



**SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures**

**INFRA-2012-2.3.1 – Third Implementation Phase of the European
High Performance Computing (HPC) service PRACE**



PRACE-3IP

PRACE Third Implementation Phase Project

Grant Agreement Number: RI-312763

**D4.5
Final Report on Training Activities**

Final

Version: 1.0
Author(s): Martti Louhivuori, CSC; Stelios Erotokritou, CaSToRC; Ondrej Jakl, VŠB-TUO/IT4Innovations; Ioannis Liabotis, GRNET; Simon Wong, ICHEC
Date: 23.01.2015

Project and Deliverable Information Sheet

PRACE Project	Project Ref. №: RI-312763	
	Project Title: PRACE Third Implementation Phase Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: < D4.5 >	
	Deliverable Nature: <DOC_TYPE: Report / >	
	Deliverable Level: PU	Contractual Date of Delivery: 31 / January / 2015
		Actual Date of Delivery: 31 / January / 2015
EC Project Officer: Leonardo Flores Añover		

* - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	Title: Final Report on Training Activities	
	ID: D4.5	
	Version: <1.0 >	Status: <i>Final</i>
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2007	
	File(s): D4.5.docx	
Authorship	Written by:	Marti Louhivuori, CSC; Stelios Erotokritou, CaSToRC; Ondrej Jakl, VŠB-TUO/IT4Innovations; Ioannis Liabotis, GRNET; Simon Wong, ICHEC
	Contributors:	
	Reviewed by:	Veronica Teodor, FZJ; Pedro Alberto, UC-LCA
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	17/Dec/2014	Draft	Initial outline
0.2	07/Jan/2015	Draft	Include contributions
0.3	08/Jan/2015	Draft	Consolidated version
0.4	09/Jan/2015	Draft	For internal review
1.0	23/Jan/2015	Final version	Revisions added

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Training
------------------	---

Disclaimer

This deliverable has been prepared by the responsible Work Package of the Project in accordance with the Consortium Agreement and the Grant Agreement n° RI-312763. It solely reflects the opinion of the parties to such agreements on a collective basis in the context of the Project and to the extent foreseen in such agreements. Please note that even though all participants to the Project are members of PRACE AISBL, this deliverable has not been approved by the Council of PRACE AISBL and therefore does not emanate from it nor should it be considered to reflect PRACE AISBL's individual opinion.

Copyright notices

© 2015 PRACE Consortium Partners. All rights reserved. This document is a project document of the PRACE project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the PRACE partners, except as mandated by the European Commission contract RI-312763 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

Table of Contents

Project and Deliverable Information Sheet	i
Document Control Sheet.....	i
Document Status Sheet	i
Document Keywords	ii
Table of Contents	iii
List of Figures.....	iv
List of Tables.....	iv
References and Applicable Documents	iv
List of Acronyms and Abbreviations.....	v
Executive Summary	1
1 Introduction	1
2 PRACE Advanced Training Centres.....	2
2.1 PATC Courses from July 2014 to January 2015	2
2.1.1 <i>Key statistics on numbers of participants and courses</i>	2
3 Seasonal Schools	3
3.1 PRACE-3IP Extension Period Seasonal Schools' Characteristics	3
4 International HPC Summer School 2014	5
5 On-line services.....	5
5.1 Completion of PRACE 3IP Video Tutorials	6
5.2 Maintenance of WP4 on-line services	6
5.3 Migration of PRACE Events website.....	7
5.4 Google Analytics PRACE Training Portal Statistics.....	7
6 Conclusions	8
7 Annex.....	9
7.1 Report on PRACE Autumn School 2014.....	9
7.1.1 <i>Basic information about the event</i>	9
7.1.2 <i>Organizational details</i>	9
7.1.3 <i>Program & content</i>	10
7.1.4 <i>Participants & feedback</i>	12
7.1.5 <i>Conclusions & lessons learned</i>	14
7.2 Report on PRACE Winter School 2015.....	14
7.2.1 <i>Basic information about the event</i>	14
7.2.2 <i>Organizational details</i>	15
7.2.3 <i>Program & content</i>	15
7.2.4 <i>Participants & feedback</i>	19
7.2.5 <i>Conclusions & lessons learned</i>	19

List of Figures

Figure 1. Visitor statistics for the Training Portal during the full year of 2014.....	8
Figure 2. Overview of feedback results from PRACE Autumn School 2014.....	13

List of Tables

Table 1. Key statistics from implementation of the 2012-13, 2013-14, and partial 2014-15 PATC joint curricula. “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.	2
Table 2. All seasonal school participants and their geographic distribution.....	4
Table 3. PRACE-3IP Video Tutorials created by partners involved in the Training Portal task.	6

