E-Infrastructures
H2020- INFRAEDI-2018-2020

INFRAEDI-01-2018: Pan-European High Performance Computing infrastructure and services (PRACE)

PRACE-6IP

PRACE Sixth Implementation Phase Project

Grant Agreement Number: INFRAEDI-823767

D7.3

Best Practice Guides for New and Emerging Architectures
Final

Version: 1.0
Author(s): Hayk Shoukourian, BADW-LRZ; Ole W. Saastad, UiO; Sebastian Lührs, JUELICH; Jorge Barbosa, UnivPorto; João Bispo, UnivPorto
Date: 18.10.2021

Project and Deliverable Information Sheet

PRACE Project	Project Ref. №: INFRAEDI-823767	
	Project Title: PRACE Sixth Implementation Phase Project	
	Project Web Site: http://www.prace-ri.eu/about/ip-projects	
	Deliverable ID: D7.3	
	Deliverable Nature: Report	
	Dissemination Level: PU*	Contractual Date of Delivery: 31 / October / 2021
		Actual Date of Delivery: 29 / October / 2021
	EC Project Officer: Leonardo Flores Añover	

* - The dissemination level is indicated as follows: **PU** – Public, **CO** – Confidential, only for members of the consortium (including the Commission Services) **CL** – Classified, as referred to in Commission Decision 2005/444/EC.

Document Control Sheet

Document	Title: Best Practice Guides for New and Emerging Architectures	
	ID: D7.3	
	Version: 1.0	Status: *Final*
	Available at: http://www.prace-ri.eu/about/ip-projects	
	Software Tool: Microsoft Word 2016	
	File(s): D7.3.docx	
Authorship	Written by:	Hayk Shoukourian, BADW-LRZ; Ole W. Saastad, UiO; Sebastian Lührs, JUELICH; Jorge Barbosa, UnivPorto; João Bispo, UnivPorto
	Contributors:	
	Reviewed by:	Lara Querciagrossa, CINECA Dirk Brömmel, JUELICH
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	07/June/2021	Draft	Initial draft and setup of the document structure Early complete draft of executive summary, introduction, chapter 2, and conclusion. Skeleton for chapter 3 and subchapters added. Draft for subchapter 3.4 regarding LUMI BPG is added.
0.2	26/June/2021	Draft	Added Application porting and code-optimization

			activities for European HPC systems section.
0.3	30/June/2021	Draft	Added summary regarding the BPG on “Modern Accelerators”.
0.4	01/October/2021	Draft	
0.5	01/October/2021	Draft	Layout updates
0.6	07/October/2021	Draft	1. Review comments
0.9	11/October/2021	Draft	Integration of reviewer remarks
1.0	18/October/2021	Final	2. Review comments

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Preparatory Access, SHAPE, HLST

Table of Contents

List of Figures

References and Applicable Documents

[1] https://eurohpc-ju.europa.eu
[2] https://prace-ri.eu/hpc-access/preparatory-access/
[3] https://prace-ri.eu/hpc-access/shape-access/
[4] https://prace-ri.eu/hpc-access/deci-access/deci-access-information-for-applicants/
[5] https://www.lumi-supercomputer.eu/about-lumi/

[6] https://www.lumi-supercomputer.eu/lumi-pilot-user-projects/
[7] https://docbook.org/
[8] https://bscw.zam.kfa-juelich.de/
[9] https://www.scientific-computing.com/
[10] https://prace-ri.eu/training-support/best-practice-guides/
[11] https://prace-ri.eu/training-support/best-practice-guides/application-porting-and-code-optimization-activities-for-european-hpc-systems/
[12] https://prace-ri.eu/training-support/best-practice-guides/modern-processors/
[13] https://prace-ri.eu/training-support/best-practice-guides/modern-accelerators/

