E-Infrastructures
H2020-EINFRA-2014-2015

EINFRA-4-2014: Pan-European High Performance Computing Infrastructure and Services

PRACE-4IP

PRACE Fourth Implementation Phase Project

Grant Agreement Number: EINFRA-653838

D4.2
Final Report on Training Activities

Final

Version:	1.0
Author(s):	Jussi Enkovaara, CSC, Ioannis Liabotis, GRNET, Simon Wong, ICHEC, Eli Shmueli, IUCC
Date:	22.04.2017

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: EINFRA-653838**	
	Project Title: PRACE Fourth Implementation Phase Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: < **D4.2**>	
	Deliverable Nature: < Report >	
	Dissemination Level: PU	**Contractual Date of Delivery:** 30 / April / 2017
		Actual Date of Delivery: 30 / April / 2017
	EC Project Officer: Leonardo Flores Añover	

* - The dissemination level are 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 2991/844/EC.

Document Control Sheet

Document	**Title: Final Report on Training Activities**	
	ID: D4.2	
	Version: <1.0>	**Status:** ***Final***
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2010	
	File(s): D4.2.docx	
Authorship	**Written by:**	Jussi Enkovaara, CSC, Ioannis Liabotis, GRNET, Simon Wong, ICHEC, Eli Shmueli, IUCC
	Reviewed by:	Leon Kos, UL Veronica Teodor, FZJ
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	07/Feb/2017	Draft	Outline
0.9	03/April/2017	Draft for internal review	
1.0	23/April/2017	Final version	

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Training, PATC, Seasonal School, MOOC, Centres of Excellence, CoE, International HPC Summer School

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

[1] PRACE-4IP D4.1: "*Interim PRACE Training Report*", http://www.prace-ri.eu/IMG/pdf/D4.1-4ip_v02.pdf

[2] PRACE-4IP D4.3: "*Assessment of the PRACE Advanced Training Centres*".

[3] PRACE-4IP D4.4. "*MOOC Pilot for HPC*"

List of Acronyms and Abbreviations

aisbl	Association International Sans But Lucratif (legal form of the PRACE-RI)
CoE	Center of Excellence
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture (NVIDIA)
DoA	Description of Action (formerly known as DoW)
EC	European Commission
EESI	European Exascale Software Initiative
EoI	Expression of Interest
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, i.e. DP) 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 European HPC community to shape the European HPC Research Infrastructure. Produced the scientific case and valuable groundwork for the PRACE project.
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
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
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, i.e. DP) 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
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE 2	The upcoming next phase of the PRACE Research Infrastructure following the initial five year period.
PRIDE	Project Information and Dissemination Event
PTC	PRACE Training Centres
RI	Research Infrastructure
TB	Technical Board (group of Work Package leaders)
TB	Tera (=10^{12}) Bytes (= 8 bits), also TByte
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, i.e. DP) per second, also TF/s
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

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 UYBHM)
BSC	Barcelona Supercomputing Center - Centro Nacional de Supercomputacion, Spain
CaSToRC	Computation-based Science and Technology Research Center, Cyprus
CCSAS	Computing Centre of the Slovak Academy of Sciences, Slovakia
CEA	Commissariat à l'Energie Atomique et aux Energies Alternatives, France (3 rd Party to GENCI)
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 (3 rd Party to GENCI)
CNRS	Centre National de la Recherche Scientifique, France (3 rd 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)
EPCC	EPCC at The University of Edinburgh, UK
ETHZurich (CSCS)	Eidgenössische Technische Hochschule Zürich – CSCS, Switzerland
FIS	FACULTY OF INFORMATION STUDIES, Slovenia (3rd Party to ULFME)
GCS	Gauss Centre for Supercomputing e.V.
GENCI	Grand Equipement National de Calcul Intensiv, France
GRNET	Greek Research and Technology Network, Greece
INRIA	Institut National de Recherche en Informatique et Automatique, France (3 rd Party to GENCI)
IST	Instituto Superior Técnico, Portugal (3rd Party to UC-LCA)
IT4Innovations	VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA, IT4Innovations National Supercomputing Center, Czech Republic
IUCC	INTER UNIVERSITY COMPUTATION CENTRE, Israel
JKU	Institut fuer Graphische und Parallele Datenverarbeitung der Johannes Kepler Universitaet Linz, Austria
JUELICH	Forschungszentrum Juelich GmbH, Germany
KIFÜ	Governmental Information Technology Development Agency, Hungary
KTH	Royal Institute of Technology, Sweden (3 rd Party to SNIC)
LiU	Linkoping University, Sweden (3 rd Party to SNIC)
NCSA	NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS, Bulgaria
NTNU	The Norwegian University of Science and Technology, Norway (3rd Party to SIGMA)

NUI-Galway	National University of Ireland Galway, Ireland
PRACE	Partnership for Advanced Computing in Europe aisbl, Belgium
PSNC	Poznan Supercomputing and Networking Center, Poland
RISCSW	RISC Software GmbH
RZG	Max Planck Gesellschaft zur Förderung der Wissenschaften e.V., Germany (3 rd Party to GCS)
SIGMA2	UNINETT Sigma2 AS, Norway
SNIC	Swedish National Infrastructure for Computing (within the Swedish Science Council), Sweden
STFC	Science and Technology Facilities Council, UK (3rd Party to EPSRC)
SURFsara	Dutch national high-performance computing and e-Science support center, part of the SURF cooperative, Netherlands
UC-LCA	Universidade de Coimbra, Labotatório de Computação Avançada, Portugal
UCPH	Københavns Universitet, Denmark
UHEM	Istanbul Technical University, Ayazaga Campus, Turkey
UiO	University of Oslo, Norway (3rd Party to SIGMA)
ULFME	Univerza v LjubljanI, Slovenia
UmU	Umea University, Sweden (3 rd Party to SNIC)
UnivEvora	Universidade de Évora, Portugal (3rd Party to UC-LCA)
UPC	Universitat Politècnica de Catalunya, Spain (3rd Party to BSC)
UPM/CeSViMa	Madrid Supercomputing and Visualization Center, Spain (3rd Party to BSC)
USTUTT-HLRS	Universitaet Stuttgart – HLRS, Germany (3rd Party to GCS)
WCNS	Politechnika Wroclawska, Poland (3rd party to PNSC)

Executive Summary

The Training Work Package (WP4) of the PRACE Fourth Implementation Phase (PRACE-4IP) project has continued the PRACE training programme by offering large variety of face-to-face training events as well as by developing on-line training services. During Feb 2015 – Apr 2017 PRACE-4IP has delivered 569 training days as part of 198 training events that have reached an audience of 4,492 participants. The training team has also started the conception of PRACE Training Centres (PTCs) for complementing the PRACE Advanced Training Centres (PATCs). Six Seasonal Schools have been organised within PRACE-4IP, and in collaboration with international partners the series of International HPC Summer Schools has continued to prosper. In order to better serve research communities such as the Centres of Excellence (CoEs), a new type of On-Demand training event has been introduced, and five such events have been organized. As a completely new training activity, PRACE has developed two Massively Open Online Courses (MOOCs). By mid-April 2017, 5,907 participants have enrolled in the two MOOCs. As a new online training activity PRACE has also piloted CodeVault, an open repository of high performance computing code samples. Overall, PRACE-4IP WP4 has delivered a vital service to foster the growth of computational skills in Europe and to strengthen the competitiveness of European science and industry in the field.

1 Introduction

One of the main aims of PRACE, since it has been initiated in 2008 by the PRACE Preparatory Phase project (PRACE-PP), has been to offer a training programme suitable for the needs of the European students and researchers. An objective of PRACE-4IP, in particular, is to build up European human resources skilled in HPC and HPC applications. This can be achieved by, among others, organising highly visible events, enhancing the state-of-the-art training provided by the PRACE Advanced Training Centres (PATCs), providing a better articulation with the offerings of the CoEs and the ETP4HPC, while improving online training.

In this document we describe the work done in the PRACE-4IP project, by the Training WP4, for implementing these objectives. Major part of the effort has focused on the face-to-face trainings, which are of the following types:

- PRACE Advanced Training Centres (PATC) Courses,
- PRACE Seasonal Schools,
- PRACE On-demand Events,
- International HPC Summer School.

In addition, as new activities WP4 has been developing two Massively Open On-line Courses (MOOCs) as well as an open repository of HPC code samples called CodeVault.

This document is organised as follows: Section 2discusses the operation and development of PATCs in PRACE-4IP including the work done in preparing the launch of PRACE Training Centres (PTCs). Section 3 presents the Seasonal Schools carried out during the project, section 4 On-Demand Events, and section 5 the International HPC Summer School. Section 0 discusses the experiences in piloting the MOOCs, Section 7 presents the CodeVault repository and finally section 8 provides conclusions. The annex provide reports from PRACE Training events.

2 PRACE Advanced Training Centres

The mission of the PRACE Advanced Training Centres (PATCs) is to serve as European hubs of advanced, high-quality training for researchers working in the computational sciences. The PATCs provide and coordinate training and education activities needed to achieve the best utilisation of the PRACE Research Infrastructure by the community. The PATCs promote a common PRACE brand, representing the whole PRACE community rather than only the hosting sites, and implement a jointly developed programme of courses, designed and coordinated by PRACE with advice from external experts.

The deliverable D4.1: "*Interim PRACE Training Report"* [1] contains a report on PATC activities during the first year of the PRACE-4IP project, including PATC courses provided as well as an overview of PRACE-4IP D4.3: "*Assessment of PRACE Advanced Training Centres"* [2]. The latter was a significant assessment of the PATCs, examining the strengths, weaknesses, opportunities and threats of the network. So while the PATCs have been extremely successful in delivering a world-class programme of HPC-related courses to over 1,500 students and researchers on a consistent basis annually, the level of participation from those based in non-PATC hosting countries has always been disproportionately low. Hence a major recommendation from the deliverable was to establish a framework where the PATC network concept can be expanded to include new "PRACE Training Centres" to complement the activities of the PATCs and improve the level of participation across Europe; this process has already begun in the PRACE-4IP project and discussed in more detail in Section 2.2. Other opportunities and recommendations are also discussed to facilitate increased flexibility in the implementation of the PATC programme, to carry out additional feedback/impact analysis, as well as to work with other communities for better dissemination and collaboration.

	Start-up	2012-13	2013-14	2014-15	2015-16	2016-17
Duration	Mar'12 - Jul'12	Aug'12 - Jul'13	Aug'13 - Jul'14	Aug'14 - Jul'15	Aug'15 - Jul'16	Aug'16 - Apr'17
Number of courses	19	71	81	77	73	79 (60)
Total duration (days)	56	204	233	219	203	**215 (170)**
Number of participants	511	1,547	1,682	1,786	1,561	1,178
Number of participant-days	1,715	4,702	5,187	5,384	4,583	3,476
Female (%)	-	12.9%	14.4%	16.3%	17.4%	19.3%
Non-academic (%)	-	9.9%	12.3%	15.6%	20.8%	20.5%
Non-host country (%)	-	20.6%	25.4%	29.5%	16.1%	13.4%
Non-PATC country (%)	-	13.8%	17.7%	19.9%	8.9%	12.0%
Feedback response rate (%)	**-**	**63%**	**64%**	**53%**	**52%**	**54%**
Average overall rating (0-10)	-	8.5	8.4	8.4	8.4	8.5

Table 1. Key statistics of PATC programmes, each comprised of courses delivered within a year aligned with the academic year, since the establishment of the PATCs in 2012. "Non-host country" indicates the proportion of participants affiliated with institutions from anywhere outside the country hosting the PATC course. "Non-PATC country" indicates the proportion of participants affiliated with institutions from non-PATC hosting countries. For the 2016-17 programme, statistics are based on partial implementation of the programme, with numbers in parentheses indicating the number of courses, participants and training days delivered so far out of the total numbers planned for the programme.

2.1 PATC Courses from February 2015 to April 2017

This subsection presents the PATC activities during the PRACE-4IP project period from February 2015 to April 2017, encompassing multiple PATC programmes (partial 2014-15 programme, entire 2015-16 programme, and partial 2016-17 programme). In order to maintain consistency with the prior training report [1] that examined the entire 2014-15 programme after its completion, although it overlapped with the PRACE-3IP project, we focus here on the 2015-16 programme (supported by the PRACE-4IP project in its entirety and completed in July 2016). However, statistics are shown for all PATC courses until the end of the PRACE-4IP project at the end of April 2017.

2.1.1 *Key statistics on numbers of participants and courses*

Table 1 shows the key output and statistics from the network of six PATCs since their establishment in 2012. Focusing on the 2015-16 programme, the statistics show that the level of participation compares well relative to previous years, where 1,561 participants attended 73 PATC courses. The feedback from these courses continue to be very positive, with an overall average rating of 8.4 out of 10 for the PATC courses during this period; one of the most consistent statistics over the years. Both the female representation at PATC courses and non-academic participation have increased slightly. However, there has been a considerable drop, relative to previous years, in the percentage of participants from non-host countries (16.1%) and non-PATC (8.9%). One explanation is that some of the prior PATC programmes, in order to widen geographical reach, hosted PATC courses in non-PATC hosting countries (a course in Stockholm during the 2013-14 programme, two courses in Dublin and Amsterdam during the 2014-15 programme), hence providing small boosts to the level of non-PATC participation. However, PATCs did not deliver courses in non-PATC countries in the 2015-16 programme. This can sometimes occur as course delivery in non-PATC sites often require collaboration with another project or partner who may not always have the necessary funding and resources.

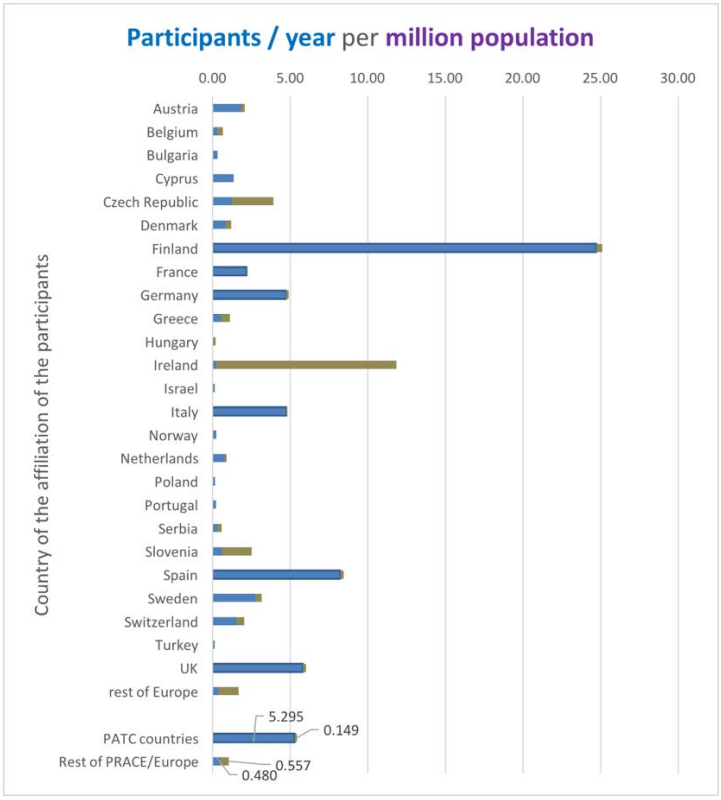

Figure 1: Geographic distribution of PATC course participants, normalised by national population sizes. Blue bars indicate participation in PATC courses, brown bars indicate participation in PRACE seasonal schools and other training events.

2.1.2 *Geographic distribution of the participants*

Consistent with past trends, the geographic distribution of the participants tend to be highly skewed towards PATC countries, i.e. low level of participation by those from outside PATC countries. Shown in Figure 1 is the geographic distribution of the 2015-16 PATC programme participants. Subsequent to normalisation with national population sizes, the level of participation from PATC-hosting countries is 10 times that of non-PATC countries. This has long been a documented issue for the PATCs to achieve wide geographic distribution of participants by themselves. One solution to expand such geographic reach was to establish further PRACE training centres, ensuring more PRACE courses in non-PATC countries, which are then covered in Section 2.2.

2.2 Conception of PRACE Training Centres

The PRACE-4IP project deliverable D4.3: "*Assessment of the PRACE Advanced Training Centres*" [2] carried out a SWOT analysis on the network of PATCs, and produced a list of key recommendations, including an argument to establish new PRACE Training Centres (PTCs) to expand the geographic reach of PRACE training, complementing the activities of the PATCs. While the financial support and formal selection/establishment of such PTCs had subsequently been included in the activities of the forthcoming PRACE-5IP project, preliminary work was carried out in the PRACE-4IP project to prepare for timely establishment of the PTCs, at an early stage of the PRACE-5IP project, so that the pilot PTC course programme would coincide with the PATC 2017-18 programme.

