ability to debug a program using simple techniques such as inserting debugging output statements into the source code e.g. using `printf` – also against the background of potential problems with the ordering of the (stdout) output that may exist in parallel environments like MPI

Ability to use sophisticated debuggers such as GDB, DDT, and TotalView

## SD3 Parallel Programming

*Short Background:* Parallel programming of shared memory systems and message passing systems as well as load balancing is addressed.

*Description:*

Ability to assess the parallel nature of algorithms

### SD3.1 Parallel Algorithms

*Description:*

Ability to understand that some algorithms are embarrassingly (i.e. trivially) parallelizable while their parallelization will vary from easy to hard in practice

Ability to assess that there are algorithms having a so-called sequential nature that have been notoriously difficult to parallelize, for example alpha-beta game-tree search

Ability to determine the computational complexity of algorithms

### SD3.2 Programming Shared Memory Systems

*Short Background:* The parallel concepts of threads and processes are introduced and their impacts on performance are outlined.

*Description:*

Ability to understand race conditions and to use synchronization mechanisms to avoid them

Ability to understand the problems that may result from erroneous use of synchronization mechanisms (e.g. deadlocks)

Ability to assess parallel concepts typically used for shared memory systems, e.g. to exploit temporal locality by data reuse with an efficient utilization of the memory hierarchy

Ability to assess concepts like software pipelining, e.g. to optimize loops by out-of-order execution, and vectorization principles

Ability to assess data dependency situations, i.e. an instruction reading the data written by a preceding instruction in the source code, and anti-dependencies, i.e. an instruction having to read data before a succeeding instruction overwrites it, and output dependencies, i.e. instructions writing to the same memory location

Ability to assess the influence of control dependencies by jumps, branches, and function calls, e.g. on pipeline filling

Ability to use data parallelism, e.g. applying parallel streams of identical instructions to different elements of appropriate data structures such as arrays

Ability to understand the concept of functional parallelism, i.e. executing a set of distinct functions possibly using the same data

Ability to assess the applicability of parallel language extensions like OpenMP, CUDA, OpenACC, and C++ AMP, as well as their interoperability (e.g. combining OpenACC and CUDA)

### SD3.3 Programming Message Passing Systems

*Short Background:* Communication plays a central role in message passing systems. When parallel processes cannot or should not exchange information via shared memory, they typically send messages to each other to communicate.

*Description:*

Ability to understand the various communication modes (e.g. blocking vs. non-blocking, point-to-point vs. collective) and the concept of overlay networks

Ability to develop programs using MPI as the de-facto standard for parallelizing programs in distributed environments like HPC cluster systems

Ability to understand how race conditions and deadlocks may occur in MPI parallelized programs and how they can be avoided, namely by reordering send and receive operations or using non-blocking communication combined with waiting for completion of the communication operations concerned

Ability to assess the impact of communication and synchronization on the performance of a parallel program (see also K3.2 Parallelization Overheads)

### SD3.4 Load Balancing

*Description:*

Ability to apply domain decomposition strategies (see also K3.3 Domain Decomposition)

Ability to apply simple scheduling algorithms like task farming to achieve an appropriate distribution of the workloads across the multiple computing resources of the HPC system

Ability to apply more sophisticated approaches e.g. based on tree structures like divide-and-conquer or work-stealing to achieve an appropriate distribution of the workloads across the multiple computing resources of the HPC system

### SD3.5 I/O Programming

*Description:*

Ability to assess general concepts of HPC I/O systems (e.g. parallel file systems, see also K1.3 I/O Architectures) and how to map the data model to the storage system, e.g. by using appropriate I/O libraries and middleware architectures

## SD4 Object Oriented Approach

*Description:*

Ability to apply object oriented methods, i.e. object oriented analysis (OOA), design (OOD), and programming (OOP) (particularly to scientific and parallel programming)

Ability to apply design patterns to HPC, e.g. patterns for coding of parallel algorithms and their mapping to various architectures

#### SD5 Agile Methods

*Short Background:* The advantages of a test-driven development and of applying automated testing (e.g. using unit and integration tests) as well as coding guidelines and code refactoring are addressed. The idea of continuous integration and tools like Jenkins and Buildbot are presented. Portability aspects are taken into account e.g. for the source code of programs and job scripts to avoid typical compiler-, linker-, and MPI-issues.