List of Acronyms and Abbreviations

aisbl	Association International Sans But Lucratif (legal form of the PRACE-RI)
ALCF	Argonne Leadership Computing Facility
BCO	Benchmark Code Owner
BPG	Best Practice Guide
CFM	Cubic Feet per Minute
CMU	CPU Memory Unit
CoE	Centre of Excellence
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture (NVIDIA)
CVC	Calibrated Vectored Cooling
DARPA	Defense Advanced Research Projects Agency
DEISA	Distributed European Infrastructure for Supercomputing Applications EU project by leading national HPC centres
DoA	Description of Action (formerly known as DoW)
EAP	Early Access Platform
EC	European Commission
ECMWF	European Centre for Medium-Range Weather Forecasting
EESI	European Exascale Software Initiative
EFlop/s	Exa (= 10^{18}) Floating-point operations (usually in 64-bit) per second, also EF/s
EoI	Expression of Interest
EtS	Energy-to-Solution
ESFRI	European Strategy Forum on Research Infrastructures
EuroHPC JU	The European High-Performance Computing Joint Undertaking
EWHPC	European Workshops on HPC Infrastructures
FinFET	Fin Field-Effect Transistor
Flop/s	Floating Point Operations Per Second
GB	Giga (= 2^{30} ~ 10^9) Bytes (= 8 bits), also GByte
Gb/s	Giga (= 10^9) bits per second, also Gbit/s
GB/s	Giga (= 10^9) Bytes (= 8 bits) per second, also GByte/s
GÉANT	Collaboration between National Research and Education Networks to build a multi-gigabit pan-European network. The current EC-funded project as of 2015 is GN4.
GFlop/s	Giga (= 10^9) Floating-point operations (usually in 64-bit) per second, also GF/s
GHz	Giga (= 10^9) Hertz, frequency =10^9 periods or clock cycles per second
GPU	Graphic Processing Unit
HET	High Performance Computing in Europe Taskforce. Taskforce by representatives from the European HPC community to shape the European

	HPC Research Infrastructure. Produced the scientific case and valuable groundwork for the PRACE project.
HLST	High-Level Support Team
HPC	High-Performance Computing; Computing at a high-performance level at any given time; often used synonym with Supercomputing
HPE	Hewlett Packard Enterprise
HPL	High Performance LINPACK
HW	Hardware
ISC	International Supercomputing Conference; European equivalent to the US-based SCxx conference. Held annually in Germany.
kB	Kilo (= 2^{10} ~10^{3}) Bytes (= 8 bits), also kByte
LINPACK	Software library for Linear Algebra
MB	Management Board (highest decision-making body of the project)
MB	Mega (= 2^{20} ~ 10^{6}) Bytes (= 8 bits), also MByte
MB/s	Mega (= 10^{6}) Bytes (= 8 bits) per second, also MByte/s
MFlop/s	Mega (= 10^{6}) Floating-point operations (usually in 64-bit) per second, also MF/s
MOOC	Massively open online Course
MoU	Memorandum of Understanding.
MPI	Message Passing Interface
NDA	Non-Disclosure Agreement. Typically signed between vendors and customers working together on products prior to their general availability or announcement.
PA	Preparatory Access (to PRACE resources)
PATC	PRACE Advanced Training Centres
PB	Tera (= 2^{50} ~ 10^{15}) Bytes (= 8 bits), also Pbyte
PFlop/s	Peta (= 10^{15}) Floating-point operations (usually in 64-bit) per second, also PF/s
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE 2	The current phase of the PRACE Research Infrastructure following the initial five year period.
PRIDE	Project Information and Dissemination Event
RI	Research Infrastructure
SW	Software
TB	Technical Board (group of Work Package leaders)
TB	Tera (= 2^{40} ~ 10^{12}) Bytes (= 8 bits), also TByte
Tb/s	Tera (= 10^{12}) bits per second, also Tbit/s
TB/s	Tera (= 10^{12}) Bytes (= 8 bits) per second, also TByte/s
TCO	Total Cost of Ownership. Includes recurring costs (e.g. personnel, power, cooling, maintenance) in addition to the purchase cost.
TDP	Thermal Design Power
TFlop/s	Tera (= 10^{12}) Floating-point operations (usually in 64-bit) per second, also TF/s
TGG	The Green Grid
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context, the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1
UNICORE	Uniform Interface to Computing Resources. Grid software for seamless access to distributed resources.