Much of the preparatory work involved meetings and discussions, involving WP4 members, PATC representatives and PRACE management (e.g. Board of Directors), to conceive the various responsibilities of the PTCs, and the process for their selection.

Essentially, PTCs will implement a collective training programme jointly coordinated by PRACE, where each PTC will focus on subject areas that may be of particular relevance or interest to industry and/or the research communities of the local and surrounding region. Each PTC is expected to deliver 2-3 training events annually, representing up to 9 training days. Finally all PTC courses are to be open to all researchers, free of charge (for participation) and preferably given in English. The responsibilities of each PTC includes planning, dissemination, delivery, and reporting of its courses, working in collaboration with other PTCs and PATCs under direction from PRACE.

A number of factors were also identified that would form the basis for the selection of PTCs. These criteria include:

- Ability to carry out the requested number of training events.
- Consistency and track record of training activities.
- Availability, range and expertise of local trainers.
- Justification for a tentative training programme to be proposed by each PTC.
- Suitability of the hosting site for attracting students from local and surrounding regions.

Having derived these conditions and selection criteria, it was anticipated that the call for applications for PTCs would take place in February 2017, followed by selection and establishment of four PTCs in April/May 2017 in the early stage of the PRACE-5IP project.

3 PRACE Seasonal Schools

The PRACE Seasonal Schools have been running since 2008 as part of the PRACE educational programme offering top-quality face-to-face training events organised around Europe, aiming to improve the skills necessary for the use of the PRACE ecosystem. The Seasonal School topics range from generic intermediate to advanced programming techniques to more specialised topical schools that e.g. focus on a specific topic, such as big data, or offer discipline specific parallel tracks.

Since 2012 Seasonal Schools have run in parallel with the PRACE Advanced Training Centres offering training opportunities mainly in countries where PATCs are not in operation. With the introduction of the PTCs in PRACE-5IP, the Seasonal School will also run in parallel to the PTC. To maximize PRACE training coverage around Europe, the countries selected to host Seasonal Schools should not be PATC or PTC hosting countries. In PRACE-4IP Seasonal Schools run also in parallel with the On-Demand Events offering a complete portfolio of training services distributed across the different countries of Europe members of the PRACE association. Among the aims of the Seasonal Schools, particularly within the scope of the PRACE-4IP project, is their extension focusing on new skills with a strong link to user communities.

3.1 PRACE-4IP Seasonal Schools Planning

Since the end of the PRACE-3IP project it was evident that there is more interest in hosting Seasonal Schools than the planned maximum number of schools in the scope of the projects, and with many countries already having hosted a school the selection of hosts was not trivial

anymore. Thus an open process to ensure fairness and transparency has been established since the beginning of the PRACE-4IP project. The process has been described in detail in D4.1 [1]

3.2 Selected Seasonal Schools Characteristics

The final plan for the PRACE-4IP Seasonal School schedule was proposed by the selection panel and accepted by the PRACE-4IP MB in July 2015. The following table summarises the Seasonal Schools dates, location and main subject. More details on the events are available in the relevant section of the PRACE agenda system that can be accessed via https://events.prace-ri.eu/category/6/

School	Location	Date	Main Subject
Winter 2015/2016	Bratislava, Slovakia	25–28 January 2016	Density Functional Theory (DFT)
Spring 2016	Ireland	16-20 May 2016	Material Science, e-CAM CoE
Autumn 2016	Austria	27 - 30 September 2016	Modern HPC Development for Scientists and Engineers
Winter 2016	Israel	6 - 9 February 2017	PRACE 2017 Winter School-Fuelling Scientific Discovery with HPC Infrastructure
Early Spring 2017	Sweden	10 - 13 April 2017	Biomolecular modelling techniques, BioExcel CoE
Spring 2017	Cyprus	25 - 27 April 2017	HPC compute, data management and domain specific services

Table 2: PRACE-4IP Seasonal School schedule

The following sections provide a summary of the contents and results of the Seasonal Schools that where organised since April 2016 (when D4.1 was submitted). Information on the Winter school 2015-2016 in Bratislava, Slovakia has been provided in D4.1. The information we have for the last two Spring schools organise in April 2017 are limited since the deliverable had to be prepared during the dates of the school.

Spring 2016: Ireland

The school was aimed at researchers who wish to gain a better understanding of methodologies and best practices in exploiting molecular and atomic modelling applications on HPC systems. The programme contained a mixture of scientific talks (HPC challenges in the field), sessions on HPC skills (parallel programming, numerical libraries) as well as application-oriented sessions with a large emphasis on hands-on practical exercises (e.g. classic molecular dynamics packages such as DL_POLY and Gromacs, electronic structure calculation packages such as CP2K and Quantum Espresso, covering example calculations, scalability and performance

considerations, suggestions for development such as implementing custom functions and the Python-based Atomic Simulation Environment).

The school was jointly organised and funded by PRACE and the Horizon 2020 E-CAM project, an e-Infrastructure for software, training and consultancy in simulation and modelling. The local organisers are the Irish Centre for High-End Computing (ICHEC) and CECAM-IRL (the Irish node of CECAM and partner in E-CAM).

The programme was designed in collaboration with representatives from CECAM-IRL (partner of the E-CAM CoE), focusing on methodologies and best practices in exploiting molecular and atomic modelling applications on HPC systems such as those available via the PRACE research infrastructure. It was conceived early on that the programme would contain a mixture of scientific talks (HPC challenges in the field), sessions on HPC skills (parallel programming, numerical libraries, computer architectures) as well as application-oriented sessions (e.g. ways to implement custom functions, running on different architectures), with a large emphasis on hands-on practical exercises. The major learning outcome for the audience is better insight into working with molecular and atomic modelling codes by gaining more in-depth understanding of HPC concepts, programming techniques and architectures. The school was also intended to serve as a short introduction to participants to scientific coding practices central to E-CAM and CECAM Extended Software Development Workshops, which are the principal generator of advanced simulation software for massively parallel systems and training scheme of E-CAM.

In total 32 participants attended the four (4) day school. The participants where from the following countries: Austria, Belgium, France, Ireland, the Netherlands, Spain and the UK. A total of 13 participants responded to the feedback survey as developed by PRACE for most of its training events. The school achieved an overall rating of 8.6 (0: worst, 10: best). Many aspects of the school, from information about the event, registration, quality of the venue, catering and organisation were all deemed to be either good or excellent by the participants. There were general agreement (>85%) that the topic being taught were relevant to the participants' research; most (>90%) found that the lectures were clearly presented and comprehensible, and all respondents agreed that the pace of teaching was about right for the school.

Autumn 2016: Austria

The goal of this seasonal school was to make participants familiar with modern techniques and tools for the development of scalable applications on modern computer architectures (massively parallel systems as well as many-integrated-core architectures). We pursued an integrated approach where various topics were presented that complement each other. In order to make the programme attractive for participants with varying backgrounds, the programme offered both, more basic topics (parallel programming with OpenMP and MPI, parallel I/O, profiling techniques and tools) and more advanced/specialised ones (Xeon Phi programming, GPU programming, Portable Extensible Toolkit for Scientific computing) from which the participants could individually select that portfolio that was most suitable for their background.

The workshop ran for four days, ending in the middle of the last day.

- Day 1: introduction, presentations of invited academic and industrial speakers on the current state-of-the-art and advanced topics, poster session with presentations of the participants, panel discussion. Common lunch and coffee breaks; free evening for social interactions and professional networking.

- Day 2: common invited talk, then two parallel lecture modules, each with an about 2.5h morning and a 3.5h afternoon session (with coffee breaks); common lunch between the sessions. Sessions include combinations of presentations and hand-on work. Free evening for social interactions and professional networking.
- Day 3: like second day. Late afternoon/evening was reserved for a common social programme.
- Day 4: two parallel tracks with a 3h morning session; close-up of school, common lunch.

Lecture topics were organized in six modules (two parallel modules per day from day 2 on) that could be booked individually. Attendance was limited to max. 25 persons per module. Attendees had marked their preference for module A or B for each training day individually. The assignment of applicants to training modules was done after the application deadline by the selection committee such that the given preferences was respected as well as possible.

During the application period (4 April - July 2016), the school received 58 applications from 13 countries (Austria 21, Bosnia and Herzegovina 1, Czech Republic 5, Denmark 1, France 1, Germany 5, Hungary 6, Israel 1, Italy 2, Romania 5, Spain 4, Turkey 5, UK 1) which were periodically evaluated by the admission committee, all considered eligible, and thus accepted in multiple batches. Due to withdrawals after acceptance (in particular, 4 Turkish applicants did not manage to get their visa), we finally prepared the school for 44 trainees. Ultimately, due to some last minute cancellations and no-shows 38 trainees were participating in the school. Together with 15 lecturers/trainers and 3 organisers there were thus in total 56 participants at the school.

Winter 2016/2017: Israel

The winter school's in Israel programme committee, consisted of experts from the local HPC academic community, computational scientists close to the true need of local and regional users and computational scientists who had participated in a PRACE "train the trainer" programme supported via the PATCs (i.e. https://events.prace-ri.eu/event/510/). The programme was designed to accommodate trainees with varying expertise level and professional backgrounds and to give a very thorough overview and training in fundamental HPC topics and to introduce topics of importance and interest to the computational science community targeted. Input was also solicited from younger students who had participated in a PRACE training programme and was in tune with the needs of the geophysics community.

The topics covered various parallel programming techniques:

- Intro + MPI + Open MP: Frontal lectures and hands-on exercises on parallel computing in distributed and shared memory systems.
- Parallel I/O: This session opened with a general overview about parallel I/O strategies, parallel I/O bottlenecks, followed by an introduction to MPI-I/O (or SIONlib), including a hands-on exercise. The second session included an introduction to parallel, portable data formats with HDF5, including a hands-on exercise.
- GPU: This session focused on GPU programming and included a hands-on exercise, as it differs more from traditional shared memory programming than the models usually used for programming the Xeon Phi.
- Intel XEON Phi: Lecture and hands on exercise.
- Performance Tools: This session opened with an introduction on methodology, followed by introducing a well-known portable, free, open-source tool set and closed with a one hour hands-on exercise.

- GPU Algorithm Design: Lecture and hands-on exercise on GPU architecture and programming, algorithms, and methods for adapting algorithms for efficient GPU utilisation.
- LAMMPS: Lecture and tutorial on using the LAMMPS molecular dynamics simulator.
- Multiprocessing in Pythion: An introduction and hands-on exercises in multiprocessing in python both at node level and cluster level and technologies in Python that allow scientists to parallelise work on a cluster and multi-core (also in MPI), compared to traditional tools in C or Fortran.

The school received a total of 75 applications which was narrowed down after requiring firm confirmation and cancelations. The final number of training participants was 58 originating from the following 5 countries: Israel (54), Austria (1), Italy (1), Spain (1) and Turkey (1).

Early Spring 2017 Sweden

Life Science research has become increasingly digital, and this development is accelerating rapidly. Biomolecular modelling techniques such as homology modelling, docking, and molecular simulation have advanced tremendously due to world leading European research, resulting in extreme demands for better computational performance and throughput as these tools are used in applied research and industrial development. This research has direct influence on our daily life in areas such as health and medical applications, the development of new drugs, efficient drug delivery, biotechnology, environment, agriculture and food industry. Life Science is one of the largest and fastest growing communities in need of high-end computing, and it is a critically important industrial sector for Europe. The dedicated CoE for Biomolecular Research (BioExcel) has recently been established in Sweden, Stockholm where KTH is the coordinating partner. Thus the mission of this school was to give a comprehensive curriculum on several widely used life science software on PRACE HPC systems. Discuss their scalability and performance issues and teach best practices on HPC systems that have been observed during PRACE research activities. Extensive hands-on sessions covered more than half of the school period. 69 students participated event from 17 different countries including Austria (2), China (1), Denmark (1), Finland (2), France (1), Germany (3), India (1), Italy (3), Norway (1), Poland (1), Russian Federation (2), Slovakia (2), Spain (1), Sweden (43), Switzerland (1), Ukraine (1) and the United Kingdom (2).

Mid Spring 2017: Cyprus

The mission of the Seasonal School in Cyprus was to have two parallel sessions to cater for the training requirements of different individuals in the PRACE community.

One of these sessions was the HPC system's administrator track. This is something that PRACE had not yet offered in its Seasonal School training and has been considered beneficial to PRACE and other HPC sites. Given the large number of systems present and available to users through PRACE Project Calls and PRACE DECI calls, it is important for the system administrators of these sites to have up-to-date, state-of-the-art education on the latest technologies and tools which could help in their duties and allow them to better help PRACE and their respective users who gain access to their systems. Furthermore, experienced administrators from PRACE where able to share their expertise, network with each other, discuss, advice and share ideas.

The second session was co organised in collaboration with the regional (South Eastern Europe and Mediterranean) Virtual Research Environment (VI_SEEM [https:vi-seem.eu]) that supports the fields of life sciences, climatology and cultural heritage. The focus was not only

on HPC compute but also data management and domain specific services deployed over different computational resources including HPC, cloud and grid. Parts of the second session where relevant to PRACE communities as apart from HPC the track provided them with training on data management and more specifically training for data generation, processing, preparation of the data for the simulation step and data analysis.

The total number of participants to this event was 69, with 28 participants in the developer track and 41 in the system administrator's track. The event attracted students from 20 different countries including those from central and Western Europe as well as those of the south Eastern Europe and eastern Mediterranean areas.

4 PRACE On-Demand Events

On-demand events organised in the context of the PRACE-4IP project are events specifically organised with the collaboration of research communities that have special needs for training and the expertise of PRACE trainers. Such targeted communities are the Centres of Excellence (CoEs). PRACE-4IP made special effort to collaborate with the CoE in many different areas, including training. To better understand the requirements for the different CoEs, PRACE-4IP WP4 and WP7 prepared a technical discussion with different members of the CoEs. The results and recommendations have been documented in D4.1 [1].

Based on those recommendation WP4 partners proposed, planed and organised a series of such events. The following table presents the six On-Demand Events schedule for the duration of the PRACE-4IP project.

Location	Date	Main Subject
Sofia, Bulgaria	25 – 28 April 2016	Code Modernisation for Intel Multi Core and Xeon Phi Architectures
Stuttgart, Germany	17 – 18 October 2016	Parallel Programming with MPI and OpenMP for Beginners, PATC-on-demand course for project CoeGSS
Ostrava, Checz Republic	14 – 15 December 2016	Performance Optimization and Productivity Tutorial (POP CoE)
Bologna, Italy	5 – 6 December 2017	Material Science codes on innovative HPC architectures: targeting exascale (MaX CoE)
Sofia, Bulgaria	22 – 24 March 2017	Practical Programming Models and Skills on INTEL Xeon for Scientific Research Engineers
Barcelona, Spain	24 – 27 April 2017	3rd Training Workshop EoCoE – POP

Table 3: PRACE-4IP On-Demand Events schedule

Follow some more detailed information on the organisation of the On-Demand Events.

Code Modernisation for Intel Multi Core and Xeon Phi Architectures

The 4 day school focused on software modernisation techniques needed for the next generation of supercomputers with highly dense parallel architectures, both homogeneous (Intel Xeon) and hybrid with acceleration co-processor (Intel Xeon Phi). The school programme was comprised of lectures and training exercises to address the crucial aspects of both the implementation of new HPC applications as well as the re-factoring of existing ones. These software engineering techniques for high productivity languages complement the more traditional lectures on parallel programming, to allow the implementation and continual modernisation of applications that need to be maintained across complex and fast evolving HPC architectures. The total number of participants in this school was 20. The overall score of the school was 9.1 out of 10.

Parallel Programming with MPI and OpenMP for Beginners, PATC-on-demand course for project CoeGSS

The CoeGSS on-demand, 2 day course, was designed as introductory part providing all the necessary skills to use MPI and OpenMP. It was also designed as prerequisite for the PRACE PATC advanced parallel programming course that was following the CoeGSS course. The week is designed to learn from the beginning most of MPI and OpenMP that is needed for an efficient use of HPC systems. The course is also designed to teach participants with very different previous knowledge. This course starts with very basic concepts and tries to reach 100% of the needed parallelisation skills for HPC. The participants were grouped in pairs for the exercises. Each group used one node on Hazel Hen, the PRACE Tier-0 system at HLRS.

Performance Optimisation and Productivity Tutorial (POP CoE)

The EU funded Performance Optimisation and Productivity Centre of Excellence (POP) offers services in the area of performance analysis and performance optimisation. Therefore, POP mainly uses performance analysis tools developed by POP partners, namely Extrae, Paraver, Dimemas, Score-P, Cube and Scalasca. These tools are used to generate uniform performance audits for parallel applications to provide a good overview of the applications execution.

