Best Practice Guide - Application porting and code-optimization activities for European HPC systems

Sebastian Lührs, JSC, Forschungszentrum Jülich GmbH, Germany

Björn Dick, HLRS, Germany

Chris Johnson, EPCC, United Kingdom

Luca Marsella, CSCS, Switzerland

Gerald Mathias, LRZ, Germany

Cristian Morales, BSC, Spain

Lilit Axner, KTH, Sweden

Anastasiia Shamakina, HLRS, Germany

Hayk Shoukourian, LRZ, Germany

30-04-2020

Table of Contents

1. Introduction

With the increasing number of European High Performance Computing (HPC) systems and the targeted road towards Exascale computing, HPC applications always faced a diversity of different system types and environment approaches. Typically, one major task of each HPC application is the porting and code-optimization activity to take advantage of all system capabilities and to run the application with the highest efficiency. Normally the porting of an optimization work can not be achieved just by the application developer. HPC experts are typically involved during the certain stages of application development to (i) help users in resolving all the possible issues arisen due to application porting as well to (ii) guide users in achieving the maximum possible performance from the target high-end IT platform.

Due to this support demand, various HPC support activities were established during the last years to offer general HPC knowledge as well as scientific area specific optimization strategies towards the research and industrial community.

This guide provides an overview about the larger running application porting and code-optimization activities for European HPC systems. It does not cover specific technical details of certain application but guides application developers and users to the right project or activity, which can help in context of certain HPC codes or scientific areas.

This guide will start with an introduction of EuroHPC, which deals a framework for several upcoming application support programs. The third chapter covers the porting and optimization activities within the PRACE-6IP project such as Preparatory access, SHAPE or DECI. This section is followed by an overview about the current generation of Centres of Excellence on HPC. Finally, examples of national HPC support activities are given.

2. EuroHPC Joint Undertaking

The EuroHPC Joint Undertaking is a joint initiative between the EU and European countries to develop a world-class supercomputing ecosystem in Europe.

The Joint Undertaking has the aim to acquire and provide a world-class petascale and pre-exascale supercomputing and data infrastructure for Europe's scientific, industrial and public users, matching their demanding application requirements by 2020 and to support an ambitious research and innovation agenda to develop and maintain in the EU a world-class HPC ecosystem. Activities are supported through procurement and open Calls. The Joint Undertaking will initially operate from 2019 until 2026.

The first level of funding will kick-start the activities to address the overall European HPC strategy, mainly: to acquire two pre-exascale machines and several petascale systems by 2020, and R&I actions covering the full HPC ecosystem launched in 2019 and 2020, including the support for HPC Competence Centres.

More information: https://eurohpc-ju.europa.eu

3. Application porting and optimization activities in PRACE

PRACE itself runs and offers a wide variety of different porting and optimization activities. In the following sections the preparatory access program, the SHAPE initiative, the DECI call, the high level support teams as well as the "forward looking software solutions" activity are presented.

3.1. Preparatory Access

3.1.1. Overview

Access to PRACE Tier-0 systems is managed through PRACE regular calls, which are issued twice a year. To apply for Tier-0 resources the application must meet technical criteria concerning scaling capability, memory requirements, and runtime set up. There are many important scientific and commercial applications that do not meet these criteria today. To support researchers PRACE offers the opportunity to test and optimise their applications on the envisaged Tier-0 system prior to applying for a regular production project. This is the purpose of the Preparatory Access (PA) Call. It is possible to choose between four different types of access:

- Type A is meant for code scalability tests, the outcome of which is to be included in the proposal in a future PRACE regular call. Users receive a limited number of core hours; the allocation period is two months.
- Type B is intended for code development and optimisation by the user. Users get also a small number of core hours; the allocation period is six months.
- Type C is also designed for code development and optimisation with the core hours and the allocation period being the same as for Type B. The important difference is that Type C projects receive special assistance by PRACE experts from the PRACE-IP project to support the optimisation. As well as access to the Tier-0 systems, the applicants also apply for one to six person months (PMs) of supporting work to be performed by PRACE experts.
- Type D allows PRACE users to start a code adaptation and optimisation process on a PRACE Tier-1 system. PRACE experts help in the system selection process. In addition to Tier-1 computing time, the PRACE user will also receive Tier-0 computing time towards the end of the project in the form of a PA Type A project to test the scalability improvements. The work is supported by PRACE experts similar to Type C. The maximum duration of Type D projects is twelve months.

More information: https://prace-ri.eu/hpc-access/preparatory-access/

3.1.2. Application and review processes

The Call for Proposals for PRACE Preparatory Access Type A and Type B schemes is a continuously open call, with a maximum time-to-resources-access (start-date) of two weeks after the date of submission.

Requests for Type C and D scheme are evaluated and granted on a quarterly schedule, with cut-offs for evaluations to be set on the first working day of March, June, September and December. Awarded proposals for Type C will have a maximum time-to-resources-access (start-date) of two months after the relevant cut-off date.

All proposals will be evaluated on a technical level. In addition, Type C and Type D proposals are evaluated by members of the PRACE-IP project to verify the requested level of support.

3.1.3. Project white papers

Most of the supported preparatory access projects (Type C and Type D) publish a white paper which covers the technical details and results. The following overview lists the white papers which were published online [4] during the last years as part of the PRACE-4IP, PRACE-5IP and PRACE-6IP projects. These guides can provide helpful information how to improve or adapt certain applications or general algorithms to the different HPC system architectures:

Table 1. PRACE preparatory access white papers

White paper title	Link
WP207: Hybrid MIMD/SIMD High Order DGTD Solver for the Numerical Modeling of Light/Matter Interaction on the Nanoscale	https://prace-ri.eu/wp-content/uploads/WP207.pdf
WP208: Large Scale Parallelized 3d Mesoscopic Simulations of the Mechanical Response to Shear in Disordered Media	https://prace-ri.eu/wp-content/uploads/WP208.pdf
WP209: Optimising PICCANTE – an Open Source Particle-in-Cell Code for Advanced Simulations on Tier-0 Systems	https://prace-ri.eu/wp-content/uploads/WP209.pdf
WP210: Parallel Subdomain Coupling for non-matching Meshes in Alya	https://prace-ri.eu/wp-content/uploads/WP210.pdf
WP211: Numerical Simulation of complex turbulent Flows with Discontinuous Galerkin Method	https://prace-ri.eu/wp-content/uploads/wp211.pdf
WP221: Parallel I/O Implementation in Hybrid Molecular Dynamics-Self Consistent Field OCCAM Code using fpioLib Library	https://prace-ri.eu/wp-content/uploads/WP221.pdf
WP222: Optimising UCNS3D, a High-Order finite-Volume WENO Scheme Code for arbitrary unstructured Meshes	https://prace-ri.eu/wp-content/uploads/WP222.pdf
WP223: Enabling Space Filling Curves parallel mesh partitioning in Alya	https://prace-ri.eu/wp-content/uploads/WP223.pdf
WP224: Parallel curved mesh Subdivision for flow Simulation on curved Topographies	https://prace-ri.eu/wp-content/uploads/WP224.pdf
WP225: Particle Transport in a Fluid interacting with an immersed Body with Alya	https://prace-ri.eu/wp-content/uploads/WP225.pdf
WP258: Optimization of REDItools package for investigating RNA editing in thousands of human deep sequencing experiments	https://prace-ri.eu/wp-content/uploads/WP258.pdf
WP266: A numerical code for the study of water droplets growth, collision, coalescence and clustering inside turbulent warm cloud-clear air interfaces	https://prace-ri.eu/wp-content/uploads/WP266.pdf
WP269: Optimisation of EC-Earth 3.2 model	https://prace-ri.eu/wp-content/uploads/WP269.pdf
WP270: Extending the scalability and parallelization of SimuCoast code to hybrid CPU+GPU supercomputers	https://prace-ri.eu/wp-content/uploads/WP270.pdf
WP271: Optimizing ART: Radiative Transfer Forward Modeling code for Solar Observations with ALMA	https://prace-ri.eu/wp-content/uploads/WP271.pdf
WP274: Automation of High-Fidelity CFD Analysis for Aircraft Design and Optimization	https://prace-ri.eu/wp-content/uploads/WP274.pdf
WP284: Scalable Delft3D Flexible Mesh for efficient modeling of shallow water and transport processes	https://prace-ri.eu/wp-content/uploads/WP284.pdf
WP287: Performance Assessment of Pipelined Conjugate Gradient method in Alya	https://prace-ri.eu/wp-content/uploads/WP287.pdf