References and Applicable Documents

- [1] PRACE Training Portal:
<http://www.training.prace-ri.eu>
- [2] D4.2: “*Report on Training Events*”; PRACE-3IP public deliverable,
http://www.prace-ri.eu/IMG/pdf/d4.2_3ip.pdf
- [3] D4.1: “*Continuity Plan for the PRACE Advanced Training Centres*”; PRACE-3IP public deliverable,
<http://www.prace-ri.eu/IMG/pdf/d4.1.pdf>
- [4] D4.4: “*Training Collaboration*”; PRACE-3IP public deliverable,
http://www.prace-ri.eu/IMG/pdf/d4.4_3ip.pdf
- [5] PRACE Events website:
<http://events.prace-ri.eu>
- [6] D4.3: “*Report on Training Portal Development*”; PRACE-3IP public deliverable,
http://www.prace-ri.eu/IMG/pdf/d4.3_3ip.pdf
- [7] PRACE video tutorials:
<http://www.training.prace-ri.eu/tutorials/index.html>

List of Acronyms and Abbreviations

AISBL	Association sans but lucrative (legal form of the PRACE RI)
BSC	Barcelona Supercomputing Center (Spain)
CaSToRC	Computation-based Science and Technology Research Center, The Cyprus Institute (Cyprus)
CERN	European Organization for Nuclear Research
CSC	Finnish IT Centre for Science (Finland)
CUDA	Compute Unified Device Architecture (NVIDIA)
EPCC	Edinburg Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
GPGPU	General Purpose GPU
GPU	Graphic Processing Unit
GRNET	Greek Research and Technology Network (Greece)
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
InDiCo	Open source Event Management software originating from CERN
I/O	Input/Output
IPB	Institute of Physics, Belgrade (Serbia)
LRZ	Leibniz Supercomputing Centre (Garching, Germany)
MOOC	Massively Open On-line Course
MPI	Message Passing Interface
NIIF	National Information Infrastructure Development Program (Hungary)
OpenMP	Open Multi-Processing
PATC	PRACE Advanced Training Centre
PDF	Portable Document Format
PGAS	Partitioned Global Address Space
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE-PP	PRACE Preparatory Phase project
PRACE-1IP	PRACE First Implementation Phase project
PRACE-2IP	PRACE Second Implementation Phase project
PRACE-3IP	PRACE Third Implementation Phase project
PRACE-4IP	PRACE Fourth Implementation Phase project (submitted)
RIKEN	RIKEN Advanced Institute for Computational Sciences (Japan)
VŠB-TUO	VŠB-Technical University of Ostrava (Czech Republic)
WP4	Work Package 4 (Training)
XSEDE	U.S. National Science Foundation's eXtreme Science and Engineering Discovery Environment project

Executive Summary

The Training Work Package (WP4) of the PRACE Third Implementation Phase (PRACE-3IP) project has continued the highly appreciated PRACE training programme by offering additional training events during the extension period of the PRACE-3IP project¹ and by maintaining the existing on-line training services. Including the extension period, PRACE-3IP has delivered 307 training days as part of 107 training events that have reached an audience of 2,286 participants. In the academic year of 2014-2015 (until January 2015), PRACE has offered 26 courses organised by the PRACE Advanced Training Centres and two seasonal schools in the form of PRACE Autumn School 2014 and PRACE Winter School 2015. The PRACE Training Portal has continued to attract visitors from around the world and in collaboration with international partners the series of International HPC Summer Schools has continued to prosper. Overall, PRACE-3IP 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 key components of PRACE, throughout its existence, has been to offer user training in the form of courses, workshops, and seasonal schools as well as on-line resources for independent study. PRACE-3IP WP4 (Training) has continued this work, which began in 2007 by PRACE-PP and was carried over by PRACE-1IP and PRACE-2IP. In this document we describe the work done by the Training Work Package (WP4) since previous reports and especially during the PRACE-3IP extension period from July 2014 to January 2015 presenting up-to-date statistics of PRACE training events.

Due to the rapid pace of technological development connected to supercomputing systems and the ever-increasing importance of computational science, there exists a significant need for skills relevant to High-Performance Computing (HPC) in the scientific community. Training offered by PRACE aims to build up the HPC skill set of European scientists not only to ensure that PRACE computational resources are used effectively, but also to support the competitiveness of European science and industry in the field.