The aim of this tutorial, especially during its first day, was to introduce participants to the performance analysis tools and the methodologies used to create POP performance audits. The second day was devoted to large hands-on sessions where application developers, who brought their own code, were doing a performance analysis with guidance of POP performance experts.

The tutorial was in some sense a follow-up event after the course Practical Parallel Performance Analysis on Salomon held at IT4Innovations on 20-21 October 2016, therefore most participants were pre-selected as they had participated in this event.

The total number of participants were 23, of them 21 where from Czech Republic and 2 from Germany. The overall rating of the event got a good average mark of 8.38 out of 10.

Material Science codes on innovative HPC architectures: targeting exascale (MaX CoE)

The course was organised in collaboration between CINECA and CSCS focusing on several of the main codes used by the material science community (namely, Quantum Espresso, Siesta, Yambo, Fleur), introducing their architecture and usage. The course presented how the codes exploit petascale systems, targeting in particular to PRACE Tier-0 systems, and how they are preparing for the exascale era, with specific attention to hybrid and accelerated architectures.

Hands-on were proposed to let the attendees to get familiar with the codes and their optimal usage on HPC systems.

The course addressed in particular the MAX community, more specifically the users of the selected codes and the software developers starting contributing to those codes with new components and functionalities targeting the support of novel architectures and the performance and scalability improvement.

The course was planned during three days and hosted at CINECA on 5-7 December 2016. Teachers were provided by CINECA and ETH Zurich PRACE partners and by the MAX project members. The total number of participants was 16 with 1 from Germany and 15 from Italy. The overall feedback was very positive. General evaluations show that the attendees liked the course, its organisation and the subject of the lectures. In the suggestions, different thoughts were expressed, mostly according the personal background of the attendees. There were no indications for major corrections in the organisation of this course, or actions that need to be undertaken to correct some issues.

Practical Programming Models and Skills on INTEL Xeon for Scientific Research Engineers

The school (4-days event) was a continuation of the first on demand event in Bulgaria and was organized by NCSA in assistance with Science and. The Technology Facilities Council (STFC) and Bayncore. The event focused on training practical skills by demonstrations on how to test performance scaling on Xeon Phi of a few project codes. The event targeted the software engineers from the previous event and trained in the skills already obtained by those via demonstrating how these are applied in a number of case studies whose various modes' performance scaled on the NCSA heterogeneous cluster Avitohol. The total number of participants was 20. The CoeGSS course had 5 participants from Sweden (1), Italy (1), France (1), and Spain (2). The overall evaluation of the event was 8.1 out of 10.

3rd Training Workshop EoCoE – POP

This training event (3 days event) similarly to the "**Performance Optimisation and Productivity Tutorial (POP CoE)"** addresses the needs of the POP CoE community and in addition to the EoCoE community. The course provides information on accessing specific HPC resources, specific tools such as Score-P, Scalasca, Vampir, Extrae and Paraver as well as courses on integration and benchmarking. For all the tools and methodologies the students are given the opportunity to do hands on exercised on HPC systems.

5 International HPC Summer Schools

The International HPC Summer School on Challenges in Computational Sciences is a series of events started in 2010 and taking place annually. Initiated by NSF's TeraGrid and EU's DEISA project, the first school was organised in Italy in 2010, and was then continued by NSF's XSEDE, EU's PRACE, Japan's RIKEN and Canada's Compute/Calcul projects with locations alternating between North America and Europe. The objective of the Summer School is to familiarise the best students of the respective continents or countries in computational sciences with a strong bond to supercomputing with all major state-of-the-art aspects related to HPC for a broad range of scientific disciplines, catalyse the formation of networks, provide mentoring

through faculty members and supercomputing experts from renowned HPC centres, and to facilitate international exchange and open further carrier options.

5.1 International HPC Summer School 2016

Graduate students and postdoctoral scholars from institutions in Europe, Canada, Japan and the United States participated to the seventh International Summer School on HPC Challenges in Computational Sciences, that was held between 26 June – 1 July 2016, in Ljubljana, Slovenia. The summer school was sponsored by Compute/Calcul Canada, the Extreme Science and Engineering Discovery Environment (XSEDE) with funds from the U.S. National Science Foundation, the Partnership for Advanced Computing in Europe (PRACE) and the RIKEN Advanced Institute for Computational Science (RIKEN AICS).
Leading European, Canadian, Japanese and American computational scientists and HPC technologists offered training on a variety of topics, including:

- HPC challenges by discipline (e.g, earth, life and materials sciences, physics)
- HPC Programming Proficiencies
- Performance analysis and profiling
- Algorithmic approaches and numerical libraries
- Data-intensive computing
- Scientific visualisation
- Canadian, EU, Japanese and U.S. HPC-infrastructures

The International HPC Summer School series typically utilises a combined organising/programme committee that oversees both logistics and the programme of the school; it also includes members of the local organising committee. It is a large committee with many members from all partners (PRACE, XSEDE, RIKEN AICS, Compute Canada), although there are typically person(s) assigned for specific tasks such as speaker/programme development, mentoring activities, practical session preparations, evaluation, etc.

79 students were selected and overall a total of 126 participants (students, staff, mentors, presenters) were in Ljubljana. The 2016 International HPC Summer School has been a hugely successful event. Building on experience from the past 7 years, the programme continues to deliver untold benefits to participants. Feedback has been overwhelmingly excellent for the school; there were particular praise for the tremendous efforts put into the event by the local organisers.

5.2 International HPC Summer School on 2017

The eighth International Summer School on HPC Challenges in Computational Sciences, will be held 25 - 30 June 2017, in Boulder, Colorado, United States of America. PRACE collaborated with partners from the United States, Canada and Japan for the organisation of this summer school.

6 PRACE MOOCs

The development of PRACE MOOCs and the selection of the MOOC platform were reported in Deliverable 4.4. "MOOC Pilot for HPC" [3]. Two courses have been developed: The first is "Supercomputing", which started 6 March 2017, and the second is "Managing Big Data with R and Hadoop" which started 3 April 2017. The courses are hosted by FutureLearn, the leading European MOOC platform, in English. The duration of the courses is five week each.

The "Supercomputing" course introduces to the layman what supercomputers are, how they are used and how one can exploit their full computational potential to make scientific breakthroughs. The course is designed for anyone interested in leading-edge computing technology, supercomputers or the role that computer simulation takes in modern science and engineering.

All of the technical aspects will be covered at a conceptual level and there is no requirement to be able to write computer programs. However, anyone with existing programming experience will learn how programming modern supercomputers differs from programming a home PC.

The objective of "Managing Big Data with R and Hadoop" is to introduce individuals with limited programming knowledge to various HPC facilities for big data analysis. At the end of the course, they will be able to use them, avoiding common pitfalls and thus saving them money and time. The course is especially suitable for participants interested in data science, computational statistics and machine learning. The course may be also useful for advanced undergraduate students and 1st year PhD students in data analysis, statistics, or bioinformatics who wish to gain a basic understanding in data management and HPC computing.

The main advantage and comparative variety to other similar MOOCs is that the R environment for statistical computing and graphics is being used. R is a programming language renowned for its simplicity, elegance, and the support of an outstanding community. R augmented with Hadoop allows data scientists to quickly utilise the enterprise-grade capabilities of Hadoop with the analytic capabilities of R.

The courses are designed according to FutureLearn's pedagogic philosophy. FutureLearn emphasizes social constructivist learning, in which the learner community plays a key role. FutureLearn courses are shorter, and use less video and more written material than the MOOC platforms Coursera and edX. FutureLearn favours the courses to be designed as a series of short steps each comprising a short article, quiz or discussion topic. Although use of video is encouraged, it is normally used to highlight specific topics and not as the major method of content delivery. Any videos are usually only a few minutes long and do not resemble traditional lectures. This design not only provides more flexibility to many learners, but also decreases production costs per course, which will enable us to provide more in the long run.

6.1 Syllabi

Supercomputing:

Week 1: Supercomputers

This covers what modern supercomputers look like and why they are designed that way. There is a short history of supercomputers over the years to illustrate how performance has increased over time

Week 2: Parallel Computers

The concepts of shared and distributed memory architectures are explained. We cover the similarities and differences between specialist supercomputers and more general purpose computers such as PCs, laptops and games machines.

Week 3: Parallel Computing

The basic parallel programming models (shared and distributed memory) are explained at a conceptual level using a traffic modelling example. This is used to motivate the basics of parallel performance, including Amdahl's and Gustafson's laws, but without using any equations.

Week 4: Computer Simulation

This covers what supercomputers are used for, e.g. the kinds of simulations that are done in the areas of nanotechnology, engineering and climate research. The traffic model serves as a very basic example of computer simulation, illustrating how computers effectively run virtual experiments where scientists choose the input parameters to model real situations.

Week 5: Case Studies

The final week comprises a number of case studies of real scientific problems being tackled using supercomputers. Ideally, these have associated material that makes the talks interesting, e.g. visualisations of real simulations. We explain the approaches taken in terms of the concepts introduced in the previous weeks' lectures.

Managing Big Data with R and Hadoop:

Week 1: Welcome to BIG DATA MOOC

This covers introduction to big data and distributed systems. Parallel computing systems and parallel databases are explained in great details. Students are introduced to big data, Linux, virtual machine provided for the course, terminal window, streaming, Hadoop, and basic Unix commands.

Week 2: First steps in R and RStudio

In this lecturewe cover an introduction to R and Rstudio. R is a leading programming language for computational statistics and graphics and could be augmented for big data processing. Rstudio is a powerful and productive graphical user interface for R, make R easier and more efficient.

Week 3: Working with Apache Hadoop I (Fundamentals)

This covers massively parallel computing with Hadoop. We present Hadoop as a system for distributed computing and data storage, show how to use "ravro" library for reading and writing files in avro format and present powerful "plyrmr" library for data processing. We also explain how to set up local virtual environment and prepare an image of virtual machine that the students will download and locally install.

Week 4: Working with Apache Hadoop II (RHadoop)

This lecture covers introduction to functions providing management of the Hadoop Distributed File System (HDFS) using the "rhdfs" library. We teach how R programmers can browse, read, write, and modify files stored in HDFS. We also present how to perform statistical analysis via Hadoop MapReduce functionality on a Hadoop cluster. Finally, we introduce HBASE distributed database and present how to browse, read, write, and modify tables stored in it.

Week 5: Statistical learning

This covers introduction to machine learning and to classification framework in particular. Wepresent principles of supervised and unsupervised learning, the most frequent algorithms in linear regression, classification (discriminant analysis) and clustering (hierarchical clustering and k-means). For all these algorithms we demonstrate how to implement them using R and RHadoop to handle big data.

6.2 Learning outcomes

The courses have well-defined expected learning outcomes.

Supercomputing

- Understand how the performance of modern supercomputers is measured and achieved;
- Explain why they are built from thousands of simple processors;
- Understand the differences between of shared-memory and distributed-memory computers;
- Compare the architecture of a typical modern supercomputer with a desktop PC;
- Explain why computer simulation is a fundamental component of modern scientific discovery;
- Work with simple cellular automaton models;
- Analyse simple problems and look for opportunities for parallel processing;
- Explain the limitations of parallel computing;
- Give examples of scientific areas where computer simulation is used.

Managing Big Data with R and Hadoop

- Understand how the performance of modern supercomputing is achieved;
- Understand the basic functionality of Bash terminal window;
- Understand the basic functionality of Apache Hadoop for scalable, distributed computing;
- Understand the basic functionality of RHadoop;
- Understand the basic problems of supervised and unsupervised learning; Perform basic clustering, regression and classification with RHadoop.

6.3 Participations

The two MOOCs have turned out to be highly successful. For “Supercomputing” there has been 3020 joiners, and 1487 learners out of which 1001 have been active learners. These are complete statistics for the Supercomputing course, accurate up to midnight on 8 April 2017. In these statistics, joiner is defined as user having account in the FutureLearn platform and having enrolled in the course (includes also educators and admins). Learners are users (of any role)

who have at least viewed at least one step at any time in any course week, and active learner are those who have completed at least one step at any time in any course week.

“Managing Big Data with R and Hadoop” has just startedat the time of this writing, and has currently 2887 joiners at week one, accurate up to midnight on 8 April 2017.

7 PRACE CodeVault

The PRACE CodeVault is an open repository containing various high performance computing code samples. The project aims to support self-learning of HPC programming and will be used as an Open platform for the HPC community to share example code snippets, proof-of-concept codes and so forth. CodeVault currently uses a Gitlab platform. The project is a collaboration of WP4, WP6 and WP7. WP4 is has been responsible for coordination of the project and for content generation, WP6 has assisted with technical issues, and both WP4 and WP7 have been responsible for the content generation. WP7 has provided HPC kernels, suitable for utilization in real-world applications, while WP4 works on code samples that can be used in learning HPC programming concepts.

Currently, WP4 has provided examples of parallel programming concepts with OpenMP and MPI, both with Fortran and C programming languages.

A pilot version of CodeVault was launched in January 2016 and promoted to PRACE members and HPC community with help from WP3. The repository is currently available at https://gitlab.com/PRACE-4IP/CodeVault.git, hosted by Gitlab.com, but transition to PRACE hosted service is currently in progress. The repository is open for anonymous access, and the code samples are made available with open source licenses. For contributing, registration is needed.

8 Conclusions

The Training WP4 of the PRACE-4IP project has continued to provide high quality training by offering large variety of face-to-face training events as well as by developing on-line training services. In total, PRACE-4IP has delivered 569 training days as part of 198 training events that have reached an audience of 4,492 participants. Six Seasonal Schools have been organised within PRACE-4IP, and in collaboration with international partners the series of International HPC Summer Schools has continued to prosper. The training program has been highly appreciated as shown by the feedback of various events. As new activities, PRACE has started the conception of PRACE Training Centres (PTC) for complementing the PRACE Advanced Training Centres. In order to better serve research communities such as the Centres of Excellence (CoEs), a new type of On-Demand training event has also been introduced, and five such events have been organised. In order to support better self-learning, PRACE has developed two Massively Open On-line Courses (MOOCs), and piloted the CodeVault repository for HPC code samples. The MOOCs have been highly successful, as 5,907 participants have enrolled in the two courses, and 1487 participants have been active in the “Supercomputing” so far (“Managing Big Data with R and Hadoop” has not yet started). Overall, PRACE-4IP WP4 has delivered a vital service to foster the growth of computational skills in Europe and to strengthen the competitiveness of European science and industry in the field.

9 Annex

Detailed reports of PRACE Seasonal Schools, On-Demand Events, and the International HPC Summer School 2016.

9.1 Report on PRACE Spring School 2016

9.1.1 *Basic information about the event*

Name: PRACE Spring School 2017 and E-CAM Tutorial on Molecular and Atomic Modelling

Dates: 16-20 May 2016

Location: Dublin, Ireland

Organizing sites: ICHEC (NUI Galway)

Mission: The school is aimed at researchers who wish to gain a better understanding of methodologies and best practices in exploiting molecular and atomic modelling applications on HPC systems. The programme contained a mixture of scientific talks (HPC challenges in the field), sessions on HPC skills (parallel programming, numerical libraries) as well as application-oriented sessions with a large emphasis on hands-on practical exercises (e.g. classic molecular dynamics packages such as DL_POLY and Gromacs, electronic structure calculation packages such as CP2K and Quantum Espresso, covering example calculations, scalability and performance considerations, suggestions for development such as implementing custom functions and the Python-based Atomic Simulation Environment).

The school was jointly organised and funded by PRACE and the Horizon 2020 E-CAM project, an e-infrastructure for software, training and consultancy in simulation and modelling. The local organisers are the Irish Centre for High-End Computing (ICHEC) and CECAM-IRL (the Irish node of CECAM and partner in E-CAM).

Event URL: https://events.prace-ri.eu/event/497/

9.1.2 *Organizational details*

Local organizing committee

Simon Wong	**Committee chair**
Donald MacKernan	**General assistance (E-CAM)**
Kate Collins	**General assistance (E-CAM)**
Lisa Walkin	**General assistance**

Venue

A suitable room on the campus of University College Dublin (UCD), also host of the CECAM-IRL node and a partner of E-CAM, was chosen as the venue for the school. The location is easily accessible by public transport to Dublin city centre, as well as to/from Dublin international airport. Designated hotel accommodation was arranged prior to the school to house participants from Europe.