White paper title	Link
WP288: Optimisation of the Higher-Order Finite-Volume Unstructured Code Enhancement for Compressible Turbulent Flows	https://prace-ri.eu/wp-content/uploads/WP288.pdf

3.2. SHAPE

3.2.1. Overview

The small and medium enterprises (SME) HPC Adoption Programme in Europe (SHAPE) runs under the Applications Enabling and Support work package under the PRACE-6IP project. The use of HPC can be very powerful for SMEs in terms of increasing competitiveness via the development of new products or services, reduction in the time-to-market and costs of research and development or an increase in quality of products or services. However, the costs and risks of SMEs trying out HPC can be large, particularly if the SME has little or no experience of HPC. The SHAPE programme exists to help SMEs get a foot on the HPC ladder by working with an SME to develop the use of HPC within their workflows. This may simply be giving them access to an appropriate HPC system and helping them to try out an existing HPC code, or it may be taking their software and enabling it for HPC to see if running on an HPC works for the SME. If successful, an SME is provided with effort from an HPC expert at the nearest PRACE centre and access to appropriate HPC resources is provided. In exchange the SME provides domain expertise and helps contribute a white paper produced at the end of the project.

The programme began with a pilot call in 2013 and has now issued 10 separate calls with 54 SMEs across 14 different countries having been awarded effort. Since the fifth call in 2017 calls have been issued every 6 months receiving a total of 73 proposals with 55 projects awarded funding.

Within PRACE we are keen to increase the range of countries from which proposals originate. To help facilitate this a new SHAPE+ initiative has been launched which provides an allocation of effort made available to PRACE sites for assisting projects from countries new to SHAPE.

More information: https://prace-ri.eu/prace-for-industry/shape-access-for-smes/

3.2.2. Application and review processes

Calls are published on the PRACE website and publicised via appropriate channels, which consist of flyers, presence at trade associations and workshops, as well as direct contact with SMEs. Applicants then apply via a form on the PRACE website and can either submit via an on-line form [5] or via an MS Word document which contains the same information as the on-line form. It is recognised that many SMEs have very little experience with HPC and so several sections of the application form are optional. SMEs who are assisted by a PRACE centre to fill in the forms may well be able to provide more information than those who have not and this is taken into account by the panel when assessing proposals. No personal details are given when applying other than name and email address of an applicant.

When the call closes, the proposals are assessed by a panel consisting of SHAPE staff, PRACE technical experts, the chair of the PRACE Industrial Advisory Committee, business development officers from PRACE sites around Europe and PRACE peer review officers. The panel meets and ranks the proposals based on the expected business outcomes of a given project, also taking into account the technical viability of the project.

After the panel meeting, the list of recommendations is taken to the PRACE Board of Directors who then formally decide which projects to approve. The SMEs are then informed of the decisions and are provided with feedback from the panel, particularly for any SMEs not successful. Any conditions on acceptance will be given to the SMEs at this time. Unsuccessful applicants are invited to apply to the next call, or pointed to other appropriate calls (e.g. PRACE Preparatory Access calls).

3.2.3. Project Outcomes

At the end of the project the SMEs together with the PRACE centres involved produce a white paper [4] following a template provided. This explains the context of the project and gives some technical information about what was achieved and how this has benefited the SME's business. A list of these white papers is given below.

In addition, around a year after a project has finished, both SMEs and PRACE the centres involved are invited to fill in a survey which provides information on any new or enhanced software features, technical impact (e.g. increased speed of code, or greater data processing capabilities, overall performance improvements, etc.), any change in take up of HPC or improved HPC skills within the SME, previous and present HPC access, future plans, return on investment, business outcomes, etc.

Table 2. SHAPE white papers

White paper title	Link
WP291: Vision-e: Deep Tile Inspection	https://prace-ri.eu/wp-content/uploads/WP291.pdf
WP290: Axyon: a scalable HPC Platform for AI Algorithms in Finance	https://prace-ri.eu/wp-content/uploads/WP290.pdf
WP283: Shiloh Industries Italia: Multiphasic Simulation in High-Pressure Die Casting Process	https://prace-ri.eu/wp-content/uploads/WP283.pdf
WP273: High-fidelity Simulation of an industrial swirling Combustor	https://prace-ri.eu/wp-content/uploads/WP273.pdf
WP221: Parallel I/O Implementation in Hybrid Molecular Dynamics-Self Consistent Field OCCAM Code using fpioLib Library	https://prace-ri.eu/wp-content/uploads/WP221.pdf
WP272: AXESSIM – CINES Partnership: HPC for connected Objects	https://prace-ri.eu/wp-content/uploads/WP272.pdf
WP268: Numerical Modelling of 3D Printed Cranial Orthoses	https://prace-ri.eu/wp-content/uploads/WP268.pdf
WP256: Airinnova: Automation of High Fidelity CFD Analysis in Aerodynamic Designs	https://prace-ri.eu/wp-content/uploads/WP256.pdf
WP267: Milano Multiphysics: Evaluation of the Intel Xeon Phi performances for high fidelity nuclear applications	https://prace-ri.eu/wp-content/uploads/WP267.pdf
WP258: Optimization of REDItools package for investigating RNA editing in thousands of human deep sequencing experiments	https://prace-ri.eu/wp-content/uploads/WP258.pdf
WP263: Scienomics Development of Chameleon Monte Carlo code for HPC: Toward Realistic Modelling of Complex Composite Systems	https://prace-ri.eu/wp-content/uploads/WP263.pdf
WP264: Renuda: Optimising 2D flow for faster, better steam turbine design	https://prace-ri.eu/wp-content/uploads/WP264.pdf
WP259: AmpliSIM DemocraSIM: Enable air quality Simulation as a Commodity	https://prace-ri.eu/wp-content/uploads/WP259.pdf
WP250: Anemos SRL: SUNSTAR Simulation of UNSteady Turbulent flows for the AeRospace industry	https://prace-ri.eu/wp-content/uploads/WP250.pdf
WP260: Creo Dynamics: Scale Resolving CFD and CAA processes for Ground Vehicles based on Open Source	https://prace-ri.eu/wp-content/uploads/WP260.pdf
WP249: ACOBIOM – CINES Partnership: MARS: Matrix of RNA-Seq	https://prace-ri.eu/wp-content/uploads/WP249.pdf
WP251: WB-Sails: Modelling Sail Boat Performance in HPC	https://prace-ri.eu/wp-content/uploads/WP251.pdf
WP252: FDD Engitec S.L.: Pressure drop simulation for a compressed gas closed system	https://prace-ri.eu/wp-content/uploads/WP252_SHAPE-Corrected-patent.pdf
WP253: Pharmacelera S.L.: HPC Methodologies for PharmS	https://prace-ri.eu/wp-content/uploads/WP253.pdf