List of Project Partner Acronyms

BADW-LRZ	Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften, Germany (3rd Party to GCS)
BILKENT	Bilkent University, Turkey (3rd Party to UHEM)
BSC	Barcelona Supercomputing Center - Centro Nacional de Supercomputacion, Spain
CaSToRC	The Computation-based Science and Technology Research Center (CaSToRC), The Cyprus Institute, Cyprus
CCSAS	Computing Centre of the Slovak Academy of Sciences, Slovakia
CEA	Commissariat à l'Energie Atomique et aux Energies Alternatives, France (3rd Party to GENCI)
CENAERO	Centre de Recherche en Aéronautique ASBL, Belgium (3rd Party to UANTWERPEN)
CESGA	Fundacion Publica Gallega Centro Tecnológico de Supercomputación de Galicia, Spain, (3rd Party to BSC)
CINECA	CINECA Consorzio Interuniversitario, Italy
CINES	Centre Informatique National de l'Enseignement Supérieur, France (3rd Party to GENCI)
CNRS	Centre National de la Recherche Scientifique, France (3rd Party to GENCI)
CSC	CSC Scientific Computing Ltd., Finland
CSIC	Spanish Council for Scientific Research (3rd Party to BSC)
CYFRONET	Academic Computing Centre CYFRONET AGH, Poland (3rd Party to PNSC)
DTU	Technical University of Denmark (3rd Party of UCPH)
EPCC	EPCC at The University of Edinburgh, UK
EUDAT	EUDAT OY
ETH Zurich (CSCS)	Eidgenössische Technische Hochschule Zürich – CSCS, Switzerland
GCS	Gauss Centre for Supercomputing e.V., Germany
GÉANT	GÉANT Vereniging
GENCI	Grand Equipement National de Calcul Intensif, France
GRNET	National Infrastructures for Research and Technology, Greece
ICREA	Catalan Institution for Research and Advanced Studies (3rd Party to BSC)
INRIA	Institut National de Recherche en Informatique et Automatique, France (3rd Party to GENCI)
IST-ID	Instituto Superior Técnico for Research and Development, Portugal (3rd Party to UC-LCA)
IT4I	Vysoka Skola Banska - Technicka Univerzita Ostrava, Czech Republic
IUCC	Machba - Inter University Computation Centre, Israel
JUELICH	Forschungszentrum Jülich GmbH, Germany
KIFÜ (NIIFI)	Governmental Information Technology Development Agency, Hungary
KTH	Royal Institute of Technology, Sweden (3rd Party to SNIC-UU)
KULEUVEN	Katholieke Universiteit Leuven, Belgium (3rd Party to UANTWERPEN)
LiU	Linkoping University, Sweden (3rd Party to SNIC-UU)
MPCDF	Max Planck Gesellschaft zur Förderung der Wissenschaften e.V., Germany (3rd Party to GCS)
NCSA	NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS, Bulgaria

NTNU	The Norwegian University of Science and Technology, Norway (3rd Party to SIGMA2)
NUI-Galway	National University of Ireland Galway, Ireland
PRACE	Partnership for Advanced Computing in Europe aisbl, Belgium
PSNC	Poznan Supercomputing and Networking Center, Poland
SDU	University of Southern Denmark (3rd Party to UCPH)
SIGMA2	UNINETT Sigma2 AS, Norway
SNIC-UU	Uppsala Universitet, Sweden
STFC	Science and Technology Facilities Council, UK (3rd Party to UEDIN)
SURF	SURF is the collaborative organisation for ICT in Dutch education and research
TASK	Politechnika Gdańska (3rd Party to PNSC)
TU Wien	Technische Universität Wien, Austria
UANTWERPEN	Universiteit Antwerpen, Belgium
UC-LCA	Universidade de Coimbra, Labotatório de Computação Avançada, Portugal
UCPH	Københavns Universitet, Denmark
UEDIN	The University of Edinburgh
UHEM	Istanbul Technical University, Ayazaga Campus, Turkey
UIBK	Universität Innsbruck, Austria (3rd Party to TU Wien)
UiO	University of Oslo, Norway (3rd Party to SIGMA2)
UL	UNIVERZA V LJUBLJANI, Slovenia
ULIEGE	Université de Liège; Belgium (3rd Party to UANTWERPEN)
U Luxembourg	University of Luxembourg
UM	Universidade do Minho, Portugal, (3rd Party to UC-LCA)
UmU	Umeå University, Sweden (3rd Party to SNIC-UU)
UnivEvora	Universidade de Évora, Portugal (3rd Party to UC-LCA)
UnivPorto	Universidade do Porto, Portugal (3rd Party to UC-LCA)
UPC	Universitat Politècnica de Catalunya, Spain (3rd Party to BSC)
USTUTT-HLRS	Universitaet Stuttgart – HLRS, Germany (3rd Party to GCS)
WCSS	Politechnika Wroclawska, Poland (3rd Party to PNSC)