The venue, situated at the state-of-the-art O'Brien's Centre for Science building, also offered a well laid out room (with capacity for up to 50 students) for hands-on training delivery where students are seated in groups at a number of tables. Each student was required to bring along

individual laptops to carry out hands-on exercises, where they can work in groups and instructors have easy access to any student who require assistance. The venue also offer fast wireless internet access, including Eduroam login for visiting participants. Catering for the event was also facilitated by an on-site restaurant, where a section was reserved for breaks.

9.1.3 *Programme and content*

Program committee

Simon Wong	**Chair**
Donal MacKernan	**Member**
Alin Elena	**Member**
Michael Lysaght	**Member**
Iain Bethune	**Member**

Final program

Monday 16th May		
09:30 - 09:50	Welcome and opening	Simon Wong & Donal MacKernan
09:50 - 10:30	Introduction to parallel programming concepts	Michael Lysaght
10:30 - 11:00	*BREAK*	
11:00 - 12:30	OpenMP overview I (main concepts, parallel loops)	Alin Elena
12:30 - 13:30	*LUNCH*	
13:30 - 15:00	OpenMP overview II (synchronisation, performance)	Alin Elena
15:00 - 15:30	*BREAK*	
15:30 - 17:00	MPI overview I (main concepts, point-to-point communication)	Michael Lysaght
Tuesday 17th May		
09:30 - 11:00	MPI overview II (collective communication)	Ruairi Short
11:00 - 11:30	*BREAK*	
11:30 - 12:30	Numerical libraries (overview of libraries, linking)	Alin Elena
12:30 - 13:30	*LUNCH*	
13:30 - 15:00	Numerical libraries (parallel libraries, performance considerations)	Alin Elena
15:00 - 15:30	*BREAK*	
15:30 - 17:30	Electronic poster session	
Wednesday 18th May		
09:30 - 10:30	Scientific & HPC challenges: Molecular simulation in the exascale era	Erik Lindahl
10:30 - 11:00	*BREAK*	
11:00 - 12:30	Overview of classical molecular dynamics packages (DL_POLY, Gromacs) and introduction to practical exercises	Alin Elena (DL_POLY) Erik Lindahl (Gromacs)
12:30 - 13:30	*LUNCH*	
13:30 - 15:00	Practical session on DL_POLY & Gromacs Free-form tutorial with example calculations, scalability tests, performance considerations, suggestions for development, etc.	Alin Elena (DL_POLY) Erik Lindahl (Gromacs)
15:00 - 15:30	*BREAK*	

15:30 - 17:00	Practical session on DL_POLY & Gromacs Free-form tutorial with example calculations, scalability tests, performance considerations, suggestions for development, etc.	Alin Elena (DL_POLY) Erik Lindahl (Gromacs)
19:00 - 21:30	Social dinner, Dublin city centre	
Thursday 19th May		
09:30 - 10:30	Scientific & HPC challenges: Simulations from first principles	Emilio Artacho
10:30 - 11:00	*BREAK*	
11:00 - 12:30	Overview of electronic structure calculation packages (CP2K, QuantumEspresso) and introduction to practical exercises	Iain Bethune (CP2K) Stefano de Gironcoli (QuantumEspresso)
12:30 - 13:30	*LUNCH*	
13:30 - 15:00	Practical session on CP2K and QuantumEspresso Free-form tutorial with example calculations, scalability tests, performance considerations, suggestions for development, etc.	Iain Bethune (CP2K) Stefano de Gironcoli (QuantumEspresso)
15:00 - 15:30	*BREAK*	
15:30 - 17:00	Practical session on CP2K and QuantumEspresso Free-form tutorial with example calculations, scalability tests, performance considerations, suggestions for development, etc.	Iain Bethune (CP2K) Stefano de Gironcoli (QuantumEspresso)
Friday 20th May		
09:30 - 10:30	Atomic Simulation Environment I	Jussi Enkovaara
10:30 - 11:00	*BREAK*	
11:00 - 12:30	Atomic Simulation Environment II	Jussi Enkovaara
12:30 - 13:30	*LUNCH*	
13:30 - 14:30	Atomic Simulation Environment III	Jussi Enkovaara

List of trainers

Michael Lysaght	**ICHEC, Ireland**
Alin Elena	**STFC Daresbury, UK**
Ruairi Short	**ICHEC, Ireland**
Erik Lindahl	**Stockholm University, Sweden**
Emilio Artacho	**Nanogune, Spain**
Iain Bethune	**EPCC, UK**
Stefano de Gironcoll	**SISSA, Italy**
Jussi Enkovaara	**CSC, Finland**

Designing the program

The programme was designed in collaboration with representatives from CECAM-IRL (partner of the E-CAM CoE), focusing on methodologies and best practices in exploiting molecular and atomic modelling applications on HPC systems such as those available via the PRACE research infrastructure. It was conceived early on that the programme would contain a mixture of scientific talks (HPC challenges in the field), sessions on HPC skills (parallel programming, numerical libraries, computer architectures) as well as application-oriented sessions (e.g. ways to implement custom functions, running on different architectures), with a large emphasis on hands-on practical exercises. The major learning outcome for the audience is better insight into working with molecular and atomic modelling codes by gaining more in-depth understanding of HPC concepts, programming techniques and architectures. The school was also intended to

serve as a short introduction to participants to scientific coding practices central to E-CAM and CECAM Extended Software Development Workshops, which are the principal generator of advanced simulation software for massively parallel systems and training scheme of E-CAM.

Computer resources

Since practical sessions constituted a major part of the school's programme, students were asked to each bring a laptop to the school for logging into HPC systems. ICHEC provided its national HPC system, Fionn, for the purpose of this school. The main partition of Fionn is a 7,680 processor SGI ICE X cluster with 24 compute cores and 48 GB memory per node.

9.1.4 *Participants & feedback*

Number of participants by country

Austria	**1**
Belgium	**1**
France	**1**
Ireland	**25**
Netherlands	**1**
Spain	**1**
UK	**2**
Total	**32**

Process for selecting the participants

Participants were accepted on a first come, first served basis.

Statistics and analysis of the feedback

A total of 13 participants responded to the feedback survey as developed by PRACE for most of its training events. The school achieved an overall rating of 8.6 (0: worst, 10: best). Many aspects of the school, from information about the event, registration, quality of the venue, catering and organisation were all deemed to be either good or excellent by the participants. There were general agreement (>85%) that the topic being taught were relevant to the participants' research; most (>90%) found that the lectures were clearly presented and comprehensible, and all respondents agreed that the pace of teaching was about right for the school.

From some of the free-form feedback collected from the evaluation survey, the respondents highlighted the following positives about the school: (i) the hands-on exercises, (ii) good overview of trends and recent developments in chemistry, i.e. HPC challenges, (iii) the ability to interact with simulation code experts and developers, (iv) the competency of the speakers/instructors. When asked about any negatives about the school, the only major criticism is that there were insufficient time to complete some of the practical exercises. Admittingly, the programme was somewhat ambitious in that the committee wanted to provide participants with a broad overview of different simulation packages; to stimulate students in learning, leading to their attendance at more in-depth courses/tutorials on specific packages/methods that are organised by PRACE, and also CECAM/E-CAM.

Conclusions & lessons learnt

The conclusion, having evaluated the feedback from participants and based on feedback from the trainers, is that the school has been a real success. It provided the participants with better insight into working with molecular and atomic modelling codes by interacting with experts and developers, and led to better understanding of HPC concepts, programming techniques and architectures.

9.2 Report on PRACE Autumn School 2016

9.2.1 *Basic information about the event*

Name PRACE 2016 Autumn School "Modern HPC Development for Scientists and Engineers"

Dates September 27-30, 2016

Location Hagenberg, Austria

Organizing sites RISC Institute/Johannes Kepler University Linz, RISC Software (Hagenberg), VSB-Technical University of Ostrava /IT4Innovations.

For detailed information, see

http://events.prace-ri.eu/e/PRACE-Autumn-School-2016
http://www.risc.jku.at/projects/prace/AutumnSchool2016

9.2.2 *Organizational details*

Local organizing committee

- Wolfgang Schreiner, RISC/JKU Linz
- Michael Krieger, RISC Software
- Thomas Ponweiser, RISC Software
- Ondrej Jakl, VSB-Technical University of Ostrava
- David Horak, VSB-Technical University of Ostrava

Venue

The school took place in the castle of Hagenberg, a medieval castle 20 km from Linz in the beautiful landscape of the Mühlviertel, see http://www.risc.jku.at/about/castle. The castle has its roots in the 12th century and was in the following centuries repeatedly extended. At the end of the 20th century, it has been renovated and adapted for the needs of a modern computer science institute; since 1989, RISC is located in the castle. The ingenious combination with historical elements brought the architects Prof. Peter Riepl and Thomas Moser an award of the government of Upper Austria.

Together with an extension building opened in 2013, the facilities available for the seasonal school include two seminar rooms for 40-50 persons each, smaller meeting rooms, Internet access and wireless LAN, office and secretary infrastructure. Next to the castle there is the "Schlossrestaurant" which offers lunches, dinners, and catering services on demand. Other

restaurant facilities are nearby. The location can be internationally reached by car, train, or plane via Linz, the capital of Upper Austria, see http://www.risc.jku.at/about/map/ for more information. Like for the Spring School 2014, participants we plan to accommodate participants in the Sommerhaus Hotel Linz, see http://www.sommerhaus-hotel.at/en/index.php and provide bus shuttle services to/from the venue of the seasonal school.

The site was already chosen for the PRACE 2014 Spring School; it was very well received by the participants of the event.

9.2.3 *Program & content*

Program committee

- Wolfgang Schreiner, RISC/JKU Linz
- Michael Krieger, RISC Software
- Thomas Ponweiser, RISC Software
- Ondrej Jakl, VSB-Technical University of Ostrava
- David Horak, VSB-Technical University of Ostrava
- Karoly Bosa, JKU/CDCC
- Dieter Kranzlmüller, LMU/LRZ

Final program

For details, see http://www.risc.jku.at/projects/prace/AutumnSchool2016/Programme.pdf

The workshop ran for four days, ending in the middle of the last day.

Day 1: introduction, presentations of invited academic and industrial speakers on the current state of the art and advanced topics, poster session with presentations of the participants, panel discussion on a "hot topic". Common lunch and coffee breaks; free evening for social interactions and professional networking.

Day 2: common invited talk, then 2 parallel lecture modules, each with an about 2.5h morning and a 3.5h afternoon session (with coffee breaks); common lunch between the sessions. Sessions include combinations of presentations and hand-on work. Free evening for social interactions and professional networking.

Day 3: like second day. Late afternoon/evening is reserved for a common social program (to be defined).

Day 4: 2 parallel tracks with a 3h morning session; close-up of school, common lunch.

Lecture topics were organized in 6 modules (2 parallel modules per day from day 2 on) that could be booked individually. Attendance was limited to max. 25 persons per module. Attendees had marked their preference for module A or B for each training day individually. The assignment of applicants to training modules was done after the application deadline by the selection committee such that the given preferences was respected as well as possible.

Day 1	Introduction, Presentations, Poster session, Panel discussion	
Day 2	Module 2a: **Parallel Programming** • OpenMP (morning) • MPI (afternoon)	Module 2b: **Intel Xeon Phi Programming** • Native mode, Intel offload, OpenMP • Hands-on sessions
Day 3	Module 3a: **Parallel IO** • SIONlib • MPI-IO • HDF5 / (P)NetCDF	Module 3b: **PETSc Tutorial** • Vectors, Matrices • Linear Solvers, etc. • Hands-on sessions
Day 4	Module 4a: **Tools for Performance Analysis** • Performance Analysis: HPCToolkit • IO-Profiling: Darshan, Vampir, TAU • Hands-on sessions	Module 4b: **Advanced Parallel Programming** • Levels of Parallelism on Multi-Core CPUs • Memory Hierarchies • Vectorization

List of trainers and presenters

- Dr. techn. Christoph ANTHES, Leibniz Supercomputing Centre (LRZ)
- DI Dr. Christoph GONIVA, DCS Computing GmbH
- Bernhard GRUBER, BSc., RISC Software GmbH
- Vaclav HAPLA, M.Eng., VŠB-Technical University of Ostrava
- Michael HAVA, MSc., RISC Software GmbH
- Dr. David HORAK, VŠB-Technical University of Ostrava
- DI (FH) Alexander LEUTGEB, RISC Software GmbH
- Prof. Bernhard MANHARTSGRUBER, Johannes Kepler University Linz
- DI Andreas MAYR, Johannes Kepler University Linz
- Dr. Sandra MENDEZ, Leibniz Supercomputing Centre (LRZ)
- DI Thomas PONWEISER, RISC Software GmbH
- Prof. Ulrich RÜDE, Friedrich-Alexander Universität Erlangen-Nürnberg
- Ing. Radim SOJKA, VŠB-Technical University of Ostrava
- DI Thomas UNTERTHINER, Johannes Kepler University Linz
- Dr. Volker WEINBERG, Leibniz Supercomputing Centre (LRZ)

Designing the program

The goal of this seasonal school was to make participants familiar with modern techniques and tools for the development of scalable applications on modern computer architectures (massively parallel systems as well as many-integrated-core architectures). We pursued an integrated approach where various topics are presented that complement each other. In order to make the program attractive for participants with varying backgrounds, the program offered both, more basic topics (parallel programming with OpenMP and MPI, parallel I/O, profiling techniques and tools) and more advanced/specialized ones (Xeon Phi programming, GPU programming, Portable Extensible Toolkit for Scientific computing) from which the participants could individually select that portfolio that was most suitable for their background.

Description of the contents

Module 2a: Parallel Programming

This module focuses on two state-of-the-art programming models for HPC applications: OpenMP and MPI. For both programming models, basic as well as certain selected advanced topics (e.g. for MPI new features such as non-blocking or sparse neighborhood collectives) are presented. For the MPI session, special attention is also put on best practices for achieving good program performance, based on the presenter's experience from the support of recent PRACE Preparatory Access Type C projects.

Module 2b: Intel Xeon Phi Programming

In this module, Intel's Many Integrated Core (MIC) architecture is introduced. The session covers various programming models for Intel Xeon Phi coprocessors (like native mode vs. offload mode, OpenMP and MPI parallelization etc.) as well as some selected optimization techniques. Hands-on sessions are planned to take place on LRZ's Intel Xeon Phi based system SuperMIC.

Module 3a: Parallel-IO

This module covers parallel IO concepts related with parallel file systems, IO techniques and performance analysis. Furthermore, it introduces the IO libraries MPI-IO, SIONlib and high level libraries HDF5 and NetCDF. The theoretical part will be complemented by practical exercises for each presented library.

Module 3b: PETSc Tutorial

The Portable Extensible Toolkit for Scientific computing (PETSc) is a modular library for linear algebra, nonlinear solvers, time integrators, optimization, and spatial discretization. Solver configuration and diagnostics are valuable skills for users, whether calling PETSc directly or via one of many higher level packages that access PETSc solvers. The tutorial will start with the fundamental linear algebra components then proceed to principles of preconditioning and Krylov solvers, convergence diagnostics, performance analysis, and the higher level solver interfaces. It will contain hands-on exercises to build the skills necessary to evaluate methods and design solvers for complex problems in science and engineering.

Module 4a: Tools for Performance Analysis

This module gives an introduction to effective strategies for analyzing performance and IO behavior of HPC applications. The focus will lie on HPCToolkit for performance analysis as well as on Darshan, Vampir and TAU for IO Profiling and IO Tracing. Hands-on sessions shall lower the threshold for attendees to actually using these tools in the course of their every-day work.

Module 4b: Advanced Parallel Programming

Exploiting Parallelism on Multi-Core CPUs Considering Memory Hierarchies

This session gives at first the motivation for parallelization in central processing units (CPUs). The different levels of parallelism implemented in hardware are presented in the case study of the Intel x86 Sandy Bridge architecture. These are namely task level, control level, data level, and instruction level parallelism. For each level the implementation in hardware is illustrated. We analyze the relevance of each level from a programmer's point of view. Because parallel algorithms are not only bound by the computing power (peak floating point performance) of the CPU, but also limited by the memory bandwidth between CPU and the main memory, memory hierarchies are presented in more detail. We show the motivation for Caching in

Hardware and what kind of problems arise from Caching in a parallel context. Finally the Roofline Model is presented. It allows for any parallel algorithm with a certain operational density to calculate the attainable floating point performance considering the peak floating point performance and the peak memory bandwidth of the target machine.

Introduction to Vectorization

This session gives an introduction into vectorization on Intel x86 CPUs from a programmer's point of view. We will present the potential of vector units. The different levels of the programmers control over the vectorization are shown. Special focus lies in the auto vectorization support of the Intel C/C++ compiler. The inhibitors of auto vectorization and techniques how to get rid of them will be presented. Finally we will talk about portability issues, if we use the latest vector unit features of Intel CPUs.