White paper title	Link
WP254: BAC Engineering Consultancy Group: Numerical Simulation of Accidental Fires with a Spillage of Oil in Large Buildings	https://prace-ri.eu/wp-content/uploads/WP254.pdf
WP255: Algo'Tech: A Successful Switch to HPC	https://prace-ri.eu/wp-content/uploads/WP255.pdf
WP256: Airinnova: Automation of High Fidelity CFD Analysis in Aerodynamic Design	https://prace-ri.eu/wp-content/uploads/WP256.pdf
WP217: Cybeletech: Parallelization and optimisation for plant selection with Cybeletech	https://prace-ri.eu/wp-content/uploads/WP217.pdf
WP218: Design Methods:Coupled Sail and Appendage Design Method for Multihull Based on Numerical Optimisation	https://prace-ri.eu/wp-content/uploads/WP218.pdf
WP220: Ingenieurburo Tobias Loose: HPCWelding: Parallelized Welding Analysis with LS-DYNA	https://prace-ri.eu/wp-content/uploads/WP220.pdf
WP213: Open Ocean: High Performance Processing Chain – faster on-line statistics calculation	https://prace-ri.eu/wp-content/uploads/WP213.pdf
WP215: Optimad Engineering srl.: RAPHI: Rarefied Flow Simulations on Xeon Phi Architecture	https://prace-ri.eu/wp-content/uploads/WP215.pdf
WP216: Vortex Bladeless: Parallel multi-code coupling for Fluid-Structure Interaction in Wind Energy Generation	https://prace-ri.eu/wp-content/uploads/WP216.pdf
WP219: Hydros Innovation: Automatic Optimal Hull Design by Means of VPPApplications on HPC Platforms	https://prace-ri.eu/wp-content/uploads/WP219.pdf
WP214: Ergolines – ITU Partnership: HPC-based Design of a Novel Electromagnetic Stirrer for Steel Casting	https://prace-ri.eu/wp-content/uploads/WP214.pdf
WP189: AMET s.r.l.: ROSPA: Robustness in Safety Performances Analysis	https://prace-ri.eu/wp-content/uploads/WP189.pdf
WP190: OPTIMA pharma GmbH: Enhanced Airflow Simulations around Filling Machines in Clean Rooms	https://prace-ri.eu/wp-content/uploads/WP190.pdf
WP195: NSilico Lifescience Ltd: Novel HPC Technologies for Rapid Analysis in Bioinformatics	https://prace-ri.eu/wp-content/uploads/wp195.pdf
WP194: Monotricat SRL: Hull resistance simulations for an innovative hull using OpenFOAM	https://prace-ri.eu/wp-content/uploads/wp194-updated.pdf
WP197: Juan Yacht Design: Testing LES turbulence models in race boat sails	https://prace-ri.eu/wp-content/uploads/wp197.pdf
WP193: ENTARES Engineering: Electromagnetic simulation for large model using HPC	https://prace-ri.eu/wp-content/uploads/wp193-entares-shape.pdf
WP191: Albatern: Numerical Simulation of Extremely Large Interconnected Wavenet Arrays	https://prace-ri.eu/wp-content/uploads/wp191_albatern.pdf
WP192: AUDIONAMIX: A benchmark of linear algebra libraries for HPC	https://prace-ri.eu/wp-content/uploads/wp192-audionamix-shape.pdf
WP188: Thesan srl: Design improvement of a rotary turbine supply chamber through CFD analysis	https://prace-ri.eu/wp-content/uploads/wp188-final.pdf

3.3. DECI support

3.3.1. Overview

The Distributed European Computing Initiative (DECI) began in 2005 as part of the Distributed European Infrastructure for Supercomputing Applications (DEISA), later moving into the PRACE Implementation Phase (IP) projects as well as having been run as a PRACE “Optional Programme”. Presently it runs under the Applications Enabling and Support work package within the PRACE-6IP project. The most recent DECI call (DECI-16) was the 16th call since the programme began.

DECI provides access to Tier-1 resources via a resource exchange programme using the juste retour principle. Here countries may contribute resources to DECI and will in total receive at least 70% of the amount of resources contributed to DECI by the project leader’s own country. Individual projects will get access to the most appropriate architecture for their project’s needs which will usually be housed in a different country. The remaining resources (up to 30% of the total computing time) will be awarded to projects from countries not contributing to the call. This way, all European countries can benefit from DECI even when individual countries may not have contributed themselves. This helps to provide resources more evenly across Europe and also helps to cover temporary or longer-term gaps in the national provision of resources.

Since the DECI programme began, we have received over 1000 proposals. The programme continues to be very popular with researchers across Europe with the last two DECI calls (DECI-14 and DECI-15) receiving 66 and 73 proposals respectively. In both cases 40 projects were awarded. The DECI-16 closed on 31st January 2020 receiving 68 proposals.

3.3.2. Application and review processes

Calls are published on the PRACE website (e.g. https://prace-ri.eu/call/call-for-proposals-for-deci-16/) and publicised via appropriate pan-European mailing lists as well as being promoted at a national level. Applicants apply via a dedicated submission tool hosted by the French site, CINES [11]. The tool is fully GDPR compliant and is based on the tool used for PRACE Project Access for Tier-0 resources.

After the call closes, each proposal is allocated a “home site” which is the site which will carry out the technical evaluation and, if the project is awarded, will act as the contact point for the project. Home sites are allocated to the site within the country of origin of the proposal where such a site is available. For proposals originating from countries not represented in DECI, home sites are allocated in such a way that the amount of work is balanced across the sites, but also usually taking into account geographical proximity.

During the technical evaluation, candidates may be asked clarification questions, the responses to which can be noted on the technical evaluation.

Following technical evaluation, with the exception of proposals from countries not contributing resources to the call, proposals are evaluated by national panels based on the countries of origin of the proposals. These panels follow the same processes as they would when assessing proposals for national calls. A ranking of proposals is produced for each country. From the amount of resources contributed by each country, it can be determined how many proposals from each country can be accepted.

Proposals from non-contributing countries are assessed by peer review from experts around Europe. In most cases these are reviewers who also review proposals from Tier-0 PRACE Project Access and similar conflicts of interest, confidentiality rules are followed. Again, this panel produces a ranked list of proposals. From the amount of resources available to non-contributing countries, it can be determined how many proposals from each country can be accepted.

PIs are informed of the results via email with a formal letter explaining the acceptance or otherwise of their proposal. Successful projects are then explained the requirements for the project, e.g. any acceptable use policy documents to be signed, guidelines on acknowledging DECI, final report requirements, etc. Feedback from review panels is available to all applicants.

3.3.3. Project Outcomes

At the end of each DECI project, PIs are required to produce final reports with a short template being provided. These reports are not published directly, but may be used as part of any publicity carried out by PRACE, with a number of DECI projects having been highlighted in the PRACE Digest [6] in recent years for example. The most recent two calls have provided around 150 million machine core hours each to European researchers from 80 different DECI projects. Resources have come from a range of systems from 13 different countries and have included Cray XC30 and XC40, Intel clusters (various processor and memory configurations) and hybrid systems (clusters with GPGPU accelerators or Xeon Phi Co-processors (KNC)).

3.4. High-Level Support Teams

In order for the European scientific communities to fully benefit from the performance of leadership-class Tier-0 systems provided by the PRACE 2 Hosting Members, it is necessary to provide them with high-level support teams which could help projects when it comes to code enabling and optimisation of scientific applications. Each team is located at a Hosting Member Tier-0 site and it is composed of HPC experts having, on the one hand, strong skills both on specific HPC architectures provided by the Hosting Members as well as applied mathematics, computer science and management of data, and on the other hand are experienced in the applications fields they will support.

In a situation where major breakthroughs are expected in terms of novel HPC architectures due to energy constraints and huge amounts of data to be managed, these teams will be in charge mainly to provide level 2 user support to PRACE user communities for code enabling and scaling out of scientific applications and models. They cooperate among each other and with similar national efforts of all PRACE Members, creating at the European level a distributed pool of expertise, able to perform as well long term user support level 3 R&D activities jointly with user’s communities.

Coordinated together, these teams established a distributed European task force, able to reach the needed critical mass. In relation with the PRACE Scientific Steering Committee (SSC) and complementing the work done by the European Centers of Excellence (CoEs), these teams, by providing specific skills in different scientific domains, work on transverse actions including the extreme scalability of selected European scientific applications and tools towards Exascale, Data Analytics, and Machine Learning as well as the industrialisation and the long-term support/dissemination across PRACE RI of scientific applications.

3.4.1. HLST activities

The following sections provide an overview of example activities that are being developed by High-Level Support Teams located at BSC, CSCS, HLRS@GCS, JSC@GCS and LRZ@GCS.

3.4.1.1. BSC: Porting NEMO to POWER 9 with GPU

Nucleus for European Modelling of the Ocean-(NEMO) is a long-standing framework with several groups working on the development of new features. There is a growing demand for computational resources on Ocean and Climate Science communities for running this framework in order to improve climate simulations using more computational power. NEMO's workload involves multiphysics thermodynamics, dynamics, transporting, and biochemical routines containing many stencil calculations. Parallel runs rely on Message Passing Interface (MPI) calls, each associated with a computation subdomain.