Executive Summary

The Work Package 7 – "Applications Enabling and Support" provides excellent support to researchers of European academic and commercial sectors in achieving an effective and efficient usage of the available high-end IT infrastructures in Europe represented by PRACE Tier-0/Tier-1 as well as European High-Performance Computing Joint Undertaking (EuroHPC JU) petascale and pre-exascale systems. For doing so, this work package, in parallel to various application enabling services for preparatory access and industry, offers a series of Best Practice Guides (BPGs) covering a wide spectrum of topics, ranging from the discussion of architectural peculiarities of systems to contemporary tools, compilers, libraries, programming techniques and best practices assisting to application porting, profiling, and overall tuning. The target audience of these BPGs are both users and application support teams at supercomputing sites who are enabling and maintaining large-scale HPC applications.

The successful series of BPGs has been initiated already at the start of PRACE-1IP, and has been strongly continued throughout subsequent implementation phases of the PRACE project. In PRACE-6IP, this series has been further extended towards new systems and technologies. More specifically, the new guides in 6IP provide an update on modern processors and accelerators, discuss currently existing major large-scale application porting and code-optimization activities in Europe as well as provide further best practices for application migration and efficient use of EuroHPC pre-exascale system LUMI hosted at CSC, Finland[1].

This document outlines the details of the BPG development process and briefly describes the guides developed within PRACE-6IP.

[1] At the time of the writing, this last BPG was still under development given the deployment timelines of the LUMI system.

1. Introduction

Knowledge of architectural peculiarities of the target HPC system, information on the available system software stack, information regarding the available application support teams and their activities, etc. are all essential components for achieving the most efficient and effective usage of the available high-end infrastructure. In order to achieve this goal, since 2012, PRACE has continuously supported European HPC researchers by also providing high quality Best Practice Guides (BPGs) that aim to assist further the European supercomputing community in their day to day porting and tuning activities of large-scale applications. Throughout these years, PRACE has developed a solid set of BPGs that looked at different aspects necessary for successful HPC application porting and tuning: PRACE systems and the specifics of accompanying processor/accelerator technologies, modern interconnects, frameworks for deep learning in HPC, etc. to name a few. For PRACE-6IP, it was agreed to continue the successful series of these BPGs and decided to further develop BPGs on:

- **Application porting and code-optimization activities for European HPC systems**
  Provides an overview on European High-Performance Computing Joint Undertaking (EuroHPC JU [1]) activities, discusses the porting and optimization activities within the PRACE-6IP project (e.g. Preparatory access [2]*, SHAPE* [3]*, and DECI* [4]*), outlines the current generation of Centres of Excellence on HPC as well as describes various national scale initiative projects available in Europe.*
- **Modern Processors**
  Provides an update on a selection of recent processors, namely: ARM64 (Huawei/HiSilicon and Marvell) and x86-64 (AMD and Intel), discusses the concomitant programming models and development environments, as well as outlines some energy-efficiency aspects for the overall increase of user awareness.
- **Modern Accelerators**
  Provides brief description of hardware for a selection of relevant accelerator technologies currently deployed at some PRACE sites, namely: GPUs; Field-Programmable Gate Arrays (FPGAs); and vector processors, discusses their suitability for different HPC applications, describes the available programming models and the development environments, as well as outlines some hints and best practices for application tuning.
- **EuroHPC pre-exascale system LUMI** [5] **deployed at CSC, Finland**
  Still under development given the deployment timeline of the system. Exact dates are undisclosed. However, the system is expected to be in full production in 2022. At the time of writing no solid timeline is published.