In 2012 PRACE-2IP established six PRACE Advanced Training Centres (PATCs) to serve as European hubs of world-class training in HPC and the courses offered by PATCs have ever since been the mainstay of the PRACE training programme. PRACE runs also a very successful series of PRACE Seasonal Schools that offer intensive training on a specific topic, akin to a scientific summer school. PRACE has also established an on-line PRACE Training Portal [1] to serve as a single source of training material, tutorials, and information of upcoming training opportunities. Additionally, in collaboration with XSEDE, Compute Canada, and recently RIKEN, PRACE has contributed to the organisation of a series of high-profile International HPC Summer Schools.

In this document, the activities of WP4 are presented in detail. Section 2 describes the activities of the PRACE Advanced Training Centres in the academic year 2014-2015 until the end of January 2015 (Task 4.1). Section 3 outlines two PRACE seasonal schools that were organised during the project extension with more detailed reports of each school included in

¹ PRACE-3IP Extension denotes the period of M25-M31 extending the work of WP2 – WP7 by seven months in order to ensure a seamless and continuous support of the project for the PRACE RI prior to the planned start of the PRACE-4IP project in H2020.

Section 7 (Task 4.2). Section 0 reports on the International HPC Summer School 2014 organised in collaboration with international partners (Task 4.4). Section 5 gives an overview of on-line training services of PRACE and their latest developments (Task 4.3).

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 curriculum, designed and coordinated by PRACE with input from user communities on an annual basis that corresponds to each academic year.

The deliverable D4.2: “*Report on Training Events*” [2] contains a comprehensive report on PATC activities during the PRACE-3IP project, including the publication of PRACE-3IP D4.1: “*Continuity Plan for the PATCs*” [3], which highlighted the operational and financial models of sustaining the PATCs taking into account lessons learnt.

2.1 PATC Courses from July 2014 to January 2015

This section presents the PATC activities in the PRACE-3IP extension period from July 2014 to January 2015, i.e. implementation of the 2014-15 joint PATC curriculum. An update to the statistics of the 2013-14 curriculum is also included, as three courses that took place in July 2014 were not included in D4.2 [2].

2.1.1 Key statistics on numbers of participants and courses

In summary, Table 1 below shows the key output from the six PATCs in the implementation of the 2013-14 curriculum (with updated numbers including three courses in July 2014 that were not previously included), and partially the 2014-15 curriculum. The same statistics from the 2012-13 curriculum are also shown.

Reporting period	2012-08 to 2013-07 “2012-13 curriculum”	2013-08 to 2014-07 “2013-14 curriculum”	2014-08 to 2015-01 “2014-15 curriculum” (partial)
Number of courses	71	81	26
Total duration (days)	204	233	74
Number of participants	1547	1682	604
Female (%)	12.9%	14.4%	12.7%
Non-academic (%)	9.9%	12.3%	13.1%
Non-host country (%)	20.6%	25.4%	36.1%
Non-PATC country (%)	13.8%	17.7%	28.1%

Table 1. Key statistics from implementation of the 2012-13, 2013-14, and partial 2014-15 PATC joint curricula. “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.

As highlighted in D4.2 [2], while PATC course participants came from all of Europe during the 2013-14 curriculum, the proportion of European participants from non-PATC hosting countries remained relatively low. The situation appears to be improving for the 2014-15 curriculum with an apparent higher proportion of participants coming from non-PATC countries (although the 2014-2015 curriculum is only partially delivered at this stage). One reason for this is the delivery of a limited number of PATC courses in non-PATC hosting countries as previously recommended (see [2] and [3]); such an example is a course entitled “Message-Passing Programming with MPI” delivered by EPCC in the Netherlands in October 2014.

However, it is still premature to draw any conclusions at this stage as comparisons of numbers should be done between annual curricula with a similar duration and number of courses. In particular, the 2014-15 curriculum is heavily “back-loaded” with a large number of courses occurring in the second half of the academic year. Hence a full analysis of the outcome of the 2014-15 curriculum will be carried out in the upcoming PRACE-4IP project. Nevertheless, the on-going activities of the PATCs continue to have a major impact in providing training for a large number of researchers. Work has already begun to devise the 2015-16 curriculum that will be supported by the PRACE-4IP project.

3 Seasonal Schools

In the frame of the PRACE-3IP project five PRACE Seasonal Schools were organised, the first three in the regular period, which were thoroughly described in the PRACE-3IP public deliverable D4.2: “*Report on Training Events*” [2], the last two in the extended period. These two, the PRACE Autumn School 2014 in Greece and the PRACE Winter School 2015 in the Czech Republic, are reported on in this section.

3.1 PRACE-3IP Extension Period Seasonal Schools’ Characteristics

In the PRACE-3IP extension period the two seasonal schools continued the trend of the previous training events in PRACE partners’ (non-PATC) countries, in fact the success stories of the previous three PRACE-3IP seasonal schools and also the preceding series of twelve seasonal schools during the PRACE-1IP and PRACE-2IP projects. They were successful despite the shorter preparation period—as the organisers were only selected in Spring 2014.