Computer resources

Participants were expected to bring their own notebooks for their hands-on work (software to be installed locally, e.g. based on virtual machines, was defined and communicated in time). Internet and wireless LAN access were available. HPC resources were available via shared memory multiprocessors available at RISC respectively via Internet connections to the computing center of IT4Innovations (for the Xeon Phi Sessions). Resources were generally adequate, only the quality of the WLAN connection in one of the seminar rooms was partially criticized.

9.2.4 *Participants & feedback*

Number of participants by country

During the application period (April 4 to July 1 2016), we received for the school 58 applications from 13 countries (Austria 21, Bosnia and Herzegovina 1, Czech Republic 5, Denmark 1, France 1, Germany 5, Hungary 6, Israel 1, Italy 2, Romania 5, Spain 4, Turkey 5, UK 1) which were periodically evaluated by the admission committee, all considered eligible, and thus accepted in multiple batches. Due to withdrawals after acceptance (in particular, 4 Turkish applicants did not manage to get their visa), we finally prepared the school for 44 trainees. Ultimately, due to some last minute cancellations and no-shows 38 trainees were participating in the school. Together with 15 lecturers/trainers and 3 organizers there were thus in total 56 participants at the school.

A list of all applicants and final participants is given in the appendix.

The programme was structured into general talks on Days 1-3 and two parallel strands (a and b) of three training modules on Days 2-4 where the planned distribution (not considering last minute cancellations and no-shows) was as follows:

Module 2a "Parallel Programming": 32 trainees (originally 43 applicants)
Module 2b "Intel Xeon Phi Programming": 12 trainees (originally 15 applicants)

- Module 3a "Parallel I/O": 27 trainees (originally 35 applicants)
- Module 3b "PETSc Tutorial": 17 trainees (originally 23 applicants)

- Module 4a "Tools for Performance Analysis": 19 trainees (originally 23 applicants)
- Module 4b "Advanced Parallel Programming": 15 trainees (originally 35 applicants)

Interest in the modules was therefore quite balanced with the exception of Module 2a "Parallel Programming" (a general introduction to OpenMPI and MPI) which was substantially more popular than the more special Module 2b "Intel Xeon Phi Programming".

Process for selecting the participants

Applications were expected to explain in the application form their background and rationale for their participation in the seasonal school. A selection committee with representatives of the organizing parties decided about the acceptance. All members of the committee received every application and a suggestion of the chair of the selection committee about the acceptance of the applicant. If there was not unanimous acceptance of the suggestion, discussions proceeded by email or tele-conferencing, until either consensus was achieved or a vote decided the issue by majority.

The application period was about 2 months; every two weeks, the newly submitted applications were be investigated. Participants that could convincingly demonstrate their suitability for the workshop, were immediately accepted. Applicants that were clearly not suitable were rejected (actually, there was no such case). Applicants that seemed in principle suitable but whose application was not very strong, were put on hold. At the end of the application period, the on hold applicants were ranked and gradually admitted.

In each application the applicant ranked his preferences among the two alternative modules for each of the days 2-4. Preferences were considered as much as possible (subject to the maximum number of 25 participants per module); if a first preference could be granted, the other was offered as an alternative (ultimately, all preferences could be granted).

Statistics of the feedback survey

After the school, we triggered the electronic evaluation, which ran for two weeks and for which we received 25 responses (return quote: 66%).

As for the most important question "Overall: how would you rate this school?", we received 10 points from 24% of the responders, 9 points from 32%, 8 points from 20%, 7 points from 24% (total: 100%) with an average of **8.**56 points (from 10).

As for all questions:

1. Information about the school was…: excellent 52%, good 40%, rest 8%.
2. Registration was…: excellent 80%, good 12%, rest 8%.
3. The venue was …: excellent 72%, good 20%, rest 8%.
4. Catering was…: excellent 56%, good 32%, rest 12%.
5. The overall organization was…: excellent 80%, good 16%, rest 4%.
6. The topics were relevant for my work/research interests: agree completely 16%, agree 68%, rest 16%.
7. I was inspired to new ways of thinking: agree completely 36%, agree 36%, rest 28%.
8. The pace of teaching was right: agree completely 20%, agree 44%, rest 36%.

9. Teaching aids used (e.g. slides) were well prepared: agree completely 44%, agree 36%, rest 20%.
10. Hands-on exercises and demonstrations were a valuable contribution of the school: agree completely 32%, agree 32%, rest: 36%.
11. Social events: were an enjoyable and important aspect of the school: agree completely 44%, agree 28%, rest: 28%.
12. Overall, how would you rate this school (scale 0-10): 10 25%, 9 32%, 8: 20%, 7: 24%, rest: 0%.
13. In the future, I will need training in general HPC programming (MPI, OpenMP)…: urgent need 4%, important need: 36%, rest: 60%.
14. In the future, I will need training in advanced HPC programming(…): urgent need 20%, important need 28%, rest: 52%.
15. In the future, I will need training in code optimisation and performance analysis: urgent need 15%, important need: 36%, rest: 48%.
16. In the future, I will need training in the porting of existing code to HPC architectures: urgent need 4%, important need: 16%, rest: 80%.
17. In the future, I will need training in specific HPC applications: urgent need 0%, important need 8%, rest: 92%.
18. In the future, I will need training in HPC programming and applications specific to my research community…: urgent need 4%, important need 28%, rest: 68%.
19. In the future, I will need training in visualisation techniques…: urgent need 12%, important need: 8%, rest: 80%.

Analysis of the feedback

In the questionnaire, some general comments stated "This is also without a doubt the best organised workshop/conference I've attended.", "The organization was top! And the idea of choosing between two courses for three days is great as well.", "I very much liked the content of all lectures", "Thank you! I am grateful for this opportunity, and have enjoyed my visit", "It was a well organized, enjoyable and useful school". Details about more specific questions are given below.

Further closed questions asked and responses were:

1. Information about the school was ...: 92% responded "excellent" or "good".
2. Registration was ...: 92% responded "excellent" or "good".
3. The venue was ...: 92% responded "excellent" or "good".
4. Catering was ...: 88% responded "excellent" or "good".
5. The overall organization was ...: 96% responded "excellent" or "good".
6. The topics were relevant for my work/research interests: 84% responded "agree completely" or "agree".
7. I was inspired to new ways of thinking: 72% responded "agree completely" or "agree".
8. The lectures were clearly presented and comprehensible: 71% responded "agree completely" or "agree".
9. The pace of teaching was right: 64% responded "agree completely" or "agree".
10. Teaching aids used (e.g. slides) were well prepared: 80% responded "agree completely" or "agree".
11. Hands-on exercises and demonstrators were a valuable contribution to the school: 64 % responded "completely agree" or "agree".

12. Social event(s) were an enjoyable and important aspect of the school: 72% responded "completely agree" or "agree".
13. In the future, I will need training in general HPC programming (MPI, OpenMP): 40% reported "urgent need" or "important need", 36% reported "some need".
14. In the future, I will need training in advanced HPC programming (hybrid OpenMP-MPI, next-gen HPC languages e.g. PGAS, GPU computing e.g. CUDA): 56% reported "urgent need" or "important need", 20% reported "some need".
15. In the future, I will need training in code optimization and performance analysis: 52% reported "urgent need" or "important need", 40% reported "some need".
16. In the future, I will need training in porting of existing codes to HPC architectures: 20% reported "urgent need" or "important need", 44% reported "some need".
17. In the future, I will need training in specific HPC applications: 8% reported "urgent need or important need", 36% reported "some need".
18. In the future, I will need training in HPC programming and applications specific to my research community: 32% reported "urgent need" or "important need", 44% reported "some need".
19. In the future, I will need training in visualisation techniques: 20% reported "urgent need" or "important need", 40% reported "some need".

The answers to the open questions can be summarized as follows:

1. What did you like most about the course? Organization and venue was highly praised, the form of the school in two parallel strands was very much appreciated, the mix of disciplines, hand-on sessions, and interaction with participants was liked.
2. What did you like least about the course? The quality of WLAN in one of the two seminar rooms was criticized twice. One participant complained about technical problems in one of the hand-on sessions, one participant complained about too few hands-on sessions (or none at all) in some lectures.
3. Additional comments on the contents, specific lectures, etc. One participant praised the Xeon Phi session and the first day's discussion panel, one liked all the lectures, one wished a more detailed walk-through of the exercises with the lecturers, one complained about lack of preparation of some lecturers for the practical sessions.
4. Are there some other fields of training you feel PRACE should provide training events in? Individual wishes were listed for "medium performance computing" (more focus of HPC to smaller clusters), CUDA, data-parallel and domain-specific languages, large scale data analysis, partial differential equations.
5. Please give any other general comments about PRACE training activities: one recommendation was given to also attract people from the programming languages community.

9.2.5 *Conclusions & lessons learned*

We generally feel that the event was a big success and well received by the applicants. A difficult problem was the high number of late/last minute cancellations, i.e., applicants who after receiving an acceptance notification (and sometimes even after initial confirmation) withdraw their application. For future events, we would recommend to keep the registration longer open (even if all places are initially already filled) to make up for those who ultimately withdraw their application.

9.3 Report on PRACE Winter School 2016/2017

9.3.1 *Basic information about the event*

Name: PRACE 2017 Winter School-Fuelling Scientific Discovery with HPC Infrastructure
Dates: February 6-9, 2017
Location: Tel Aviv University Campus, Tel Aviv, Israel
Organizing sites: Inter-University Computation Center (IUCC)
Event URL: https://events.prace-ri.eu/event/549/

9.3.2 *Organizational details*

Local organizing committee

Hank Nussbacher	IUCC, Chair
Audrey Gerber	IUCC
Einat Ben-Shushan	IUCC

Venue
The training sessions took place in the Porter School for Environmental Studies on the Tel Aviv University campus, in close proximity to IUCC offices. The smaller session of the last parallel session of the program took place in the IUCC conference room. The building provided a comfortable classroom suitable for up to 50 students. Parking facilities, as well as the all necessary audio-visual facilities were provided.

Network connectivity was provided by the Tel Aviv University Computation Center. Guest access to HPC resources via the IUCC clusters was arranged prior to the opening of the event and user names and passkeys were delivered to participants when the sessions opened.

The selection of the Campus venue proved crucial to the success of the school since the local organizers were supported by IUCC administration and the University computation centre. The location of the campus, at the entrance of the city and adjacent to a commuter train station, also enabled easy access for local trainees and lecturers and visitors from abroad those staying in Tel Aviv hotels and Tel Aviv University guest apartments.

9.3.3 *Program & Content*

Program committee
Dr. Joan Adler, Technion-Israel Institute of Technology
Dr. Avi Cohen, Bar-Ilan University
Dr. Boris Morose, Afeka College
Dr. Amir Natan, Tel Aviv University
Dr. Guy Tel-Zur, Nuclear Research Center of the Negev's (NRCN) and Gurion University of the Negev
Mr. Hank Nussbacher, IUCC

Final program

	Morning	**Afternoon**
6/2/2017	Intro+ MPI +Open MP	Parallel I/O
7/2/2017	Parallel I/O (cont.)	Performance Analysis and Tuning

8/2/2017	GPU	Intel Xeon Phi
9/2/2017	GPU algorithm design	a. LAMMPS
		b. Multiprocessing in Python

List of trainers
Mr. Mordechai (Moti) Butrashvily, Tel Aviv University
Mr. Sebastian Lührs, Jülich Supercomputing Centre
Dr. Jan H. Meinke, Jülich Supercomputing Centre
Dr.-Ing. Bernd Mohr, Jülich Supercomputing Centre
Prof. Dan Mordehai, Technion-Israel Institute of Technology
Prof. Dennis Rapaport, Bar-Ilan University
Dr. Guy Tel-Zur, Nuclear Research Center of the Negev (NRCN), Ben-Gurion University

Designing the Program
The program committee consisted of experts from the local HPC academic community, computational scientists close to the true need of local and regional users and computational scientists who had participated in a PRACE "train the trainer" program. The program was designed to accommodate trainees with varying expertise level and professional backgrounds and to give a very thorough overview and training in fundamental HPC topics and to introduce topics of importance and interest to the computational science community targeted. Input was also solicited from younger students who had participated in a PRACE training program and was in tune with the needs of the geophysics community.

Description of content
The topics covered various parallel programming techniques
Intro+ MPI +Open MP
Frontal lectures and hands-on exercises on parallel computing in distributed and shared memory systems
Parallel I/O
This session opened with a general overview about parallel I/O strategies, parallel I/O bottlenecks, followed by an introduction to MPI-I/O (or SIONlib), including a hands-on exercise. The second session included an introduction to parallel, portable data formats with HDF5, including a hands-on exercise
GPU
This session focused on GPU programming and included a hands-on exercise, as it differs more from traditional shared memory programming than the models usually used for programming the Xeon Phi.
Intel XEON Phi
Lecture and hands on exercise
Performance Tools
This session opened with an introduction on methodology, followed by introducing a well-known portable, free, open-source tool set and closed with a one hour hands-on exercise
GPU Algorithm Design
Lecture and hands-on exercise on GPU architecture and programming, algorithms, and methods for adapting algorithms for efficient GPU utilization.
LAMMPS
Lecture and tutorial on using the LAMMPS molecular dynamics simulator.
Multiprocessing in Python

An introduction and hands-on exercises in multiprocessing in python both at node level and cluster level and technologies in Python that allow scientists to parallelize work on a cluster and multi-core (also in MPI), compared to traditional tools in C or Fortran.

Computer resources
The school made use of the IUCC infrastructure facilities for training purposes. Specifically, the school participants were given accounts to two different clusters:
1. Compute Nodes (Total 55 nodes):

CentOS release 6.6
8 cores (2CPUs with 4 cores each) - Intel(R) Xeon(R) CPU E5472 @ 3.00GHz
64GB ram
1TB local storage
2. GPU cluster: NVIDIA Tesla K40c with 2880 CUDA Cores

Social events
A welcome dinner and tour was held at the Milk & Honey Distillery in Tel Aviv on Monday, February 6, 2017 for the overseas students, members of the PRACE 2017 Winter School Program Committee and IUCC staff involved in organisation and technical infrastructure. A half day historic guided tour of Tel Aviv took place on Thursday afternoon, February 9, 2017 for the three visiting faculty members from the Jülich Supercomputing Centre, accompanied by a staff member from IUCC.

9.3.4 *Participants & feedback*

Overview
Due to capacity restraints of the classroom booked, registration was closed within two weeks of opening. We received a total of 75 applications which was narrowed down after requiring firm confirmation and cancelations. A waiting list was maintained during this period with participants contacting our office up until the evening before program opened requesting permission to attend, as well as two days into the program. Classroom capacity allowed for up to 50 participants for each session.
We were both delighted and surprised by the overwhelming response, which far exceeded our expectations based on much lower interest in 2014. We regretted that the university's lack of flexibility prohibited us from securing and confirming larger facilities quickly enough to respond to this demand. The number of training participants by country is as follows. For a full listing of participants please see Appendix D.

Country	Number of
Israel	54
Austria	1
Italy	1
Spain	1
Turkey	1

Total	58

There were another three applicants from overseas (Finland, Cyprus and Germany). The applicant from Cyprus did not receive budgetary approval from his university; the applicants from Finland and German cancelled their application.

Process for selecting the participants
The school was originally planned to host up to 50 students from Israel, Europe and the Eastern Mediterranean. All registrants presented their qualifications and background clearly and all were overwhelmingly qualified and suited to the program. There was no need to be selective and registrants were accepted on a first-come, first serve basis.

Statistics of the feedback survey
The feedback survey was completed by 32 school participants who answered the following 10 questions:

1. Name and Institution
2. E-mail
3. Where did you hear about the PRACE 2017 Winter School?
4. Area of work/research?
5. Your interest in HPC is as:
    a. A Researcher in field that uses HPC
    b. Programmer
    c. Computer Scientist
    d. System administration/maintenance
    e. User support
6. Please rate the following:
    a. Information about the school
    b. Registration
    c. Venue
    d. Catering
    e. Overall organization
7. Please rate the following:
    a. The topics are relevant to my work/research interests
8. Please rate the following:
    a. The lectures were clearly presented and understood
9. The hands-on exercises and demonstrations were a valuable contribution to the school
10. Overall, how would you rate the PRACE 2017 Winter School?

The full results can be found in Appendix E.

Analysis of the feedback
We were encouraged by the relatively large number of respondents to the feedback survey – over half. The participants, were in general, quite eager to share their impressions and offered many useful comments relevant to their areas of interest, and their HPC training needs.
The responses though were as diverse as the student population. The school hosted researchers, computer engineers and programmers, system administrators and support staff. It is difficult to meet the needs of so many in a single program and this was evident in the responses. Computer scientists would have liked to receive more content on algorithm design and coding in this environment; researchers in various disciplines were interested in ore user-specific content.