The above-mentioned guides on “Modern Processors” and “Modern Accelerators”, while extensive, provide a hybrid approach of a field guide and a textbook. The aim of these two BPGs is not to replace any of the available in-depth textbooks and/or documentations of certain tools, but rather to provide a set of best practices that build upon the available literature and the expertise of authors involved to further ease the process of application porting and performance optimization. These guides showcase the usability and possibilities of further application tuning given a specific processor/accelerator technology, and do not provide any direct comparisons of different processors/accelerators involved. These guides provide a generic overview on various architectural peculiarities of current processor/accelerator technologies, discuss their accompanying programming models/environments and thus should be viewed as

complementary to the existing in-depth BPGs provided by hardware vendors that are typically specific to their own product.

The remainder of this document is organised as follows. Chapter 2 starts with the criteria description used for the selection of BPG topics in PRACE-6IP, discusses the organisational aspects and the review process, as well as outlines the external dissemination channels used to increase the visibility of the developed BPGs. Chapter 3 provides a brief description of the BPGs that are (and will be) developed within PRACE-6IP[2]. Chapter 4 provides a conclusion.

2. Approach to Best Practice Guides

2.1. Selection of Topics

In the DoA it was proposed that the team should maintain and extend the successful series of BPGs to new technologies and systems, with a special focus on new processors and accelerators as well as the accompanying memory technologies, along with new interconnects; and workflows for HPC job processing and data management. This led to the following suggestion of the initial topic list shared among partners during the kick-off meeting of the PRACE project:

- New processors/accelerators
  - Intel Skylake/Cascade Lake
  - AMD Zen2
  - Vector processors
  - Nvidia Volta/Pascal GPUs
  - AMD GPUs
  - Graphcore Intelligence Processing Unit (IPU)
- New PRACE Tier-0/Tier-1 systems
- European technologies suitable for future European exascale systems
- Application porting and code-optimization activities for European HPC systems
- Impacts of AI on HPC *(e.g. programming languages, virtualization/container technologies, etc.)*
- New Interconnects *(performance and power efficiency gains, scalability improvements, resiliency, etc.)*
- I/O for exascale *(file system, data movement, archival storage, etc.; requirements: HW/SW for exascale)*
- Modern approaches for resource and data management and scheduling strategies at exascale
- Memory Technologies *(e.g. MCDRAM, NVRAM, 3D Xpoint, etc.)*

This list was then continuously updated throughout PRACE-6IP given the feedback received by WP7 partners. The selection of exact topics was further refined based on the indicated interest and the expertise of partners involved using a voting mechanism. Since the kick-off meeting of the PRACE-6IP project back in May 2019, all partners involved were continuously asked to indicate their interests (and/or suggest new ones) across various topics suggested. This was made possible as via direct communication as well as via the help of PRACE wiki and alternative online document editing tools.

[2] Including the six-month time period covering the extension of PRACE-6IP.

2.2. Organization

As in previous phases of PRACE projects, also in 6IP, LRZ's gitlab was used as the BPG development repository. All partners involved were invited to this repository. Following the same style of the earlier PRACE-IP projects, Docbook was used to generate professional .pdf and .html versions of the guides. Key feature is having a single source format (.xml based) for a range of multiple fully cross-referenced output formats: .html, .pdf and more. For more information on Docbook XML please refer to [7].

Monthly teleconferences were continuously organized that involved discussions on pending issues, updates from partners involved, etc. Among the coverage of various organizational and communication issues, this series of task activity related teleconferences allowed to ensure that BPGs are published throughout the project and not only towards the end of PRACE-6IP and thus addressed one of the major comments relating to BPGs inherited from PRACE-5IP project review. The slide decks as well as accompanying meeting minutes and notes are constantly uploaded to the project internal collaboration platform, BSCW [8].

The whole BPG development activity was led by Hayk Shoukourian (BADW-LRZ) with the support of following lead authors:

- Sebastian Lührs (JUELICH): BPG on "*Application porting and code-optimization activities for European HPC systems*"
- Ole W. Saastad (UiO): BPG on "*Modern Processors*"
- Jorge Barbosa (UnivPorto) and João Bispo (UnivPorto): BPG on "*Modern Accelerators*"
- Ole W. Saastad (UiO): BPG regarding EuroHPC pre-exascale LUMI system hosted at CSC, Finland (currently under development).