Moreover, simulations involving dozens of thousands of cores are conventional in large HPC centers, but, at the moment, NEMO distributed parallelization strategies only includes CPUs. Regarding these characteristics, they present experience in improving NEMO's performance using CUDA. They describe their approach on a Fortran extense code regarding the port of routines and asynchronous execution opportunities.

The work involves the setup of NEMO in an IBM Power9 Machine with four V100 per node using IBM XLF compiler, including the execution routine profiles, port and benchmark routines, identification of asynchronously execution opportunities, and finally the merging of the new routines to NEMO's main code. Individually, for diagnostic prototype routines, the HLST achieve a speedup of 8x on a single node. They also use an accuracy methodology, provided by the Performance Earth Science-(PES) team on BSC, to validate the correctness of results due to float-point arithmetics differences in chaotic systems such as Earth System Models. In a larger scale simulation, the HLST accomplish an efficient execution of several MPI ranks in a single GPU using Multiple

Passing Service-(MPS), allowing a total NEMO speedup of 3.3% within a routine's time footprint of 5.4% on original CPU code. The previous speedup was obtained for an ORCA2 resolution run in single-node execution they have similar results in higher resolution as ORCA025 for up to 4 nodes execution. The HLST notice, as already expected, that the GPU performance gains depend on the model resolution. Higher-resolution corresponds to better GPU performance compared to the counterpart CPU only routine.

The role of memory transfers between host and device limits the strong scaling of the GPU diagnostic routine; this means that for a sufficient number of MPI ranks, the CPU-only diagnostic routine overcomes the GPU one. The poor strong scalability of this GPU port, together with the blocking execution, makes this approach unappealing without more code refactor. With that in mind, the HLST are working in a non-blocking asynchronous version that will be more useful in a broad range of model resolutions.

3.4.1.2. BSC: Parallelization of BioFVM

PhysiCell is a simulation software for modeling biological processes such as the release, uptake, decay, death, and diffusion etc., in cells of multicellular organisms. It allows movement, interaction of cells and the ability to study them as a function of multiple substrate diffusion and signalling factors. PhysiCell is completely written in C++ and uses Open Multiprocessing (OpenMP) to support shared memory parallelization. BioFVM is an integral component of PhysiCell and allows one to specify the external microenvironment, the substrates, and further, solves the Partial Differential Equations (PDEs) that model diffusion - after discretization using the Finite Volume Method (FVM). BioFVM uses a direct method called the Thomas algorithm that is employed for solving the tridiagonal system of linear equations resulting from the aforementioned discretization. A serious limitation of PhysiCell/BioFVM is that the simulation can only be run on a single node of an HPC cluster, as only an on-node shared memory parallelization is supported.

Their current work focuses on restructuring PhysiCell/BioFVM to support distributed parallelism using MPI. In their preliminary design, they have been able to parallelize the core kernels of BioFVM i.e. kernels that enable partition of the domain space, creation of basic agents, and the file I/O. The execution results for a chosen example show promising scaling for multiple kernels but also shed light on the non-optimal performance of their partially parallelized Thomas solver. Their efforts are directed towards an efficient distributed design for the simulation software and improving the performance of the hybrid implementation i.e. MPI + OpenMP. Specifically, they are focusing on the solver by considering a completely parallel implementation, the scalability of time-consuming kernels, and designing a strategy to account for the movement of cells across sub-domains.

3.4.1.3. CSCS: Porting EAGLE-XL on Piz Daint hybrid Cray XC50

The group of Richard Bower was supported successfully by CSCS HLST to port and run the code developed in-house to the hybrid architecture of the computing system "Piz Daint" at CSCS. SWIFT is a cosmological simulation software using particle-based algorithms to solve the equations of hydrodynamics and gravity. More specifically, Smooth Particle Hydrodynamics and a fast-multipole method are used. The software is open source and is currently mostly targeted at x86 architecture with some explicitly vectorized (AVX2) sections. The software is designed to be a drop-in replacement for the popular Gadget code and strong scaling performance on more than 250k cores has already been demonstrated [1] . A factor of 20x over Gadget is already obtained due to better algorithms. Offloading some kernels to the GPU will increase this figure. The parallelization uses a task-based framework in the form of our own C library QuickSched, which can handle conflicts, dependencies and asynchronous MPI communications between cluster nodes. The library is designed to allow for offload of some of the tasks to the GPU. Large time allocations for the use of the code have already been granted via the PRACE system.

3.4.1.4. CSCS: Performance improvement of COSMO on Piz Daint hybrid Cray XC50

The COSMO model is a rather new, flexible structure that enables to run the model on heterogeneous supercomputers, combining CPUs and GPU accelerators, with greater efficiency. There is an open collaboration with the domain scientists to improve the performance of the code on the system and to help integrate the new developments and scaling benchmarks. To better support the project on Piz Daint, the HLST is also actively working on including some relevant benchmarks in the CSCS regression testing framework (ReFrame) in order to be able to minimize the impact of Cray Linux Environment upgrades and interventions on the system. This is an ongoing

[1]https://arxiv.org/abs/1606.02738

team effort, involving different team members as COSMO is very complex code with a lot of dependencies and therefore building it on the system after every upgrade required a non-negligible effort from the users that the HLST is trying to reduce to a minimum.

3.4.1.5. HLRS@GCS: Optimization of SBLI

SBLI solves the governing equations of motion for a compressible Newtonian fluid using a high-order discretization with shock capturing. An entropy splitting approach is used for the Euler terms and all the spatial discretizations are carried out using a fourth-order explicit central-difference scheme. It should be noted that this approach was adopted from the start to enable effective and efficient parallelization based on halo-swapping. Time integration is performed using compact-storage Runge-Kutta methods with third and fourth order options. Stable high-order boundary schemes are used, along with a Laplacian formulation of the viscous and heat conduction terms to prevent any odd-even decoupling associated with central schemes. The fully parallel (MPI) version of the above scheme has been used to simulate transonic flow over a bump, turbulent spots and transitional shock-wave/boundary-layer interactions. No-slip fixed-temperature boundary conditions are applied at solid surfaces, while free stream and outflow boundaries use characteristic methods to reduce reflections. In order to remove the numerical grid-to-grid oscillations due to the central finite-difference schemes, a sixth-order standard centred explicit filter is available. Various sub-grid models are available including standard and dynamic Smagorinsky, a mixed model and a mixed time scale model that has been used previously for transitional flow calculations in supersonic boundary layer flow.

The SBLI code is a mature code, originating in 1999 and having undergone a number of updates and rolling optimizations in the meantime. One specific optimization to be tested is the use of an automatic code-generated routine, based on optimization found to be possible on the OpenSBLI project.

Right after the according PRACE project has been accepted by HLRS@GCS, the HLST did a performance audit. This audit revealed problems not only with parallelization but also with serial and I/O performance. With respect to parallelization the HLST figured out that domain decomposition was bad due to poor estimates of computational load via cell weights. Furthermore the process of assigning cells to processors was suboptimal. However, neither at HLRS nor at the user site there were enough resources to optimize both. The HLST hence focused on serial as well as I/O performance. With respect to serial performance, it was possible to speed up the code by removing data dependencies via loop splitting which allowed for auto-vectorization. This measure gave 11% speedup. THE HLST furthermore tweaked Lustre parameters and thereby reduced the time required for I/O by 37%.

3.4.1.6. JSC@GCS: Optimization of the AFiD code

In the ROHITU allocation on JUWELS, a fork of the open source CFD code AFiD was used which was based on an older, unrefactored version of AFiD. In addition, the immersed boundary method was implemented in the code version used. Due to these differences, the possible effect of those optimization developed for the ROHITU allocation were evaluated for the main stream version of AFiD. In addition, larger node counts were considered.