The programme of the last two PRACE-3IP schools was naturally designed in correspondence with local needs and competences and can be summarised as follows:

- The **PRACE Autumn School 2014** in HPC Programming Techniques focused on advanced training in programming models and optimization techniques, MPI/OpenMP and hybrid programming, profiling and benchmarking.
- In contrast to this traditional orientation, the **PRACE Winter School 2015** in HPC Tools for Data Intensive Processing chose a novel topic in the repertoire of PRACE seasonal schools, namely big data processing in HPC environment and corresponding tools, in accordance with current global developments.

The **attendance** of 54 and 45 participants in the schools, respectively, is higher than the average of 40 of previous seasonal schools. Moreover both schools had to reject several registrants to keep the number of participants adequate for efficient training. Overall, there were 99 participants in total, coming from 12 countries. The geographic distribution of these participants is shown in Table 2. As usual, we can observe that participants from the organizing country (and neighbours) clearly prevail at the events—a trend which generally holds true for all seasonal schools and face-to-face training events. Seasonal schools thus have

always a “mobilizing” effect on the hosting partner/country—this is one of their most positive consequences.

Country	Autumn School 2014	Winter School 2015	Total
Austria		1	1
Belgium		1	1
Bulgaria			
Cyprus	4		4
Czech Republic	1	37	38
Denmark		1	1
Finland			
France		1	1
Germany			
Greece	47		47
Hungary			
Ireland			
Israel			
Italy		2	2
Netherlands			
Norway			
Poland	1		1
Portugal			
Serbia			
Slovakia		1	1
Slovenia			
Spain			
Sweden		1	1
Switzerland	1		1
Turkey			
UK			
rest of Europe			
rest of the world			
Total	54	45	99

Table 2. All seasonal school participants and their geographic distribution

The PRACE Autumn School 2014 also offered a **live broadcasting** of the event. Despite its trial character and limited announcement, tens of visitors (maximum 115 during the first day) from about 20 countries worldwide “visited” the school in this way. The technology even allowed remote visitors to pose questions via text messages to be answered by the trainers. The organisers consider the live broadcasting of seasonal schools very promising, and when designed carefully, it can help to achieve a better and more geographically wide impact for the schools.

All seasonal schools ask the participants to provide **feedback** using a standard PRACE training feedback online form. The average mark on the generic question: “Overall, how would you rate this school” was high for PRACE Autumn School and PRACE Winter School, 8.0 and 7.8 respectively, which is in both cases better than the average of previous schools (7.5), in the range [0 = waste of time, 10 = excellent].

4 International HPC Summer School 2014

The International HPC Summer School 2014 was held on 1-6 June, 2014. It was a collaboration between PRACE, the U.S. National Science Foundation's eXtreme Science and Engineering Discovery Environment (XSEDE) project, Canada's Compute Canada and RIKEN Advanced Institute for Computational Sciences (RIKEN). A full report with details of the event and initial feedback statistics has been described in PRACE-3IP public deliverable D4.4: "*Training Collaboration*" [4]. Subsequent follow-up analysis of the feedback resulted in the following key findings:

- Mentors and participants found the summer school to be well organised, engaging, and of high quality.
- Participant and mentor thoughts regarding mentor matching, selection of mentors, and the role of mentoring varied. There were generally positive experiences, i.e. that participants had meaningful engagements with their assigned mentors, but the mentor matching process could be improved.
- Participants generally found technical session content combined with hands-on practice to be highly valuable and are confident they can incorporate these topics into their work.
- Social activities including outings and the poster sessions were amongst the most valued activities by participants.
- Additional suggestions for improvement made by attendees included sharing attendee contact information, offering small research stipends or compute time, further exploring social media, and conducting hands-on sessions in computer labs to reduce technical issues.

Organisation effort for the next iteration of the school in 2015—to continue an on-going collaboration since 2010—has begun during the PRACE-3IP project extension phase. The organising committee, with representatives from the same four partner organisations as the 2014 school, has designated the location of the next school to be held in Toronto, Canada. It is envisaged that announcement of the school and the acceptance of applications will begin in early 2015.

5 On-line services

In this section we summarise the activities of the Training Work Package (WP4) on-line services from January 2014 till the end of PRACE-3IP, and throughout the PRACE-3IP extension period until January 2015.