2.3. Review and publication

While this was not the case in earlier PRACE-IP projects, starting from PRACE-6IP all BPGs were additionally reviewed by PRACE-PMO prior to publication. This has further improved the quality of the published guides.

2.4. Further dissemination

To facilitate further the outreach of the developed BPGs, also in PRACE-6IP, outward dissemination activities have been continued. For achieving this, the BPG team continuously stayed in a close contact with PRACE communication office and WP3 that supported the corresponding dissemination of the guides via Scientific Computing World (SCW) [9] magazine. The following bullet list provides information regarding the so far conducted disseminations of the developed guides, while Figure 1 and Figure 2 below illustrate the corresponding web and email-based outreach via SCW:

- Best Practice Guide on "*Application porting and code-optimization activities for European HPC systems*", please refer to: *https://www.scientific-computing.com/white-paper/application-porting-and-code-optimization-activities-european-hpc-systems*
- Best Practice Guide on "*Modern Processors*", please refer to: *https://www.scientific-computing.com/white-paper/prace-best-practice-guide-modern-processors*

- Best Practice Guide on "*Modern Accelerators*", please refer to: https://www.scientific-computing.com/white-paper/best-practice-guide-modern-accelerators

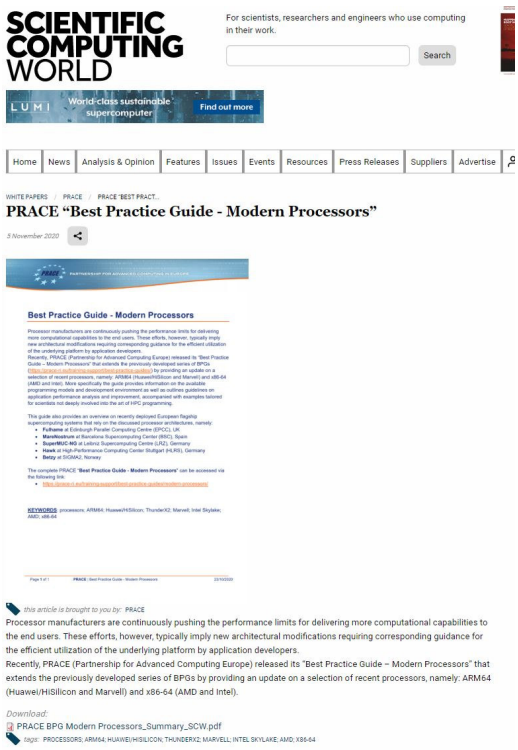

Figure 1: Snapshot illustrating the web dissemination conducted by SCW

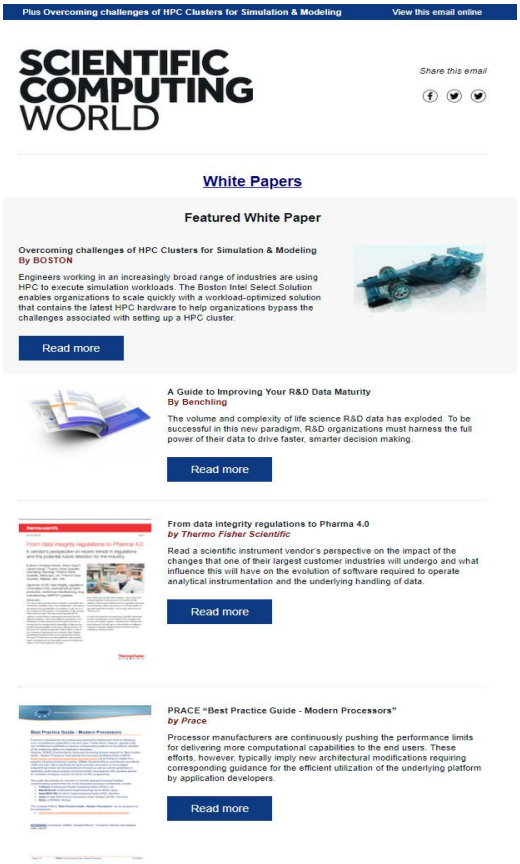

Figure 2: Snapshot illustrating the email-based dissemination conducted by SCW

3. Best Practice Guides

Given that BPGs are publicly available, this section only provides a brief overview on the BPGs that are (and will be) developed within PRACE-6IP. The complete list of all BPGs published so far, starting from 2012, can be found via [10] (see Figure 3).