And the many support and system administrative staff expressed a desire for more content relevant to their needs. The overwhelming high registration for the tutorial on multiprocessing in Python was noted and this tutorial received extremely high ratings for relevance and for the quality of its hands-on exercises.
In general though the participants expressed satisfaction with the quality and level of the content. One conclusion drawn from this is the need to try and promote topic-specific PRACE training events, despite the challenges of potential participants may face in funding travel abroad.
In terms of registration, information and organization, all responded with satisfactory to excellent ratings, with no comments or suggestions. The social event were rated similarly.
The overall impression for the school was rated quite positive, and even more so appreciative to have a local option.
Responses to questions about the necessity for future schools covering the topics of the school indicated that the school provided training on topics of great interest and usefulness to the community targeted. All this will of course be taken into account on future training activities undertaken by IUCC.

Awareness Activities/Outreach
The registration process included a query as to where applicants heard about the program, which helped us determine the efficacy of the communications channels deployed to reach potential students. Direct push notifications via relevant university departments and emails was the predominant selection, proven by the fact that, as stated previously, registration was filled within a very short time of the initial announcement sent to university departments that registration had opened.

9.3.5 *Conclusions & lessons learned*

The organization of the PRACE 2017 Winter School was a major undertaking for IUCC. Although we had done it once before, we do not manage and organize conferences of this magnitude and scope on a regular basis. We again felt fortunate to be able to rely on PRACE colleagues from other countries where seasonal schools were hosted for advice and practical solutions. Their willingness to help was most appreciated. PRACE support was essential in terms of the provision of the online registration, and of course, the pool of trainers and PRACE promotional materials much appreciated by the participants.

Our expectations to receive applicants from outside Israel were lower than the actual results. Program experience shows that seasonal schools generally attract local and regional applicants. Receipt of seven applications and four confirmations from overseas, and from locations in Northern Europe, exceeded our estimates. Original applications indicated oversubscription. Session by session analysis allowed us flexibility to reach full capacity each day, with very low overall “no show” of **only** three applicants. Requiring firm confirmation allowed us flexibility and we were able to accommodate several waiting list and several last minute applicants.

Securing trainers is a crucial part of the organization process and we note that the streamlined PRACE policies and procedures made the logistics of inviting top flight trainers from the Jülich Supercomputing Centre smooth and efficient.

9.4 Report on Code Modernisation for Intel Multi Core and Xeon Phi Architectures on demand event

9.4.1 *Basic information about the event*

Name: Code Modernisation for Intel Multi Core and Xeon Phi Architectures
Dates: 25 - 28 April 2016
Location: Institute of Information and Communication Technologies, Bulgarian Academy of Sciences, Acad. Georgi Bonchev str., Bl.25A 1113 Sofia, Bulgaria
Organizing sites
https://events.prace-ri.eu/event/571/
http://www.prace-ri.eu/code-modernisation-for-intel-multi-core-and-xeon-phi-architectures/
http://scc.acad.bg/ncsa/hpctraining/

Organizational details
Institute of Information and Communication Technologies, Bulgarian Academy of Sciences, Sofia, Bulgaria, Training Hall #2 – The Hall provides excellent facilities especially designed to accommodate training courses, lectures, labs, etc in the field of Computer Science. It can easily accommodate 30 trainees in a classroom style setting. The premises also provide enough space for registration and coffee breaks. In addition IICT's location is close enough to Sofia airport (5km) and there is a number of 3 and 4 star hotels within 1km of the training location.

9.4.2 *Program & content*

Program committee
Ilian Todorov, PhD (STFC Daresbury Laboratory)
Alin-Marin Elena, PhD (STFC Daresbury Laboratory)
Valentin Pavlov, PhD (RILA Solutions)

Final programme

Monday, 25 April 2016
09:30 - 10:15 Three Questions Everyone Keeps Asking - Stephan Blair – Chappell **10:15 - 11:00 Vectorisation Part 1 (includes lab) - Stephan Blair –Chappell** **11:00 - 11:15 Coffee break** **11:15 - 13:00 Vectorisation Part 2 (includes lab) - Stephan Blair –Chappell** **13:00 - 14:00 Lunch** **14:00 - 15:00 Case Study - Alin-Marin Elena** **15:00 - 15:15 Coffee break** **15:15 - 16:15 Vectorisation Part 3 (includes lab) - Stephan Blair –Chappell** **16:15 - 16:30 Wrap-up-Day 1**
Tuesday, 26 April 2016
09:00 - 09:30 Producing Parallel Code - Stephan Blair –Chappell **09:30 - 11:00 Profiling Parallel Programs with VTune (includes lab) - Stephan Blair –Chappell** **11:00 - 11:15 Coffee break** **11:15 - 13:00 Architectural analysis with VTune (includes lab) - Stephan Blair –Chappell** **13:00 - 14:00 Lunch**

14:00 - 15:00 Case Study 2 – Michael Seaton **15:00 - 15:15 Coffee break** **15:15 - 17:15 Programming with MKL (includes lab) - Alin-Marin Elena**
Wednesday, 27 April 2016
09:00 - 09:30 Preparing for KNL - how to check your code without having real hardware (includes lab) **09:30 - 11:00 Profiling Parallel Programs with VTune (includes lab) – Victor Gamayonov** **11:00 - 11:15 Coffee break** **11:15 - 12:15 Case Study 3** **12:15 - 13:00 Lunch** **13:00 - 17:30 Afternoon for Informal Networking or sight seeing in city centre (make your own arrangements)**
Thursday, 28 April 2016
09:00 - 09:30 Introduction to the Code Dungeon **09:30 - 10:30 Code Dungeon Part 1** **10:30 - 11:00 Coffee break** **11:00 - 12:30 Code Dungeon Part 2** **12:30 - 13:00 Code Dungeon Part -what have we discovered so far?** **13:00 - 13:45 Lunch** **13:45 - 14:30 Code Dungeon Part 3** **14:30 - 15:00 Coffee break** **15:00 - 16:00 Code Dungeon Part 4** **16:00 - 16:30 Code Dungeon Pwrap-up**

List of trainers

Stephan Blair –Chappell (Intel/Bayncore):–compiler/architecture specialist;
Victor Gamayonov (Intel) – vTune/architecture specialist;
Dr. Alin-Marin Elena, PhD (STFC Daresbury Laboratory) – Research Software engineer; Dr. Michael Seaton PhD (STFC Daresbury Laboratory) – Research Software engineer;
Dr. IlianTodorov PhD (STFC Daresbury Laboratory) – Research Software engineer;
Dr. Peicho Petkov PhD (National Centre for Supercomputing Applications)

Designing the program

The school was aimed at final year master students, PhD students, and young researchers in computational sciences and engineering, interested in applying the emerging technologies on high performance computing to their research.

Description of the contents

The school focused on software modernization techniques needed for the next generation of supercomputers with highly dense parallel architectures, both homogeneous (Intel Xeon) and hybrid with acceleration co-processor (Intel Xeon Phi). The school program is comprised of lectures and training exercises to address the crucial aspects of both the implementation of new HPC applications as well as the re-factoring of existing ones. These software engineering techniques for high productivity languages complement the more traditional lectures on parallel programming, to allow the implementation and continual modernization of applications that need to be maintained across complex and fast evolving HPC architectures.
Computer resources (listing: fluid text: comment if they were sufficient for the event):
Two INTEL XeonPhi servers, dedicated network switch and training laptops provided by INTEL, The resources were sufficient for the training purposes

9.4.3 *Participants & feedback*

Number of participants by country

Participants	Country
19	**Bulgaria**
1	**Macedonia**

1. Information about the school was...

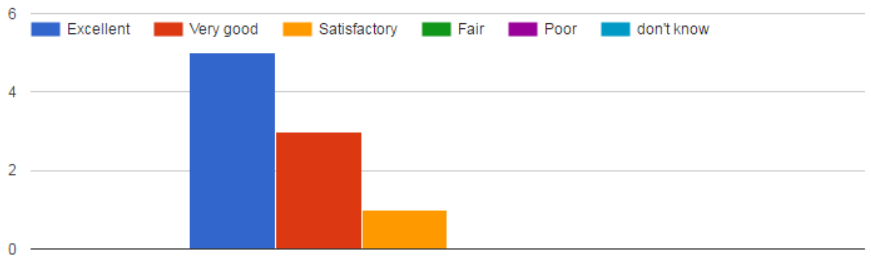

2. Registration was...

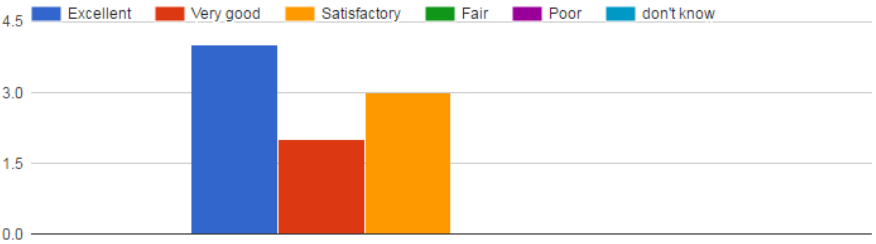

3. The venue was...

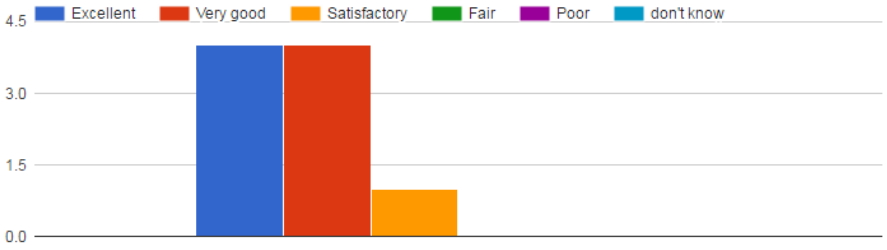

4. The overall organization was...

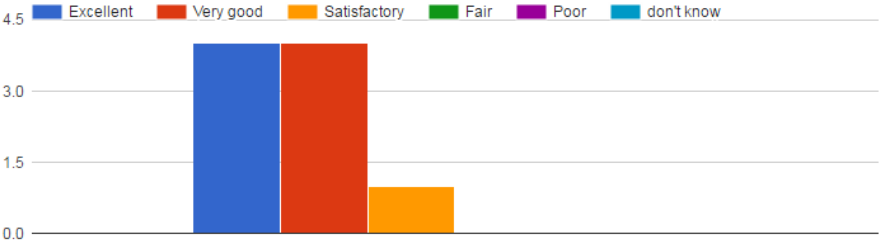

5. The topics were relevant for my work/research interests

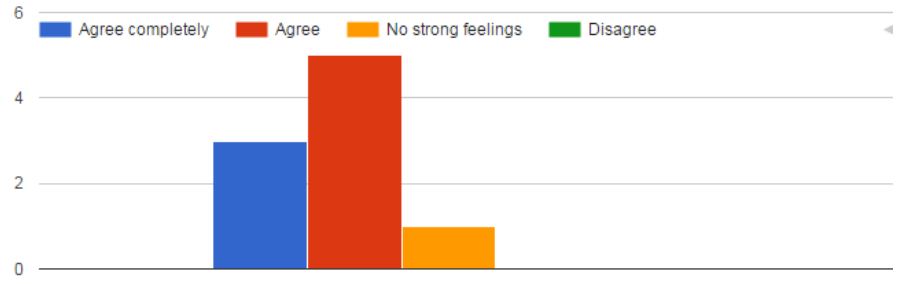

6. The lectures were clearly presented and comprehensible

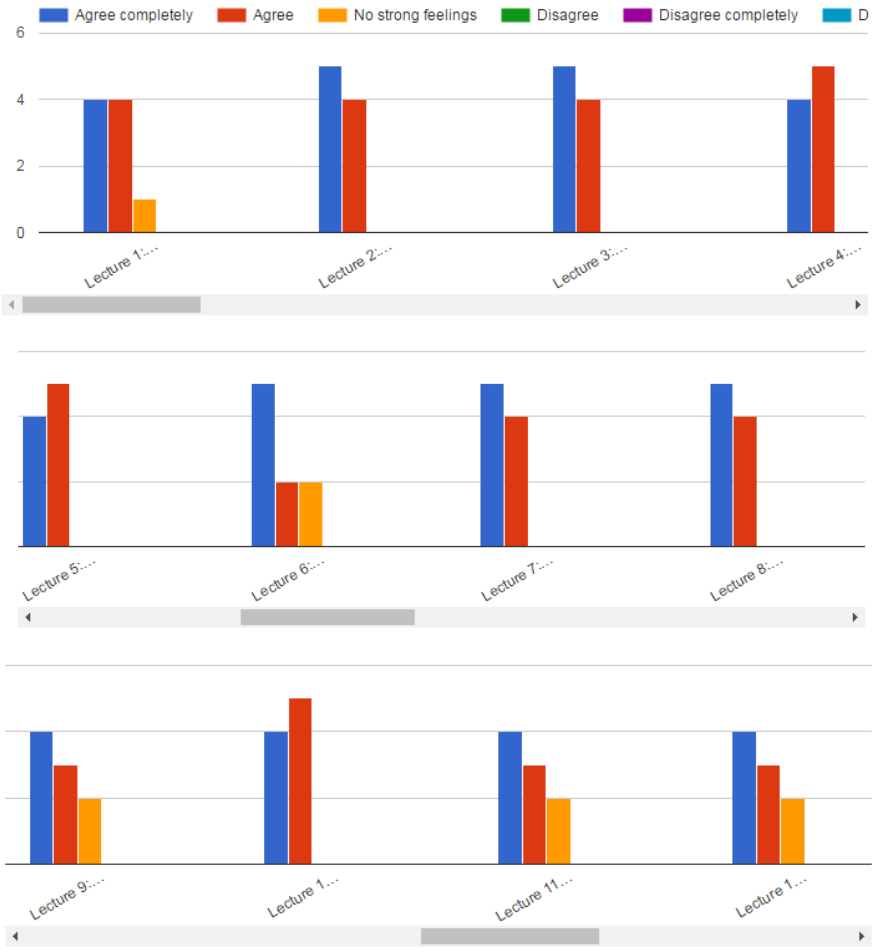

7. Hands-on exercises and demonstrations were a valuable contribution to the school

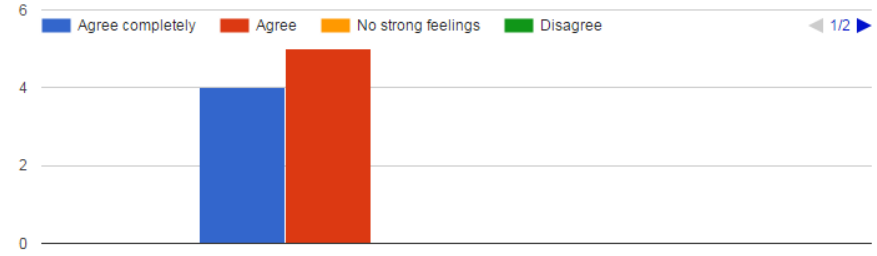

8. What aspects of this course were most useful or valuable? (9 responses)

Lecture topics, practical demonstrations and hands-on exercises.
vectorization
Lectures
The most useful was the hands on part. The chance to work with a real programs and hardware.
I do not remember such aspect.
information about options of intel compiler
VTune profiling
Programming with MKL
practical

9. Additional comments on the content, specific lectures, etc. (2 responses)

No comments
no

10. Overall, how would you rate this school?

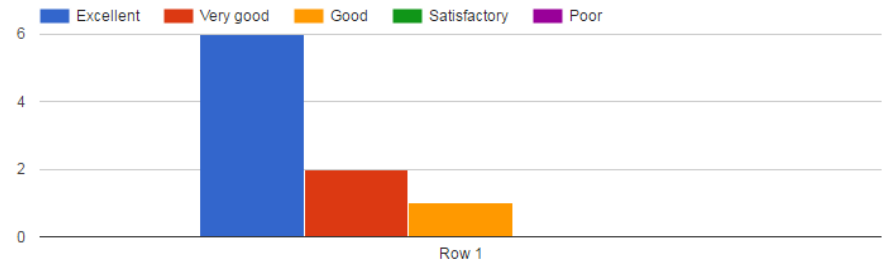

11. In the future, I will need training in

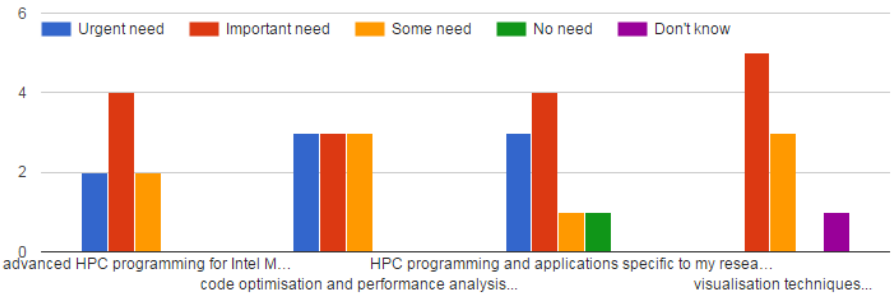

12. Please give any other general comments about PRACE training activities. (2 responses)

Well Done.
No comments

Analysis of the feedback

Over 80% of the feedback given on both the venue/organization, training quality and relevance questions falls in the “very good” to “excellent” (or “agree”-“completely agree”) range. Most importantly only one participant has rated the overall quality of the school as lower than very

good, which only confirms the positive remarks were given by the participants informally after the end of the programme.