*Description:*

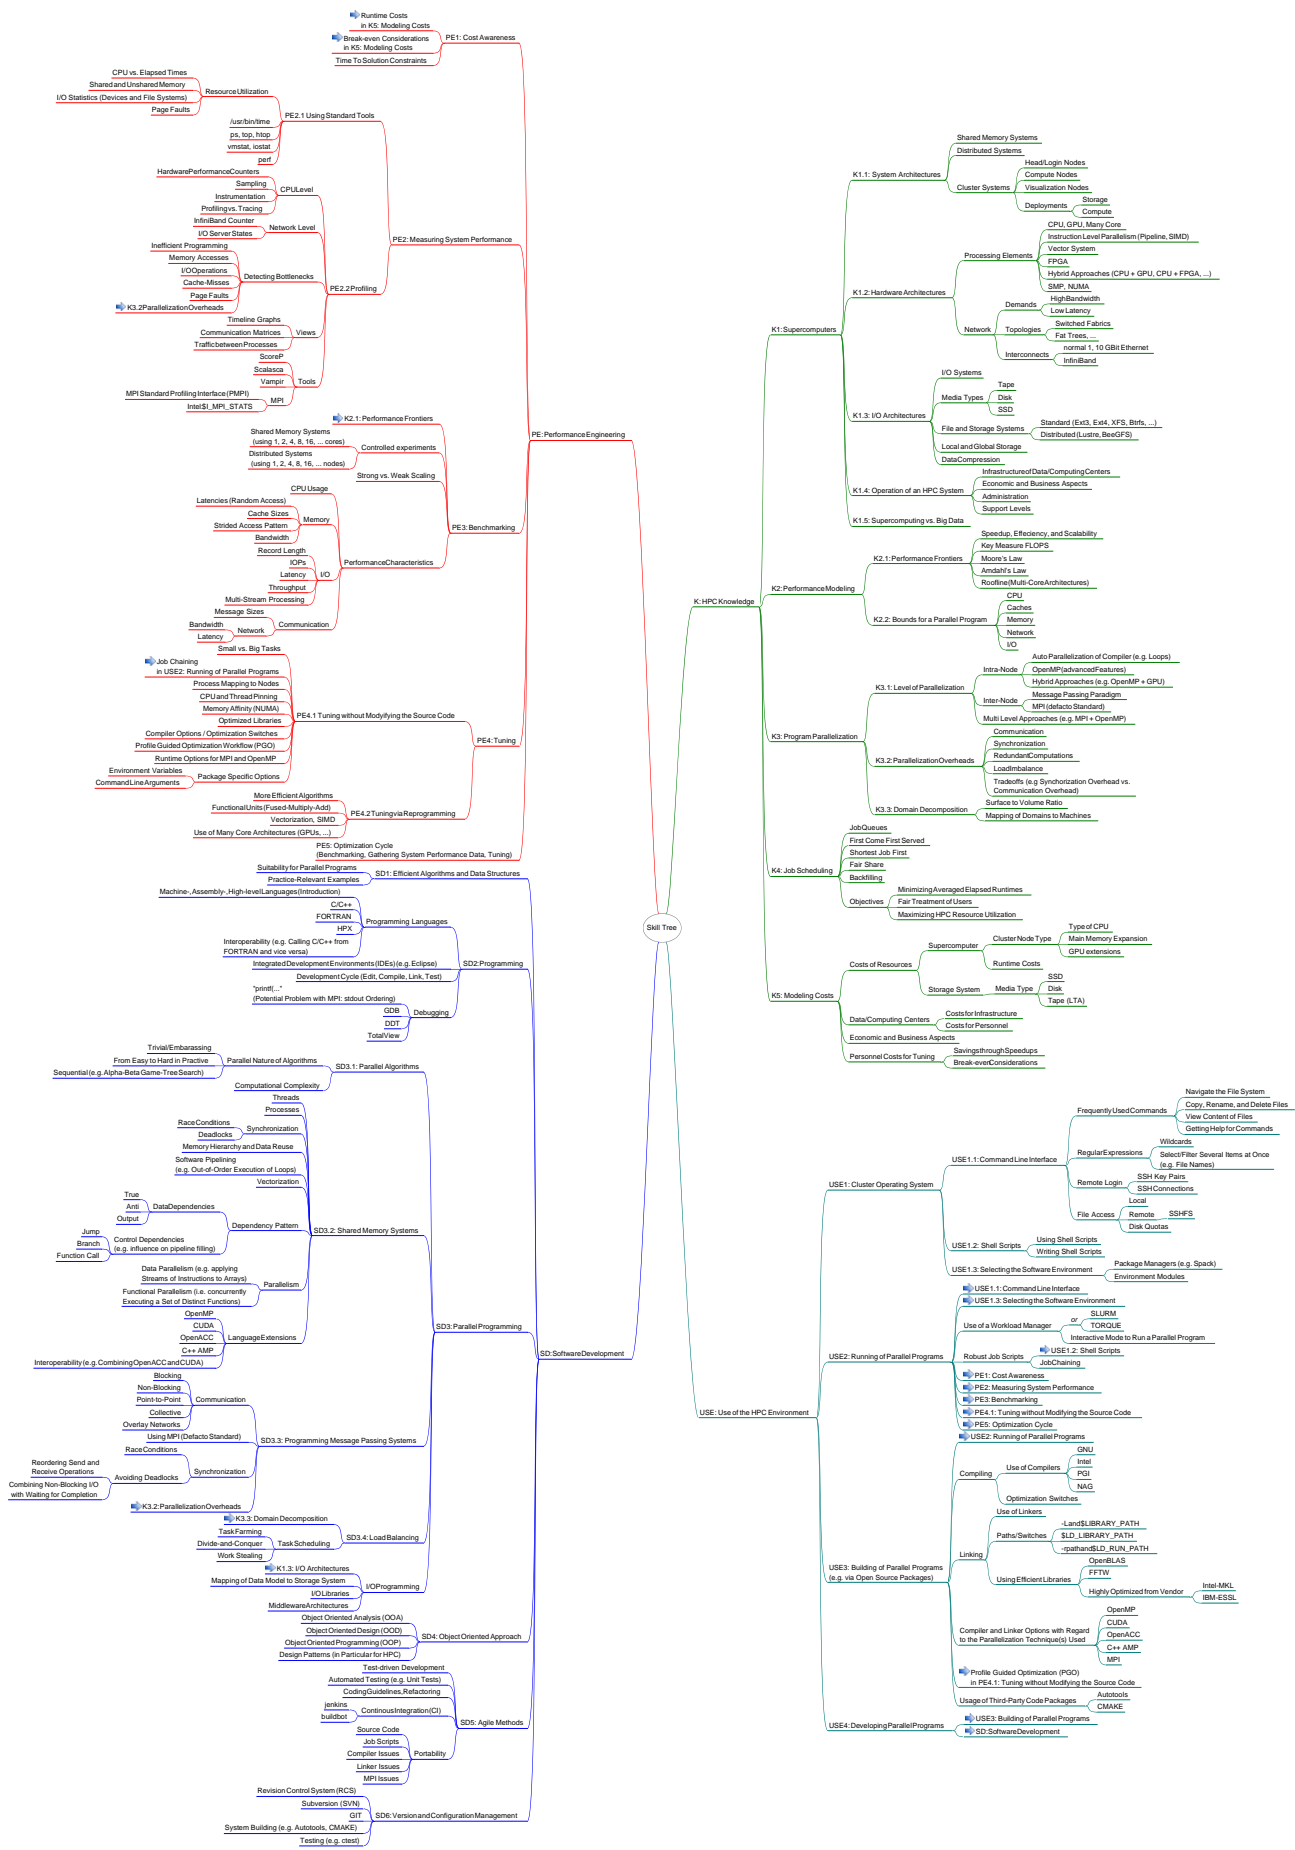
Ability to apply agile methods for scientific computing

#### SD6 Version and Configuration Management

*Short Background:* Systems like Revision Control System (RCS), Subversion (SVN), and GIT are presented as well as tools to support the building and testing of the software like Autotools, CMake, and CTest.

*Description:*

Ability to apply version and configuration management to the development of (parallel) programs in order to track and control changes in the sources, establish and maintain consistency of the program or software system throughout its life, and facilitate cooperative development



HPC Skill Tree