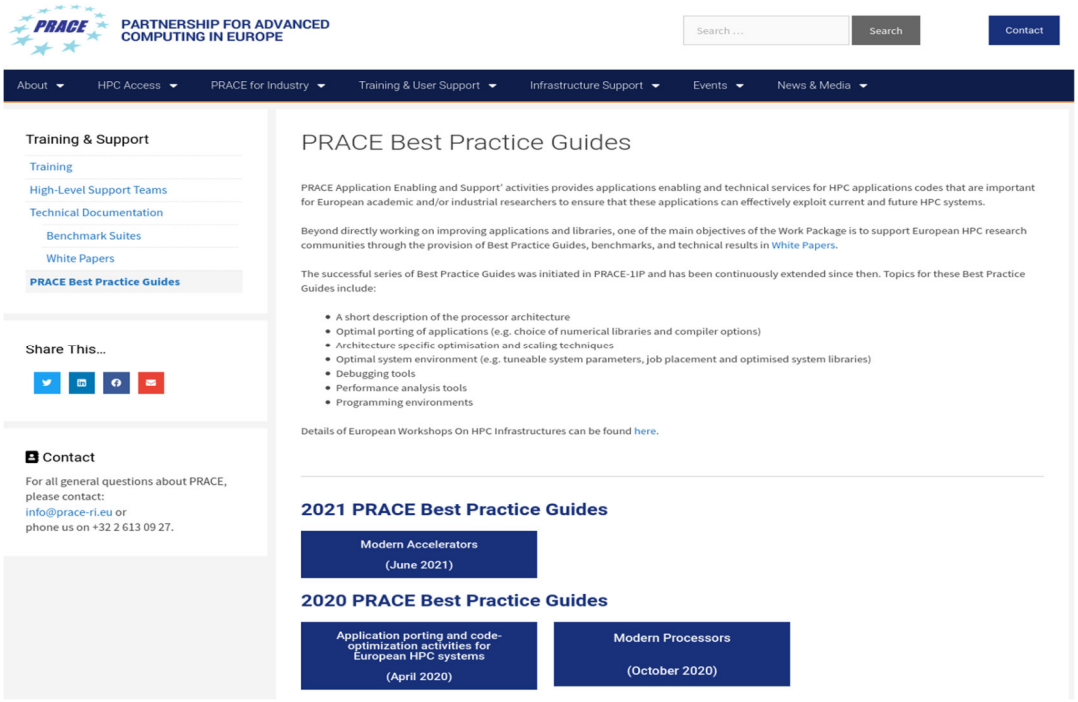

Figure 3: Snapshot illustrating PRACE web infrastructure for BPG dissemination

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

With the increasing number of European HPC systems on the road towards exascale computing, HPC applications face a diversity of different system types and technologies. 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.

Due to the necessary support demands to help with the optimisation work, various HPC support activities were established during the last years to offer general HPC knowledge as well as scientific area specific optimization strategies towards communities from research and industry.

This guide provides an overview about the larger application porting and code-optimization activities for European HPC systems currently ongoing. It does not cover specific technical details of certain applications 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 covers a range of activities in PRACE such as Preparatory access, SHAPE or DECI. In addition, it also contains an overview about the current generation of Centres of Excellence on HPC and the EuroHPC initiative. Finally, examples of national HPC support activities are given.

The complete BPG can be accessed via [11].

3.2. Best Practice Guide – Modern Processors

Efficient use of PRACE systems requires detailed knowledge of architecture specific factors influencing performance, compilers, tools and libraries. The main goal of this BPG is to investigate such issues, collect best practices on how to achieve good performance on the systems, and disseminate this knowledge to users.

This guide covers a range of the most used processors in an HPC environment, ARM, Intel & AMD and describes some experiences with the use of some common tools for these processors. Different processors have different characteristics, instructions, vector units, memory hierarchy etc.