The HLST found that for the large test case roughly 8%, 11%, and 75% of the runtime of the time marching are spent on subroutines ExplicitTerms, ImplicitAndUpdate, and SolePressureCorrection, respectively. In the analysis of the ROHITU version of AFiD, the HLST identified the solution of the linear systems in step solving the continuity equation. Here, a multitude of relatively small linear equations with differing matrices, due to the immersed boundary method, needed to be solved. the HLST showed that replacing the original LAPACK-based solvers by an integrated implementation of the Thomas algorithm resulted in a considerable speed-up. Analysis of the main stream AFiD code showed that the time for solving the linear equations within the step solving the continuity equation was rather low. The reason for this is that all linear equations share the same matrix and hence a LU-decomposition can be pre-computed once and then be applied to all linear equations. In particular, this allows a very efficient implementation using vectorization within the LAPACK routines.

Further, the HLST experimented with replacing MPI_Alltoallv communication by MPI_Alltoallw with appropriate MPI subarray data types. Similar to the ROHITU version of AFiD, the HLST found a reduction in runtime. In particular, this reduction is large (-8% to -10% runtime) for the small test case but smaller (-2.9% to -4.5% runtime) for the larger test case. Hence, this modification requires additional validation for even larger test cases. Within subroutine ExplicitTerms, the HLST tested application of non-temporal stores and, for the large test case, observed a solid reduction in runtime of this subroutine by approximately 16.5%. Loop fusioning together with non-temporal stores within the subroutine resulted in a reduction of the runtime of the subroutine by approximately 26.5%. Still, further tests for consistency are required here.

Since most of the runtime is spent on solving the pressure equation (representing the continuity equation), any reduction of the product of internal steps (stages in the Runge-Kutta method) and timesteps can be expected to reduce the overall runtime. Hence, the HLST invested time in research on improved time marching schemes.

3.4.1.7. LRZ@GCS: ECHO-3DHPC

The General Relativity, Magneto-hydrodynamical (GRMHD) code ECHO-3DHPC is an astrophysical application developed to investigate the dynamics of magnetized thick tori orbiting around black holes. The LRZ-HLST helped to modernise the code in collaboration with its main developer Matteo Bugli (CEA Saclay, France) and with Fabio Baruffa (Intel).

With the help of the Intel Software Development Tools, like Fortran compiler and Profile-Guided Optimization (PGO), Intel MPI library, VTune Amplifier and Inspector we have investigated the performance of the baseline code. It turned out that a large fraction of the CPU time is classified as imbalance or serial spinning due to insufficient concurrency of working threads. Furthermore, the code is memory-bound, with about 80% of execution pipeline slots stalled because of demand from memory loads and stores. Based on our findings, we have increased the number of OpenMP parallel regions, changed the scheduling protocol to guided, and applied PRIVATE clauses to selected variables in OpenMP parallel loops, thus removing data races. The node-level performance is improved by 2.3x and, thanks to the improved threading parallelization, the hybrid MPI-OpenMP version of the code outperforms the MPI-only, thus lowering the MPI communication overhead. The results of the project obtained so far have been published [12]. In 2019, the project was finalized running scaling tests on SuperMUC-NG. It has been confirmed that the performance benefit obtained on previous architectures (Intel Xeon and Xeon Phi) is portable on the Xeon Scalable Processor Skylake of SuperMUC-NG.

3.5. Forward looking software solutions

"Forward looking software solutions" is a activity within the current generation of the PRACE-6IP project to focus on the diversity of the incoming new hardware infrastructure and the complexity of the codes in use by the scientific communities. Multiple separate sub-projects are supported as part of this activity, which are briefly presented in the following subsections.

3.5.1. PiCKeX: Particle Kinetic codes for Exascale plasma simulation

By leveraging advanced algorithmic features (scalable I/O, space-filling curves, task-based programming models) developed for a local PIC (particle-in-cell) code (JUSPIC), the project propose to implement key algorithmic kernels as a library to be used by codes like EPOCH. They also aim to add a complementary mesh-free algorithm within this library.

3.5.2. MoPHA: Modernisation of Plasma Physics Simulation Codes for Heterogeneous Exascale Architectures

This project address the software engineering challenge of focusing on three world-leading, and widely used European simulation codes in plasma physics. The codes are complementary concerning their functionalities, but they all have in common to rely on large grids of different dimensionality, ranging from 3D (ELMFIRE) to 5D (GENE) to 6D (Vlasiator). Essential tasks include the refactoring of the codes to make them more accessible and adaptable for modern programming methodologies, the implementation of task-based parallelism to achieve performance portability and scalability, and the knowledge transfer to share lessons learned and to encourage reuse of code parts in the plasma physics community and other science areas.

3.5.3. NB-Lib: Performance portable library for N-body force calculations at the Exascale

The goal of the NB-LIB project is to address GPU acceleration and maintaining parallelisation for scientific applications to simulate particle interactions (e.g. Molecular Dynamics, Monte Carlo or multiscale simulations in life sciences or materials). They want to achieve this by separating the lower-level parts of the GROMACS code into an API of non-bonded force routines to enable re-use of widely used, highly portable, and performant HPC code. The library will be available via different APIs, including both Python3 and C++14. This will permit domain scientists to use the API to prototype and deploy solutions for new N-body simulations rapidly, leveraging existing knowledge and best practices, rather than learning how to modify existing code.

3.5.4. LoSync: Synchronisation reducing programming techniques and runtime support

The LoSync project will focus on improving the scalability of applications by removing unnecessary synchronisation and serialisation, and full realising opportunities to overlap calculation and communication. For this they plan to leverage the capabilities of using OpenMP/OmpSs tasks with data dependency clauses, MPI single-sided communication and GASPI single-sided (put-notify) communication. A significant part of the work of this task will consist of implementing these techniques in key kernels of important applications such as Gysela, Flucs, IFS, Quantum Espresso, CP2K, iPIC3D, CASTEP, LULESH, NTCHEM, as well as, some of the applications involved in the DEEP-EST project.

3.5.5. FEM/BEM based domain decomposition solvers

The aim of this project is to extend the domain decomposition library ESPRESO to support highly scalable solution of problems in complex domain using finite/boundary element tearing and interconnecting (FETI/BETI) non overlapping domain decomposition method, thus enabling solution of large scale sound scattering and harmonic analysis problems. The goal is to provide a modern, modular and portable code written in C++, parallelised on all possible levels, and capable of utilizing the most powerful supercomputers. The code will be parallelised in a hybrid manner in shared and distributed memory and will benefit from the SIMD vector capabilities of modern CPUs. The parallelization in shared memory and vectorization will be done using the OpenMP pragmas. The computationally most demanding parts of the code will be accelerated using state of the art GP-GPUs.

3.5.6. Performance portable linear algebra

The goal of this project is to provide a modern and efficient distributed linear algebra package based on HPX, that can replace ScaLAPACK in scientific applications, and help the developers of scientific applications in the process of adopting modern, performance portable, and distributed linear algebra libraries. Since independent tasks can be executed easily in parallel, libraries built on task-based runtime can improve significantly the parallel efficiency of single function calls, and solve the fork-join mechanism issue of ScaLAPACK.

3.5.7. GHEX: Generic Halo-Exchange for Exascale

GHEX will strive to provide a clean and performance-portable implementation of halo exchange for the most important grid types independently of the particular applications characteristics. Priority will be given to grids used by the strategic scientific collaborators, among which MeteoSwiss and the COSMO weather and climate model, and the Rosseland Centre for Solar Physics (RoCS) and the DISPATCH modelling framework. GHEX will be used as a communication layer in GridTools - a set of open-source libraries to develop next generation weather and climate simulation applications in a common infrastructure.

3.5.8. LyNcs: Linear Algebra, Krylov methods, and multi-grid API and library support for the discovery of new physics

LyNcs will implement cutting-edge sparse linear solver algorithms, develop and prototype new Krylov and block Krylov solvers, and optimize existing parallel codes that implement a range of preconditioners for these solvers, such as multi-grid. At the lowest level, the project will further develop the efficient sparse matrix support software librsb, its APIs and adapter libraries. This library development and optimization activity will be accompanied by a coherent effort in designing, implementing, documenting, and maintaining an API, enabling various scientific user communities to build full-fledged scientific applications on top of these libraries.