9.4.4 *Conclusions & lessons learned*

The school was very popular and received very well which necessitated the organisation of a second school to affirm the newly trained skills, expand them and train further skills by targeting better the attendees' demand. It is worth mentioning that among the participants were researchers with academic record (from BAS, Technical University – Sofia), as well as university students (from TU-Sofia, Sofia University - Faculty of Physics) and representatives of software companies which goes to show the broad interest that the School managed to respond to.

9.5 Report on Parallel Programming with MPI and OpenMP for Beginners, PATC-on-demand course for project CoeGSS on demand event

9.5.1 *Basic information about the event*

Name: Parallel Programming with MPI and OpenMP for Beginners, PATC-on-demand course for project CoeGSS.

Dates: Oct. 17-18, 2016

Location: HLRS, University of Stuttgart, Seminar room, Allmandring 30, 70569 Stuttgart

Organizing sites: HLRS, University of Stuttgart, GCS PATC

9.5.2 *Organizational details*

Local organizing committee: Course organizer: Rolf Rabenseifner

Venue: HLRS, University of Stuttgart, Seminar room, Allmandring 30, 70569 Stuttgart. The room is the same as for the subsequent PATC course on "Advanced Parallel Programming with MPI and OpenMP" and has space for 60 participants.

Synergetic events: The course was provided in parallel with an HLRS course for many other participants. The course was held on Monday+Tuesday, directly before the advanced PATC course on Parallel Programming, in the same room. Most participants attended both courses, i.e. could reduce travelling costs to only one travel.

9.5.3 *Program & content*

Program committee: The course program was designed by Rolf Rabenseifner.

Final program

1st Day

08:30 Registration

09:00 Introduction

09:15 Parallel Architectures and Programming Models

10:15 Coffee break

10:30 Message Passing Interface (MPI), Introduction
10:40 Overview
11:15 Process Model
11:55 Process Model (practical)
12:45 Lunch
13:45 Messages and Point-to-point Communication
14:25 Messages and Point-to-point Communication (practical)
15:15 Coffee break
15:30 Nonblocking Communication
16:15 Nonblocking Communication (practical)
16:45 Coffee break
17:00 Collective Communication
17:45 Collective Communication (practical)
18:00 End

2nd Day

08:30 Shared Memory Parallelization with OpenMP
08:30 Overview
09:00 Execution model
09:30 Execution model (practical)
09:50 Coffee break
10:05 Worksharing directives
10:55 Worksharing directives (practical)
11:20 Coffee break
11:35 Data environment
11:55 Data environment (practical)
12:05 Summary
12:25 Pitfalls
12:45 Lunch
13:45 Pitfalls (continued)
14:35 Coffee break
14:50 Verifying an OpenMP Parallelization with the Intel Inspector XE
15:25 Verifying an OpenMP Parallelization with the Intel Inspector XE (practical)
15:50 Heat example (homework)
16:10 Coffee break

16:25 OpenMP-4.0 Extensions

17:30 Computer room tour and CAVE visualization demo

18:30 End

The complete agenda of the whole 5 day course (including the 3 PATC days), see https://fs.hlrs.de/projects/par/events/2016/agenda_2016-PAR.pdf

List of trainers:

Rolf Rabenseifner: All lectures on MPI and OpenMP. He is member of the MPI Forum (the MPI standardization body)

Uwe Wössner and Martin Aumüller: CAVE visualization demo

Designing the program: The CoeGSS on-demand course was designed as introductory part providing all the necessary skills to use MPI and OpenMP. It was also designed as prerequisite for the PRACE PATC advanced parallel programming course that was following the CoeGSS course. The total course week is designed to learn from the beginning most of MPI and OpenMP that is needed for an efficient use of HPC systems. The course is also designed to teach participants with very different previous knowledge. This course starts at 0% and tries to reach 100% of the needed parallelization skills for HPC.

Computer resources: The participants were grouped in groups of 2 persons for the exercises. Each group used on node on Hazel hen, the Tier-0 system at HLRS.

9.5.4 *Participants & feedback*

Number of participants by country: The CoeGSS course had 5 participants from Sweden (1), Italy (1), France (1), and Spain (2)

Process for selecting the participants: No selection, all could participate. Each participant can choose, which days she or he wants to book. From the in total 58 PATC, CoeGSS and HLRS participants, only 6 visited the only the first two days, and only 4 visited only the last 3 days. All 5 CoeGSS participants stayed also for the following three PATC days.

Statistics of the feedback survey: The CoeGSS participants were not separately asked. There is only a feedback for the whole course week by all 58 participants. We received 46 feedback forms with an average score of 8.1 (out of 0=worst … 10=best). The course length of the course was rated as adequate by 76%, too short by 9%, and too long by 15%. The depth of the of the course was rated as adequate by 76%, too superficial by 7%, and too profound by 17%. 97.6% of the participants would recommend this course to others.

Analysis of the feedback: Here only a few citations:

What did you like most about the course?

- Broad and deep view of the topics. Tips & tricks related to real-world application/usage.
- The course gives a nice description of parallelization techniques starting to basis concept up to the state of the art of this technology
- The fundamental description about the concepts and exercises
- Highly prepared speakers transmitting passion for the topics of the course. Many topics have been touched giving exposure to different parallelism paradigm. Exercises of the first 2/3 days were well suited and useful to the learning process.

- Beginners days lessons were well explained giving also some insights of interesting technical details. Course material was very well done.
- The high level of expertise of the teachers, and the level of detail of the courses. Also the practical exercises are well done, plus the access to the computer room was nice.
- This course covers a wide range of knowledge on the related topics.
- The exercises which come together with the theoretical lectures.

What did you like least about the course?

- I would like to have seen more practical examples
- A lot of overlap with the Node-level Performance Optimization course in the advanced topics.
- The exercise were good but often a bit too superficial when compared to real-world applications.
- In the last two days the peace of the course was too fast, with little practical examples. A lot of time has been spent on MPI-3 shared memory. This time could have been spent on practical exercises on parallel MPI IO and hybrid programming. Hybrid programming was not covered sufficiently and it could have been.
- there were no butter brezels

As expected, for some of the participants, the course is partially too profound. As seen from the statistics, the whole 5-day course (2 beginners CoeGSS/HLRS days and 3 advance PATC days) has a good balance in length and depth.

9.5.5 *Conclusions & lessons learned*

It is a good idea to combine a beginners on-demand course with a locally provided course for other participants and to combine such a beginners course also with an advanced PATC course. About 80% of the participants booked the whole week. Especially at the beginning of a PhD, they have a deep knowledge in their field, but need such a 0-to-100 start in the area of parallel programming.

9.6 Report on Performance Optimization and Productivity Tutorial (POP CoE) on demand event

9.6.1 *Basic information about the event*

Name: Performance Optimization and Productivity Tutorial

Dates: 14-15 December 2016

Location: VŠB - Technical University Ostrava, IT4Innovations building

Organizing sites

- VŠB - Technical University Ostrava – IT4Innovations
- Performance Optimisation and Productivity Centre of Excellence in Computing Applications

9.6.2 *Organizational details*

Local organizing committee (table)

1 **Name**	2 **Affiliation**
3 Ondrej Jakl	4 IT4Innovations
5 Lucie Valeckova	6 IT4Innovations
7 Hana Valouchova	8 IT4Innovations

Venue (name, its description and why it was selected; fluid text: analysis how good the selection was for the event)

The event took place at VŠB - Technical University Ostrava, IT4Innovations building, room 207. In this training room, almost all courses of IT4Innovations take place, and it was fully appropriate for this on-demand event, too.

Synergetic events:

No

9.6.3 *Program & content*

Program committee (table)

Name	**Affiliation**
Ondrej Jakl	IT4Innovations
Dirk Schmidl	RWTH Aachen
Brian Wylie	JSC

Designing the program

The EU funded Performance Optimization and Productivity Centre of Excellence (POP) offers services in the area of performance analysis and performance optimization. Therefore, POP mainly uses performance analysis tools developed by POP partners, namely Extrae, Paraver, Dimemas, Score-P, Cube and Scalasca. These tools are used to generate uniform performance audits for parallel applications to provide a good overview of the applications execution.

The aim of this tutorial, especially during its first day, was to introduce participants to the performance analysis tools and the methodologies used to create POP performance audits. The second day was devoted to large hands-on sessions where application developers, who brought their own code, were doing a performance analysis with guidance of POP performance experts.

The tutorial was in some sense a follow-up event after the course Practical Parallel Performance Analysis on Salomon held at IT4Innovations on October 20-21, 2016, where the most participants were pre-selected.

Description of the contents

- Paraver / Extrae / Dimemas (Judit Gimenez, German LLort): practical introduction to the set of tools from BSC

- Introduction to POP services (Dirk Schmidl, Aamer Shah): POP CoE, partners, motivation, services, status, audits,...
- k-Wave audit results (Brian Wylie): example of a performance audit on an application of one of the developer teams
- Score-P library wrapping (Ronny Brendel): how user library wrapping can be achieved in Score-P?
- hands-on: individual work with 6 developer teams – one trainer per team

Computer resources

The main computing resource for the event was the Salomon supercomputer at the National Supercomputing Centre IT4Innovations, which is an SGI cluster deployed in 2015. After intensive preparatory work of the organizers, administrators and trainers, this platform could meet very well the requirements of the event, because of dedicated allocation of nodes and a high throughput queue.

9.6.4 *Participants & feedback*

Number of participants by country (table, full list of participants as an appendix)

Country	**9 Number of participants**
Czech Republic	21
Poland	2
Germany	4 (trainers)
Spain	2 (trainers)

See Table 2 in the Appendix for the full list of participants.

Process for selecting the participants

The participants were pre-selected during the course Practical Parallel Performance Analysis on Salomon held at IT4Innovations on October 20-21, 2016, based on a Call addressing application developers interested in a performance audit of their codes. Six teams (3 from IT4Innovations, 1 from Brno University of Technology, 1 from Institute of Geonics CAS, and 1 from Czestochowa University of Technology) positively answered. This number was precisely adequate to the capacity of the course and its six trainers.

Statistics of the feedback survey

The participants were encouraged to fill in a standard questionnaire for IT4Innovations training events. Its results (8 replies) are presented in Table 3 in the Appendix.

Analysis of the feedback
All eight participants who provided their opinion were satisfied with the event organization (all fully agreed with the statement "The event was well organized").

The "Overall benefit of the training for you" has got a good average mark of 8.38 on the scale (0 - waste of time ... 10 - full satisfaction). The only comment obtained was "Projector should be placed on the ceiling. It was placed on a stand where it obstructed the bottom part of the presentation from certain places", which is true.

9.6.5 *Conclusions & lessons learned*

On-demand event pilot, as introduced in PRACE-4IP, proved to be a very useful complement of the traditional Seasonal Schools in our eyes. It has great knowledge dissemination effect and helps shaping the local/national HPC training infrastructure through collaboration with and service for another bodies.

9.7 Report on Material Science codes on innovative HPC architectures: targeting exascale (MaX CoE) on demand event

9.7.1 *Basic information about the event*

Name: Material Science codes on innovative HPC architectures: targeting exascale
Dates: Dec 5th -7th 2016
Location: Cineca - Casalecchio di Reno (BO, Italy)
Organizing sites: Cineca-CSCS-MaX

9.7.2 *Organizational details*

Local organizing committee: Cineca - CSCS - MaX, the European center of excellence
Venue: The venue selected was the Cineca courses classroom, well equipped with PCs for the exercises session, connected with internet (so the connection with Cineca cluster was allowed). All the students who subscribed were admitted since the classroom had enough spots available and the trainers were in a number sufficient to guarantee a high level of support.
Synergetic events: NO

9.7.3 *Program & content*

Program committee
The course was focused of the main codes used by the material science community (namely, Quantum Espresso, Siesta, Yambo, FLEUR plus the AiiDA toolchain manager) introducing their architecture and usage. The course explained how the codes can exploit petascale systems, targeting in particular to PRACE Tier-0 systems, and how they are preparing for the exascale era, with specific attention to hybrid and accelerated architectures.

Final program
Introduction to Exascale and Tier-0 systems.
Domain specific libraries
Quantum ESPRESSO: parallelization levels and scalability
Hands-on of QE parallelization
QE+LAMMPS presentation + tutorial
AiiDA presentation
Yambo: presentation + hands-on

SIESTA: presentation + hands-on
FLEUR: presentation +hands-on
List of trainers;
Carlo Cavazzoni
Anton Kozhenikov
Fabio Affinito
Mariella Ippolito
Nicola Spallanzani
Davide Sangalli
Oswaldo Dieguez
Daniel Wortmann

Designing the program
The course has been organized in collaboration between CINECA and CSCS focusing on several of the main codes used by the material science community (namely, Quantum Espresso, Siesta, Yambo, Fleur), introducing their architecture and usage. The course presents how the codes exploit petascale systems, targeting in particular to PRACE Tier-0 systems, and how they are preparing for the exascale era, with specific attention to hybrid and accelerated architectures. Hands-on have been proposed to let the attendees to get familiar with the codes and their optimal usage on HPC systems. The course addresses in particular the MAX community, more specifically the users of the selected codes and the software developers starting contributing to those codes with new components and functionalities targeting the support of novel architectures and the performance and scalability improvement.
The course has been planned on three days and has been delivered at CINECA in the period December 5-7, 2016. Teachers have been provided by the involved PRACE partners and by the MAX project members.

Description of the contents
- Introduction to Exascale and Tier-0 systems: a general overview of the state-of-the-art of HPC architectures and trends toward the exascale challenge;
- Domain specific libraries: an introduction about the DSL approach and how it can permit to tackle the modernization of scientific software;
- Quantum ESPRESSO parallelization levels and scalability: an overview of the parallelization strategy of the Quantum ESPRESSO application and how it can permit to exploit HPC architectures;
- Hands-on of QE parallelization: all-hands lab on the topics of the previous lecture;
- QE+LAMMPS presentation: the interoperability between the MD engine of LAMMPS and the QE SCF algorithm is presented with examples;
- AiiDA overview: the AiiDA environment for large scale HTC simulation in material science is presented;
- Yambo presentation + hands-on: algorithms and parallelization strategies of Yambo with exercises;
- SIESTA presentation + hands-on: algorithms and parallelization strategies of SIESTA with exercises;
- FLEUR presentation + hands-on: algorithms and parallelization strategies of FLEUR with exercises.

Computer resources
The access to the supercomputer MARCONI-A1 was granted to all the attendees. The students were able to run on MARCONI the examples and the test codes proposed by the teachers.

9.7.4 *Participants & feedback*

Number of participants by country

Italy	15
Germany	1
Appendix	1

Process for selecting the participants
No selection was made.

Analysis of the feedback
The overall feedback is very positive. General evaluations show that the attendees liked the course, its organization and the subject of the lectures. In the suggestions, different thoughts were expressed, mostly according the personal background of the attendees. There are not indications for major corrections in the organization of this course, or actions that need to be undertaken to correct some issues.

9.7.5 *Conclusions & lessons learned*

This course presented a general overview of many different codes in the domain area of material sciences, with a particular concern to the exascale challenges. The course received a very positive feedback. In general, since the codes (and related algorithms) were quite different, some students could have been more attracted by some specific lecture than others. This was the first edition of this kind of course. In the future, we will probably need to make the content of the different lectures (i.e. related to different codes) more homogeneous.