The guide provides an overview of different architectures and covers the following tools: compilers, performance libraries, mathematical libraries, threading libraries (OpenMP), message passing libraries (MPI), debuggers and performance profilers, etc. Selected benchmarks comparing compilers and libraries have been performed. Benchmarks include processor compute performance, and memory bandwidth.

A short overview of some European systems covering the different architectures is also included.

An update of the guide involving further application tuning hints when using AMD processors with the Intel software tools was published on 05.05.2021. The complete BPG can be accessed via [12].

3.3. Best Practice Guide – Modern Accelerators

Hardware accelerators are special types of elements designed for boosting the performance of certain application regions requiring large amounts of numerical computations. Several factors contributed to broadening the use and furthering the adoption of these technologies in High-Performance Computing (HPC). One of such is the offered greater computational throughput as compared to stand-alone Central Processing Units (CPUs), which is driven by the highly parallel architectural design of accelerators.

The guide provides a generic overview on various accelerators and their accompanying programming models/environments and thus should be viewed as complementary to the existing in-depth documentation provided by hardware vendors that are typically specific to their own product. Nevertheless, it provides a set of best practices to further ease the process of application porting and possibilities of further application tuning given a specific accelerator technology.

The BPG starts with the description of the hardware for a selection of relevant accelerator technologies currently deployed at some PRACE sites, namely: GPUs, Field-Programmable Gate Arrays (FPGAs), and vector processors. This is followed by the discussion on suitability of these accelerators for different HPC applications.

Then it provides information on programming models (e.g. CUDA, SYCL, HIP, etc.) and development environments, as well as outlines some hints and best practices for application porting. Followed by strategies for application performance analysis and tuning, as well as a brief overview on the available debugging tools.

To finalize, the guide provides a brief information about various flagship and prototype HPC systems that are available at PRACE HPC sites and employ the discussed accelerator technologies.

The complete BPG can be accessed via [13].

3.4. Best Practice Guide – EuroHPC pre-exascale system LUMI

A subsequent BPG is planned to be regarding one of the three European High-Performance Computing Joint Undertaking (EuroHPC JU) [1] pre-exascale systems, namely: the LUMI system hosted at CSC, Finland [5]. According to the currently indicated timeline [6], the LUMI system is being opened to users in two phases:

- **Phase 1.** Provision of access to all partitions but LUMI-G, including the Early Access Platform (EAP) used as a migration system. This was foreseen for the late summer 2021 and scheduled for access by October 2021.
- **Phase 2.** Integration of LUMI-G partition and provision of full access, foreseen for the end of 2021.

Given the existing contractual PRACE-6IP timeline[3] together with the mentioned deployment timeline of the LUMI system, it was decided to direct the main focus of this next BPG towards the EAP. This platform should include a GPU configuration similar to what the real LUMI accelerated nodes will look like, but using the previous generation of AMD GPUs. EAP will also have all the relevant software stacks for programming the AMD GPUs.

Given that LUMI-G will have different GPUs and possibly other differences as well due to older software stack leading to possible issues, for instance, with managed memory, OpenMP offload, etc. it was further decided to dedicate the PRACE-6IP extension period[4] to update the BPG, including all peculiarities of the LUMI-G partition.

[3] The project was originally scheduled to end by December 2021.

[4] PRACE-6IP has been extended for 6 months, i.e. from 1st of January, 2022 to 30th of June, 2022.

4. Conclusion

The successful series of Best Practice Guides has been initiated in PRACE-1IP and has been continuously updated and extended since then.

Within PRACE-6IP the following three (one regular and two extended) Best Practice Guides have already been published on the PRACE website:

- "*Application porting and code-optimization activities for European HPC systems*" (regular), published in April, 2020.
- "*Modern Processors*" (extended), published in October 2020, updated in May 2021.
- "*Modern Accelerators*" (extended), published in July 2021.

Given the existing system deployment timelines, an additional publication of another regular Best Practice Guide regarding the Early Access Platform of the LUMI system (i.e. one of the three EuroHPC pre-exascale systems) is scheduled in the first half of 2022 during the extension of PRACE-6IP.