3.5.9. ParSec: Parallel Adaptive Refinement for Simulations on Exascale Computers

ParSec brings together well-known HPC CFD practitioners with the aim of sharing best practices, and collaboratively modernize the AMR implementation of three leading-edge CFD community codes for the exploitation of future (pre) Exascale machines. These codes are Alya - the multiphysics solver from BSC; Nek5000 - the scalable high-order solver for computational fluid dynamics from KTH/UIUC; and the high order multiphysics solver, Argo, from Cenaero. The objectives of the project are: i) the definition of a common separation of concerns (SoC) for the AMR process - that will allow consistent performance comparisons and software components sharing amongst simulation platforms; ii) the cross-verification and analysis of the codes performance on (pre) Exascale architectures; iii) the modernization of the codes based on disruptive solutions coming from the previous analysis; and, finally, iv) deliver self-contained OS software components solving different steps of the AMR process accordingly to the SoC defined.

3.5.10. QuantEx: Efficient Quantum Circuit Simulation on Exascale Systems

This project will develop a platform consisting of modular quantum circuit simulation tools which use tensor network contractions methods, which are capable of running efficiently on heterogeneous compute platforms and scaling to exploit pre-Exascale and Exascale compute resources. Modern development practices and software design methodologies will be used along with hierarchical layers of abstraction to encapsulate complexity and enable these tools to be easily extended and integrated into users’ circuit simulation codes.

4. Center of Excellence

The Centres of Excellence (CoE) on HPC are EU funded projects that promote the use of upcoming exascale and extreme performance computing capabilities in areas where user communities in collaboration with other HPC stakeholders can develop or scale up existing parallel codes towards exascale and extreme scaling performance. They address the fragmentation of activities for excellence in HPC applications, and fostering the widening of the use of HPC codes in the EU, by establishing a focal point for the consulting skills and training. More details and a general overview of all CoEs can also be found within the latest European HPC Handbook [3] .

The following sections should provide an overview about all established Centres of Excellence (from the second funding round, which took place in 2018) and list the applications, libraries or tools which are used, optimized or developed as part of the individual projects, to provide an easy accessible overview which CoE could provide support, training or optimization for a certain application.

More information: https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/infraedi-02-2018

4.1. BioExcel - 2 (Centre of Excellence for Computational Biomolecular Research)

BioExcel provides support to academic and industrial researchers in the use of HPC for biomolecular research. The centre improves the most popular biomolecular software packages and provides HPC expertise among the biomolecular research communities through consultancy and training.

https://bioexcel.eu/

- Applications:
  - GROMACS: https://bioexcel.eu/software/gromacs/
  - HADDOCK: https://bioexcel.eu/software/haddock/
  - CP2K: https://bioexcel.eu/software/cp2k/
  - PMX: https://bioexcel.eu/software/pmx/
  - CPMD: https://bioexcel.eu/software/cpmd/
- Workflow systems: https://bioexcel.eu/software/workflows/

4.2. Cheese (Center of Excellence In Solid Earth)

Cheese establishes a Center of Excellence in the domain of Solid Earth (SE) targeting the preparation of 10 Community flagship European codes for the upcoming pre-Exascale and Exascale supercomputer and addresses various scientific, technical and socio-economic Exascale Computational Challenges. They plan to develop 12 Pilot Demonstrators and enable services oriented to society on critical aspects of geohazards like hazard assessment, urgent computing, and early warning forecast.

https://cheese-coe.eu/

- Applications: https://cheese-coe.eu/results/flagship-codes
  - ExaHyPE
  - Salvus
  - SeisSol

- SPECFEM3D
- PARODY_PDAF
- XSHELLS
- ASHEE
- FALL3D
- T-HySEA
- L-HySEA

- Workflow systems: https://bioexcel.eu/software/workflows/

4.3. CompBioMed 2 (CoE in Computational Biomedicine)

CompBioMed is a Centre of Excellence focused on the use and development of computational methods for biomedical applications. They have users within academia, industry and clinical environments and are working to train more people in the use of their products and methods.

https://www.compbiomed.eu/

- Applications:
  - Alya: https://www.compbiomed.eu/services/software-hub/compbiomed-software-alya/
  - HemeLB: https://www.compbiomed.eu/services/software-hub/compbiomed-software-hemelb/
  - HemoCell: https://www.compbiomed.eu/services/software-hub/compbiomed-software-hemocell/
  - OpenBf: https://www.compbiomed.eu/services/software-hub/compbiomed-software-openbf/
  - Palabos - Flow Diverter: https://www.compbiomed.eu/services/software-hub/compbiomed-software-palabos/
  - PolNet: https://www.compbiomed.eu/services/software-hub/compbiomed-software-polnet/
  - InSilicoMRI: https://www.compbiomed.eu/services/software-hub/compbiomed-software-insilicomri/
  - Living Heart Human Model: https://www.compbiomed.eu/services/software-hub/compbiomed-software-living-heart-human-model/
  - BAC: https://www.compbiomed.eu/services/software-hub/compbiomed-software-bac/
  - HTMD: https://www.compbiomed.eu/services/software-hub/compbiomed-software-htmd/
  - Playmolecule: https://www.compbiomed.eu/services/software-hub/compbiomed-software-playmolecule/
  - VisualGec: https://www.compbiomed.eu/services/software-hub/compbiomed-software-visual-gec/
  - HTBAC: https://www.compbiomed.eu/services/software-hub/compbiomed-software-htbac/
  - Virtual Assay: https://www.compbiomed.eu/services/software-hub/compbiomed-software-virtual-assay/
  - CT2S: https://www.compbiomed.eu/services/software-hub/compbiomed-software-ct2s/
  - Insigneo Bone Tissue Suit: https://www.compbiomed.eu/services/software-hub/compbiomed-software-insigneo-bone-tissue-suit/

- Palabos - Vertebroplasty: https://www.compbiomed.eu/services/software-hub/compbiomed-software-palabos/

4.4. EoCoE II (Energy Oriented Center of Excellence)

EoCoE applies cutting-edge computational methods in its mission to accelerate the transition to the production, storage and management of clean, decarbonized energy. EoCoE focuses on five scientific Exascale challenges in the low-carbon sectors of energy: Meteorology, Materials, Water, Wind and Fusion.

https://www.eocoe.eu

- Meteorology for Energy: https://www.eocoe.eu/meteorology-for-energy/
  - ESIAS-Chem
  - ESIAS-Meteo
  - EURAD-IM
  - WPMS
  - SPS
- Materials for Energy: https://www.eocoe.eu/materials-for-energy/
  - Metalwalls/QMCPack
  - PVnegf
  - KMC/DMC
- Water for Energy: https://www.eocoe.eu/water-for-energy/
  - ParFlow
  - SHEMAT-Suite
  - ExaTerr
- Fusion for Energy: https://www.eocoe.eu/fusion-for-energy/
  - Gysela
- Wind for Energy: https://www.eocoe.eu/wind-for-energy/
  - Alya
  - waLBerla
- Programming Models: https://www.eocoe.eu/technical-challenges/programming-models/
  - LIKWID
  - Scalasca
  - JUBE
  - Kokkos
  - P4est
- Scalable Solvers: https://www.eocoe.eu/technical-challenges/scalable-solvers/

  - PSBLAS
  - MLD2P4
  - MUMPS
  - AGMG
  - Chameleon
  - MAPHYS
  - PaStiX
- I/O & Data Flow: https://www.eocoe.eu/technical-challenges/i-o-data-flow/
  - PDI
  - SIONlib
  - FTI
  - IME
  - Sensei
- Ensemble Runs: https://www.eocoe.eu/technical-challenges/ensemble-runs/
  - Melissa

4.5. ESiWACE2 (Centre of Excellence in Simulation of Weather and Climate in Europe)

ESiWACE2 links, organizes and enhances Europe's excellence in weather and climate modelling. The project focuses on further exploring suitable innovative technologies to improve the model performance, the development of processing tools for a more efficient data in- and output and on providing enhanced support and training of the community.

http://www.esiwace.eu/

- Applications:
  - OpenIFS program: https://www.esiwace.eu/services/support/sup_OpenIFS
  - NEMO: https://www.esiwace.eu/services/support/sup_NEMO
  - OASIS3-MCT: https://www.esiwace.eu/services/support/sup_OASIS
- Libraries:
  - Cycl: https://www.esiwace.eu/services/support/overview
  - XIOS: https://www.esiwace.eu/services/support/copy_of_overview