9.8 Report on Practical Programming Models and Skills on INTEL Xeon for Scientific Research Engineers on demand event

9.8.1 *Basic information about the event*

Name Practical Programming Models and Skills on INTEL Xeon Phi for Scientific Research Engineers
Dates 22-24 March 2017
Location Institute of Information and Communication Technologies, Bulgarian Academy of Sciences, Acad. Georgi Bonchev str., Bl.25A 1113 Sofia, Bulgaria
Organizing sites: IICT, Intel

9.8.2 *Organizational details*

Venue
Institute of Information and Communication Technologies, Bulgarian Academy of Sciences, Sofia, Bulgaria, Training Hall #2 – The Hall provides excellent facilities especially designed to accommodate training courses, lectures, labs, etc in the field of Computer Science. It can easily accommodate 30 trainees in a classroom style setting. The premises also provide enough space

for registration and coffee breaks. In addition IICT's location is close enough to Sofia airport (5km) and there is a number of 3 and 4 star hotels within 1km of the training location.

9.8.3 *Program & content*

Program committee
Dr. Ilian Todorov, PhD (STFC Daresbury Laboratory)
Dr. Alin-Marin Elena, PhD (STFC Daresbury Laboratory)
Valentin Pavlov, PhD (RILA Solutions)

Final program

Wednesday, 22 March 2017
12:00 - 17:00 Day 1 **12:00-13:00 Registration** **13:00-15:00 Sergi-EnricSiso (STFC HC, IPCC) – Performance scaling and analysis in different modes of OSPRay: A Ray Tracing Based Rendering Engine for High-Fidelity Visualization;** **Tuning Scalability and Performance on Xeon and Xeon Phi** **15:00-15:30 Coffee** **15:30-17:00 Sergi-EnricSiso (STFC HC, IPCC) - OSPRay Continued** **17:00 Wrap Up End Day 1**
Thursday, 23 March 2017
09:00 - 09:15 Welcome to Day 2. **09:05 - 10:30 Aidan Chalk (STFC HC, IPCC) - Refresher course on use of parallelization techniques and their application on XeonPhi Programming models for Xeon Phi – a quick resume** **10:30 - 11:00 Coffee** **11:00 - 12:15 Aidan Chalk (STFC HC, IPCC) - Refresher course - continued** **12:15 - 13:30 Lunch** **13:30 - 15:00 Alin-Marin Elena/Ilian Todorov (STFC HC)** Performance scaling and analysis in different modes of DL_POLY: CCP5 flagship code for Molecular Dynamics **15:00 - 15:30 Coffee** **15:30 - 16:30 PeichoPetkov (NCSA-BG)** Performance scaling and analysis in different modes of LAMMPS **16:30 Wrap up Day 2**
Friday, 24 March 2017
09:00 - 13:30 Day 3 **09:00 Welcome to Day 3** **09:05 – 10:00 Alin-Marin Elena/Ilian Todorov (STFC HC) –DL_POLY Continued** **10:00-10:30 Coffee** **10:30-12:30 Francois Fayard (Bayncore) -** Practical skills and labs on profiling and vectorization using Parallel Studio and the Xeon Phi **12:30-13:30 Break** **13:30 Wrap up End Day3**

List of trainers

- Dr. Alin-Marin Elena, PhD (STFC Daresbury Laboratory,UK) – Research Software engineer, HPC Specialist and Consultant at SCD/Hartree Centre, Theoretical/Molecular Physicist with MD (DL_POLY) and ab initio MD (CP2K) expertise;
- Francois Fayard, PhD (Bayncore) – Application Mathematician and Software Engineer; HPC consultant
- Dr. Peicho Petkov PhD (NCSA/Sofia University,Bulgaria) – Theoretical Physicist
- Sergi-Enric Siso – (STFC Hartree Centre, UK) – Computer Scientist, HPC expert, Software engineer;
- Aiden Chalk (STFC Hartree Centre, UK) – Computer Scientist, HPC expert, Software engineer)

Designing the program

The program targeted the software engineers from the previous events held at the same place with plans to train the skills already obtained by those via demonstrating of how these are applied in a number of case studies whose various modes' performance were scaled on the NCSA's heterogeneous cluster Avitohol.

Description of the contents

The training school is organized by NCSA in assistance with Science and Technology Facilities Council (STFC) and Bayncore. This is a follow up of the series of training events at NCSA that will focus on training practical skills by demonstrations on how to test performance scaling on Xeon Phi of a few project codes.

The school was designed for research software engineers (RSEs) from industry, institutes and academia. It focused on training practical skills by demonstrations of how to test, analyse and improve performance scaling on INTEL Xeon Phi of a few project codes. These included presentations on OSRay, DL_POLY and LAMMPS as well as a number of simple example codes (analised by INTEL's vTune and modified in real time).

Computer resources

User accounts to Avitohol Supercomputer for the trainers and training accounts to Avitohol for the attendees; local wireless network. The resources were sufficient for the training purposes

9.8.4 *Participants & feedback*

Number of participants by country

Participants	Country
19	**Bulgaria**
1	**Romania**

Statistics of the feedback survey (tables (appendix))

Statistics

* 1. Information about the school was...

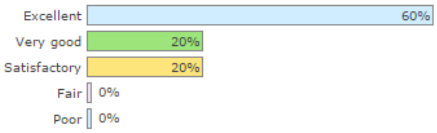

* 2. Registration was...

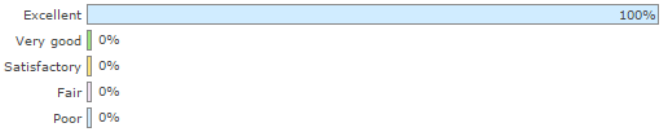

* 3. The venue was...

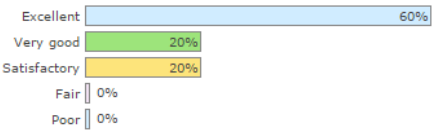

* 4. The overall organization was...

* 5. The topics were relevant for my work/research interests

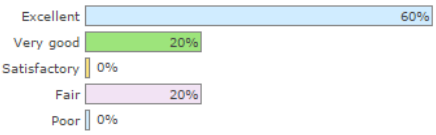

* 6. The lectures were clearly presented and comprehensible

* 7. Hands-on exercises and demonstrations were a valuable contribution to the school

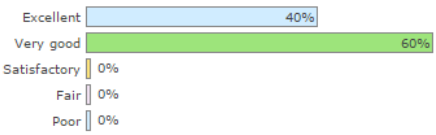

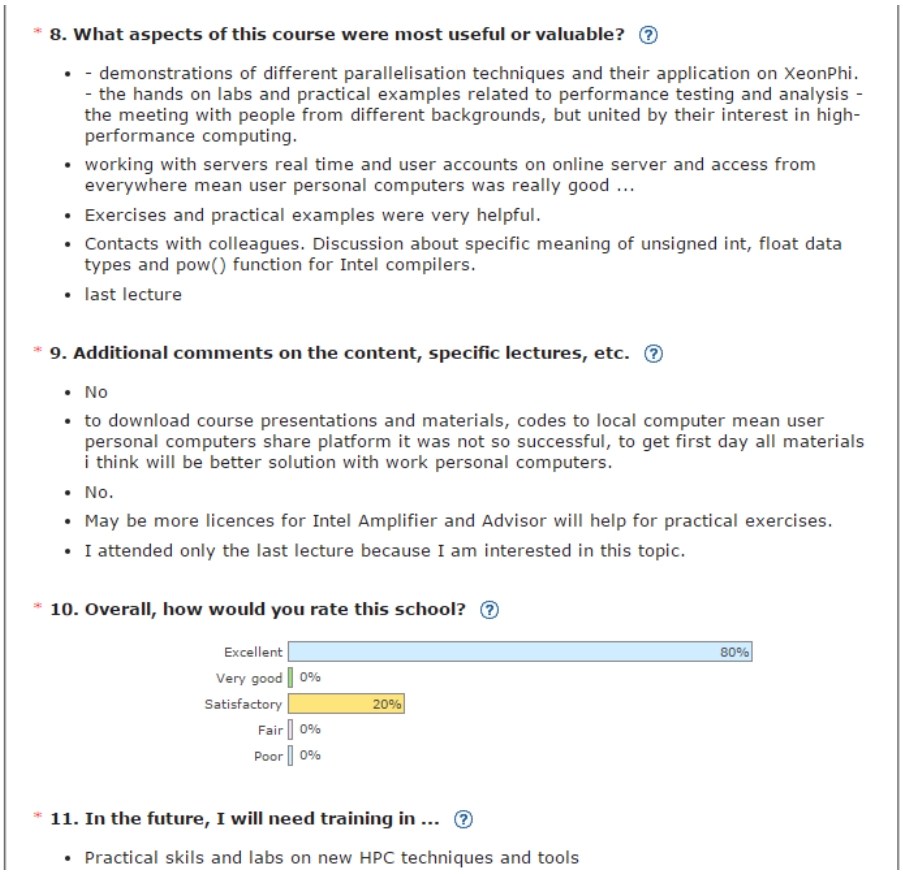

* **8. What aspects of this course were most useful or valuable?**

- - demonstrations of different parallelisation techniques and their application on XeonPhi. - the hands on labs and practical examples related to performance testing and analysis - the meeting with people from different backgrounds, but united by their interest in high-performance computing.
- working with servers real time and user accounts on online server and access from everywhere mean user personal computers was really good ...
- Exercises and practical examples were very helpful.
- Contacts with colleagues. Discussion about specific meaning of unsigned int, float data types and pow() function for Intel compilers.
- last lecture

* **9. Additional comments on the content, specific lectures, etc.**

- No
- to download course presentations and materials, codes to local computer mean user personal computers share platform it was not so successful, to get first day all materials i think will be better solution with work personal computers.
- No.
- May be more licences for Intel Amplifier and Advisor will help for practical exercises.
- I attended only the last lecture because I am interested in this topic.

* **10. Overall, how would you rate this school?**

* **11. In the future, I will need training in ...**

- Practical skils and labs on new HPC techniques and tools

Analysis of the feedback

The School received the best feedback among the courses organized by NCSA in the past year. 80% of the participants rated the overall quality as excellent and only an insignificant number of attendees rated any of the aspects as less than “very good”. Again the relevance and practicality of the hands-on labs received particularly high acclaim.

9.8.5 *Conclusions & lessons learned*

The workshop was attended by seasoned researchers and academics from BAS, the University of Technology (Sofia and Plovdiv), a few advanced students and a software engineer delegate from the Romanian Academy of Sciences. The workshop proved popular again and was received well by both the students from the universities and the researchers from the academic circles. The positive feedback given by the participants gives NCSA an incentive to continue organizing practical trainings for programming on heterogenous systems.

9.9 Report on International HPC Summer School 2016

9.9.1 *Basic information about the event*

Name: International HPC Summer School 2016

Dates: 26 June – 1 July 2016

Location: Ljubljana, Slovenia

Organizing sites: University of Ljubljana (UL), Slovenia

Mission: The mission of this annual series of summer schools, organised by PRACE and international partners (currently XSEDE, RIKEN AICS and Compute Canada), is to prepare a larger and more diverse international computational science and HPC workforce by expanding the knowledge of the attendees about HPC and its applications in multiple fields of science and engineering, and fostering new collegial friendships and partnerships (nationally and internationally) among the presenters and attendees.

Event URL: http://ihpcss2016.hpc.fs.uni-lj.si/

9.9.2 *Organizational details*

Local organizing committee

Jožef Duhovnik	**Committee chair**
Leon Kos	**Chief organisation and technical**
Mateja Maffi	**Financial**
Janez Krek	**Organisation**

Venue

The Faculty of Law building (University of Ljubljana) was chosen to host the school. It is one of the founding faculties of the University of Ljubljana and the oldest and largest law faculty in Slovenia. The premises have been recently been renovated prior to the event and provides a distinguished and creative academic setting for education at all levels. It contains all the necessary rooms and facilities for the teaching sessions of the school, as well as having appropriate space and catering facilities for breaks and meals. It is also within walking distance to accommodation organised by the event.

Program committee

The International HPC Summer School series typically utilises a combined organising/ programme committee that oversees both logistics and the programme of the school; it also includes members of the local organising committee. It is a large committee with many members from all partners (PRACE, XSEDE, RIKEN AICS, Compute Canada), although there are typically person(s) assigned for specific tasks such as speaker/programme development, mentoring activities, practical session preparations, evaluation, etc. The key members from PRACE in the organising committee (excluding the local organising committee members listed above) are:

Hermann Lederer	**RZG**
Simon Wong	**ICHEC**
David Henty	**EPCC**
Christian Feld	**JSC**

Final program

The final programme for the school can be found at: http://ihpcss2016.hpc.fs.uni-lj.si/agenda/

List of trainers

Orly Alter	University of Utah, USA
Matej Andrejašič	Pipistrel, Slovenia
Yuriko Aoki	Kyushu University, Japan
Galen Arnold	University of Illinois, USA
Ritu Arora	Texas Advanced Computing Center, USA
Philip Blood	Pittsburgh Supercomputing Center, USA
Vetria Byrd	Purdue University, USA
Scott Callaghan	Southern California Earthquake Center, USA
Thomas Cheatham (keynote)	University of Utah, USA
David Henty	EPCC at the Univeristy of Edinburgh, UK
Tomi Ilijaš	Arctur, Slovenia
Klaus Kolag	Max-Planck Institute for Astrophysics, Germany
Giovanni Lemanna	CNRS LAPP Particle Physics Lab, France
Erik Lindahl	University of Stockholm, Sweden
Takemasa Miyoshi	RIKEN AICS, Japan
Thomas Sterling	Indiana University, USA
John Urbanic	Pittsburgh Supercomputing Center, USA
Ramses van Zon	University of Toronto, Canada

Designing the program

The format of the school has remained largely the same over the years, but has always evolved slightly over the years based on evaluation from past years. The programme typically consists of a keynote lecture from a prominent speaker from the HPC and/or computational science domain, followed by a mix of scientific talks with focus on HPC challenges, practical hands-on session on HPC methodologies and techniques, as well as time for participants to interact with speakers/mentors and each other. This school's programme also included a coding competition for participants to take part in during the event. In addition, two of the invited speakers were from local industry in the HPC domain.

Computer resources

All participants were asked to bring laptops for the hands-on sessions. Course account access to both PRACE and XSEDE compute resources were provided to participants.

9.9.3 *Participants & feedback*

Number of participants by country

Belgium	1
Denmark	1
Finland	1
France	1
Germany	8
Greece	2
Hungary	1
Ireland	2

Italy	2
Serbia	1
Slovenia	1
Spain	4
Sweden	2
Switzerland	2
Turkey	1
UK	4
Canada	9
Japan	8
USA	28
Total	**79**

Process for selecting the participants

The effort in selecting the 80 participants was divided locally to the respective partner organisations, i.e. PRACE was responsible to fill its allocated 30 places from the 160 applications received from candidates based in European institutions. The European selection committee consisted of six scientists and HPC experts from France (CEA), Ireland (ICHEC), Italy (CINECA), Germany (RZG), Spain (BSC) and the UK (EPCC). Guidelines for the reviewing of applications have also been standardised across the partners (XSEDE, PRACE, RIKEN AICS, Compute Canada), where candidates were assessed based on:

- Strong scientific/technological reason for participation
- That the summer school can contribute to advancing their research goals
- Engagement in code development rather than the use of applications
- Experience with HPC programming

While the school attracted a total of 329 applications, PRACE selected 30 students to participate in the school out of 128 European applications. Another 4 European students were subsequently invited to the school after cancellations from outside Europe. Hence, there were a total of 34 European participants in the school, along with 45 participants from Canadian, Japanese and U.S. institutions.

Statistics and analysis of the feedback

A total of 74 feedback responses were received, representing a response rate of 94%. The following is a summary of the results in different categories:

- Overall: a general measure of the perceived quality of this event was the response to the statement "Overall I would rate my experience as successful"; to this 100% of respondents were in agreement that the event was a successful experience for them.
- Learning outcome: 95% of respondents indicated that their goals of attending the event were achieved; 85% stated that the skills they've learnt will significantly contribute to their research; 96% are aware of the next step to build on what they have learned.
- Organisation: 97% of respondents found the school to be well organised and 82% were satisfied with the delivery format of the school.

- International audience: 89% of respondents indicated that participation of students from other countries contributed to their learning; 93% have meaningfully engaged with other students at the school.
- Mentoring: 89% respondents have stated to have meaningfully engaged with their assigned mentor during the school.

Compute resources: 79% of respondents are interested in learning more about the resources/opportunities available through the partner organisations as a result of their experience; 69% plan on obtaining access to such resources after the school.

Conclusions & lessons learnt

The 2016 International HPC Summer School has been a hugely successful event. Building on experience from the past 7 years, the programme continues to deliver untold benefits to participants. Feedback has been overwhelmingly excellent for the school; there were particular praise for the tremendous efforts put into the event by the local organisers.