4.6. EXELLERAT

The EXCELLERAT project bundles knowledge on HPC in engineering and offers service solutions for industry and research. One focus of EXCELLERAT lies on engineering solutions for aeronautics and the automotive sector.

https://www.excellerat.eu

4.7. FocusCoe

FocusCoe supports all other EU HPC CoEs to more effectively fulfil their role within the EU HPC ecosystem and initiative: ensuring that extreme scale applications result in tangible benefits for addressing scientific, industrial or societal challenges.

http://www.focus-coe.eu

4.8. HiDALGO (HPC and Big Data Technologies for Global Systems)

HiDALGO develops novel methods, algorithms and software for HPC and HPDA to accurately model and simulate the complex processes, which arise in connection with major global challenges. The CoE enables highly accurate simulations, data analytics and data visualisation and also provides knowledge on how to integrate the various workflows as well as the corresponding data.

https://hidalgo-project.eu

4.9. MaX 2 (Materials design at Exascale)

MaX focuses on creating an ecosystem of capabilities, software applications and data workflows and analysis on HPC-oriented material simulations, designed for present and future HPC architectures. A major effort is on providing training and services for the broader HPC industrial and academic community.

http://www.max-centre.eu/

- Applications:
  - Quantum ESPRESSO: http://www.max-centre.eu/codes-max/quantum-espresso
  - SIESTA: http://www.max-centre.eu/codes-max/siesta
  - Yambo: http://www.max-centre.eu/codes-max/yambo
  - Fleur: http://www.max-centre.eu/codes-max/fleur
  - CP2K: http://www.max-centre.eu/codes-max/cp2k
  - BigDFT: http://www.max-centre.eu/codes-max/bigdft
  - AiiDA: http://www.max-centre.eu/codes-max/aiida
- Libraries:
  - CheSS: http://www.max-centre.eu/software/product/libraries#1
  - FFTXLib & LAXLib: http://www.max-centre.eu/software/product/libraries#2
  - Sirius: http://www.max-centre.eu/software/product/libraries#3

4.10. POP 2 (Performance Optimisation and Productivity)

The Performance Optimisation and Productivity Centre of Excellence in Computing Applications provides performance optimisation and productivity services for academic and industrial code(s) in all domains. They offer a portfolio of services designed to help their users optimise parallel software and understand performance issues.

https://pop-coe.eu/

- Tools: https://pop-coe.eu/partners/tools
  - Extrae
  - Paraver
  - Dimemas
  - Scalasca
  - Cube
  - Score-P
  - Extra-P
  - Vampir
  - TAU
  - MAQAO
  - SimGrid
  - Darshan

5. National Support Strategies

Beside the different European HPC support activities, different countries also run national porting and code-optimization programs as well. These programs are often connected to national calls for computing time. The following section will be extended with each new update of the guide, but can only list certain example programs of the different European countries.

5.1. SiVeGCS

SiVeGCS is a German national project, whose full title sounds "Coordination and Securing the further Availability of GCS' supercomputing resources in the framework of the national HPC infrastructure" [2] .

GCS stands here for Gauss Centre for Supercomputing (GCS) [7] – a consortium that combines three biggest German national supercomputing centres: HLRS (High Performance Computing Center in Stuttgart), JSC (Jülich Supercomputing Centre) and LRZ (Leibniz Supercomputing Centre in Munich) into Germany's Tier-0 PetaFLOP-range supercomputing infrastructure.

SiVeGCS is a common initiative of the German national (by the Federal Ministry of Education and Reserach (BMBF) [8]) and regional (states of Baden-Württemberg, North Rhine-Westphalia and Bavaria) governments that aims to establish a platform for funding future German supercomputing systems in the "exascale decade", during the time frame until 2025. With the SiVeGCS support, the centers obtain the ability to update their existing infrastructures and install new state-of-the-art supercomputing systems on their premises as well as provide the necessary operative support for them, addressing the major challenges of supercomputing application across a wide spectrum of disciplines, including engineering, bioscience, geoscience and other scientific and industrial application domains. While doing this, a special attention is put to the users – a rich set of professional training activities in the form of course (both online and physical), workshops, etc. are offered.

More information: https://www.gauss-centre.eu/for-users/user-services-and-support/

5.2. The ARCHER eCSE Programme

5.2.1. Overview

The present UK National Supercomputing Service (ARCHER) started in 2013, running through to 2020. The Service is based around a Cray XC30 supercomputer and involves the Engineering and Physical Sciences Research Council (EPSRC), the Natural Environment Research Council (NERC), EPCC, Cray Inc. and The University of Edinburgh. The Computational Science and Engineering (CSE) support on ARCHER is provided by EPCC and as part of this support an "embedded CSE" (eCSE) programme is offered. This provides funding to the ARCHER user community to develop software in a sustainable manner to run on ARCHER and on future Tier-1 services. Throughout the majority of the ARCHER Service, calls were run at 4-monthly intervals receiving a total of 222 proposals, 100 of which were selecting for funding. The scientific applications utilised on ARCHER are key to the service's on-going success, with the eCSE programme ensuring these applications are sustained for ARCHER and future Tier-1 services such as ARCHER2. Later in 2020, the ARCHER2 service will come online with a similar eCSE programme funded as part of this service.

The eCSE funding is provided to enable the employment of a Research Software Engineer (RSE) to work specifically on the relevant software to enable new features or improve the performance of the code, with an emphasis on preparing the code for existing or new Tier-1 services. The funding is made available for staff located at the institution of the PI, third parties, or can include staff from the centralised CSE support team or a mixture of the above. The aim is to embed RSE's in HPC groups across the UK, ensuring RSEs across the UK can benefit from the programme and that the HPC community can access skills in the locations that need them. Staff employed on an eCSE project are expected to play a full part in ARCHER related activities.

EPCC ran the eCSE programme as a "not for profit" programme. While there was a requirement to provide a minimum of 14 FTEs of effort over the first 5 years of service, all the money was allocated to projects, resulting in a further 32 months of funded effort for the community.

[2]In German: Koordination und **Si** cherstellung der weiteren **Ve** rfügbarkeit der Supercomputing-Ressourcen des **Gauss Centre for Supercomputing (GCS)** im Rahmen der nationalen Höchstleistungsrechner-Infrastruktur.

5.2.2. Application and review processes

Calls are published via the ARCHER website [9] and publicised via appropriate mailing lists within the UK HPC community. Applicants apply via the ARCHER SAFE [10] which they must first sign up to. ARCHER users already have a log in on the ARCHER SAFE as this is the online service management system used to administer ARCHER accounts and resources as well as running the ARCHER helpdesk. The SAFE stores all personal details securely in accordance with GDPR rules. EPCC holds ISO 27001 Information Security and ISO 9001 Quality Management certifications which cover all HPC services. Once the call has closed, technical evaluations are carried out by members of the ARCHER CSE team except where proposals include members of staff from EPCC or the University of Edinburgh. Best practice then determines that such proposals are evaluated by a dedicated team of external technical experts from Cray, The Met Office, STFC and other institutions.

During the technical evaluation, short clarification questions for candidates may arise. These are then put to candidates who are all given exactly 1 week to respond (if they wish to) or may withdraw their application if they decide this is the best course of action.

A panel of experts then review the complete set of proposals that are going forward to the final stage in the review process. Best practice dictates that this panel is independent of the ARCHER CSE team and is made up of a set of experts from the UK and beyond whose expertise spans the breadth and range of scientific disciplines of the proposals themselves. We follow a strict conflict of interest procedure for panel members which is based on that used by the UK research councils. The eCSE panel chair is non-UK based which further helps to reduce conflicts of interest.

Finally, a panel meeting takes place to rank the proposals and decide the quality threshold to determine which proposals are funded. The number of projects funded at each call is not determined in advance other than taking into account the overall effort available for the programme. This means that the panel are free to select projects based on the quality of each proposal independently. Members of the ARCHER eCSE team co-ordinate the panel meetings themselves (arranging venues, etc.) and are on hand to give ARCHER-specific advice during the panel meeting, but do not take part in the decision making. On occasions proposals may be awarded conditionally or may be partially funded if the panel feel this is appropriate.

During the ARCHER service EPCC introduced an opportunity for early career researchers to act as observers during the panel meetings. The aim here was to give early career researchers a better insight into this competitive selection process to assist in the preparation of their own funding proposals. Any such observer present follows the same conflicts of interest and confidentiality guidelines as regular panel members. Over the course of the service 10 such early career researchers attended our panel meetings, a practice we expect to continue in the future.

Following on from the panel meeting, applicants are informed of the decisions made and given feedback on their proposal as well as being told of any conditions that come with an award. Unsuccessful applicants are able to apply for a future call once the feedback provided has been addressed.

Throughout the process, guidance documents are provided for technical evaluators and panel members explaining the procedures, criteria for reviewing, and conflict of interest and confidentiality guidelines.

5.2.3. Project Outcomes

On conclusion of an eCSE project, each project is required to produce a final report. This consists of a short summary of the project's outcomes, a description of the potential scientific benefit and impact of the work and a longer freeform technical report. These are made public via the ARCHER website with each project having a dedicated area to which publications and other links to relevant material (e.g. code repositories, documentation) can also be noted. We encourage the authors of the reports to obtain a DOI for the technical report and we can assist with this. Many PIs find the technical report useful for citing in later publications.

With the final report there is a section on benefits realisation where project members can note improvements in code performance, usability or uptake in the code's usage. Using data from the ARCHER service on the overall use of a code, cost of core hours, etc. the benefit of each project can be turned into a financial figure. A summary of this is given here: https://www.archer.ac.uk/documentation/white-papers/eCSE-Impact-Paper-2018/Benefits%20of%20the%20ARCHER%20eCSE%20Programme%20Public.pdf

The overall cost of the eCSE programme was around £7M. The benefits reported up to end 2019 are almost £35M, representing almost 5 times return on investment.

5.3. SNIC

The Swedish National Infrastructure for Computing (SNIC-UU), hosted by Uppsala University, is a national research infrastructure with a mission to provide generally available national services for large-scale computing and storage/management of active data sets, aiming at supporting all areas of research in need of access to such resources. SNIC-UU also provides a coordinated effort on advanced user support and training, facilitating the efficient use of the national and relevant international services by existing and new research communities and adapted to the local needs of the consortium members.

SNIC-UU is funded by the Swedish Research Council (Vetenskapsrådet) and by the SNIC consortium, which consists of ten university partners: Umeå University, the Swedish University of Agricultural Sciences, Uppsala University, KTH Royal Institute of Technology, Stockholm University, Karolinska Institutet, Linköping University, Chalmers University of Technology, CTH, Gothenburg University and Lund University. The SNIC-UU consortium is a Joint Research Unit (JRU) which participates as the national node in international e-infrastructure projects.

SNIC-UU is the Swedish national research infrastructure that is responsible to ensure that the Swedish HPC infrastructure is adequately integrated into the European HPC ecosystem. HPC infrastructure is provided by university partners within SNIC-UU. All ten SNIC-UU partners have many years of experience in collaborating with each other related to HPC, grid, large-scale storage and cloud. All ten SNIC-UU partners have expertise in adopting and maturing software systems for leading edge technologies, systems management tools development, including accounting systems for distributed environments, user training and support, performance tools and application development.

More information: https://www.snic.se/

5.4. SeRC

The Swedish e-Science Research Centre (SeRC) is an environment within the strategic research area (SRA) of e-Science, funded by the Swedish government Strategic Research Area Initiative and based on a collaboration between four universities: Kungliga Tekniska högskolan (KTH), Stockholms universitet (SU), Karolinska institutet (KI) and Linköpings universitet (LiU). It was founded in 2010 resulting from the Swedish Government Bill on Research Policy. In this bill, a total of 24 different strategic research areas were defined, of which e-Science is one.

The mission statement of SeRC is to develop state-of-the-art e-Science tools and provide an e-infrastructure support to existing and emerging e-Science communities to achieve scientific breakthroughs in Sweden. In its most basic form, the concept of e-Science is simply the notion of using digital information and the processing of this information to gain new scientific insights. This means that much of what we could call e-Science can be conducted using standard solutions both in terms of hardware and software.

The single largest overall successes of SeRC is that it helps application researchers and computational experts to find each other and speak with one voice. Large-scale data analysis and simulation is introduced in areas where it has not been used before, the strongest computational groups in the country collaborate internationally instead of competing nationally.

SeRC also has extensive collaborations with industry and governmental agencies and thereby generates direct impact in society and contributes to societal development and commercial competitiveness. SeRC has several industrial representatives on its advisory board, and we are delighted with their interest and engagement in learning how computing can change what they do, for instance at our annual meetings.

Computing infrastructures that are part of SeRC and SNIC are PDC and NSC. SeRC has also a visualization studio at KTH called VICas its partner. VIC is funded by a grant from the Knut and Alice Wallenberg’s foundation, together with Norrköping’s Visualisation Center and Linköping University.

Visualization Cente C at Norköpping is another partner. The center hosts a large-scale arena for public visits and events including media labs, interactive exhibitions and an immersive 3D fulldome theatre. The center is a result of a close collaboration between the Municipality of Norrköping, Linköping university, Norrköping Science Park and Interactive Institute Swedish ICT.

More information: https://e-science.se/

5.5. eSSENCE

eSSENCE is an e-Science collaboration formed by the universities in Uppsala, Lund and Umeå. The eSSENCE consortium is part of a long-term Government effort on Strategic Research Areas (Strategiska forskningsområden, SFO in Swedish) initiated in the research bill from 2008.

The mission of eSSENCE is to provide an internationally top-level research environment in that it integrates e-Science research in all areas of science, as well as in education and in collaboration with society.

The eSSENCE research strategy is to:

- increase the quality and progress of digital research by providing new and more efficient models, methods, algorithms and software in computational and data-intensive application areas;
- launch new and more effective methods for utilizing e-Infrastructure;
- create a national hub for collaboration between partner universities, between different research areas and disciplines, and between society, business and academia; and
- strengthen the e-Science aspects in education at different levels.

Computer infrastructures that are both part of SNIC and part of eSSENCE: HPC2N, LUNARC and UPPMAX.

More information: http://essenceofescience.se

Further documentation

Websites, forums, webinars

[1] *EuroHPC Joint Undertaking https://eurohpc-ju.europa.eu .*

[2] *PRACE Preparatory Access https://prace-ri.eu/hpc-access/preparatory-access/ .*

[3] *European HPC Handbook 2019, ETP4HPC https://www.etp4hpc.eu/european-hpc-handbook.html .*

[4] *PRACE White Papers https://prace-ri.eu/training-support/technical-documentation/white-papers/ .*

[5] *PRACE SHAPE Program https://prace-ri.eu/prace-for-industry/shape-access-for-smes/ .*

[6] *PRACE Digest https://prace-ri.eu/news-media/publications/prace-digest/ .*

[7] *Gauss Centre for Supercomputing http://www.gauss-centre.eu .*

[8] *Federal Ministry of Education and Research, Germany https://www.bmbf.de/en .*

[9] *ARCHER website http://www.archer.ac.uk/ .*

[10] *ARCHER SAFE guide https://www.archer.ac.uk/documentation/safe-guide/ .*

[11] *National Computing Center for Higher Education https://www.cines.fr/en/ .*

Manuals, papers

[12] *Bugli, M., Iapichino, L., Baruffa, F.: ECHO-3DHPC: Advance the performance of astrophysics simulations. 2018, Intel Parallel Universe 34, 49. arXiv: 1810.04597; Cielo, S., Iapichino, L., Baruffa, F., Bugli, M., Federrath, C.: Honing and proofing Astrophysical codes on the road to Exascale. Experiences from code modernization on many-core systems. 2020, To be published in Future Generation of Computer Systems (FGCS), Special Issue on "On The Road to Exascale II: Advances in High Performance Computing and Simulations". ArXiv: 2002.08161 .*