The WP4 on-line services include the following two services:

- PRACE Training Portal
- PRACE Events website

During this time period, WP4 on-line activities have included the completion of PRACE-3IP video tutorials, maintenance of the PRACE Training Portal [1] and PRACE Events website [5] and migration of the PRACE Event website hosting location.

At the end of this section we also provide Google Analytics statistics concerning the PRACE Training Portal during the reported time period.

For a report on PRACE-3IP WP4 on-line service activities before this time period, the reader is referred to PRACE-3IP public deliverable D4.3: “*Report on Training Portal Development*” [6].

5.1 Completion of PRACE 3IP Video Tutorials

To provide an alternative channel for sharing training material, in PRACE-3IP it was decided to create a set of short video tutorials focusing on topics relevant to PRACE users drawing from the expertise of the partners involved in Task 4.3. The full set of video tutorials were completed and made available [7] on the PRACE Training Portal by April 2014. These videos support independent study with the aim that the users get a short informative introduction to a topic to get them quickly started on it. The full set of completed video tutorials is identified in Table 3.

Partner	Video Tutorial Title
NCSA (Bulgaria)	<ul style="list-style-type: none"> • Introduction to Molecular Dynamics • DL_POLY Introduction
CSC (Finland)	<ul style="list-style-type: none"> • Electronic structure calculations with GPAW • Overview of GPAW • GPAW example: Geometry optimization of CO molecule
NIIFI (Hungary)	<ul style="list-style-type: none"> • Accessing the visualization infrastructure with TurboVNC and VirtualGL • CUDA 5 in user perspective
PSNC (Poland)	<ul style="list-style-type: none"> • Classification of Sparse Matrices for most efficient SpMV computation
ULFME (Slovenia)	<ul style="list-style-type: none"> • Block Meshing with ICEM-CFD • External flow bluff body CFD simulation with ANSYS Fluent

Table 3. PRACE-3IP Video Tutorials created by partners involved in the Training Portal task.

5.2 Maintenance of WP4 on-line services

The WP4 on-line services have continued to be used for their intended purposes throughout the reporting period.

The PRACE Events website [5] has been used to host the details of new, forthcoming and past PRACE Training events including both PRACE Seasonal Schools and PATC courses. For forthcoming events it has also been used for training event registration purposes.

The PRACE Training Portal [1] has been used to advertise forthcoming PRACE training events in its homepage, to host PRACE video tutorials and to enhance the repository of PRACE training material with lectures delivered in various PRACE training events.

The PRACE Training Portal now hosts 27 PRACE video tutorials. The training material from 89 PRACE training events can be found under the “Material” tab (<http://www.training.prace-ri.eu/material/index.html>) of the portal—30 more events than previously reported in D4.3 [6]. In total, over 380 different items of training material are hosted, with each item grouped together with other material related to similar topics. Close to 120 of these include a video of the lecture where both the presenter and the slides can be viewed simultaneously.

Training material includes the slides of the presented lectures in Portable Document Format (PDF) format and, whenever a hands-on session was held, the exercises and solutions (software code etc.) are also included where possible. In total, the number of documents

hosted by the training portal exceeds nine hundred documents—three hundred more than previously reported in D4.3 [6].

In addition to the above, a link exchange between the PRACE Training Portal and the XSEDE Training Portal has been established, with a link from PRACE to XSEDE Training Portal and a reciprocal link back from XSEDE to PRACE Training Portal. This is the first step towards a possible future collaboration between PRACE and XSEDE with regards to the distribution of training material. Similarly, links to some PRACE partner local training portal links have been added to the homepage of the training portal.

5.3 Migration of PRACE Events website

Migration of the PRACE Events website (running InDiCo software) was necessary at the end of 2014. The website was hosted by IPB (Serbia), but due to IPB leaving PRACE, it became acute to find a new hosting solution for the service. NIIF (Hungary) volunteered, and was chosen, to take over the hosting from IPB forthwith. This was an ideal solution with NIIF having previous experience in hosting and developing an InDiCo server ensuring a smooth transition of responsibilities.

Since the PRACE Events website serves as the single point of entry to PRACE training events containing both detailed information of each event as well as serving as the registration tool for them, it was crucial that the migration would proceed with minimal downtime and no data loss. This was achieved with minimal disturbance to the service. The migration took place in close coordination between members of staff from NIIF and IPB in December 2014. It was successfully completed in a few days with minimal downtime (only a few hours) to the PRACE Events website.

5.4 Google Analytics PRACE Training Portal Statistics

As reported in the PRACE-3IP public deliverable D4.3: “*Report on Training Portal Development*” [6], Google analytics was incorporated to the training portal in May 2013. The visitor statistics for the Training Portal covering the full year of 2014 are summarised in Figure 1.

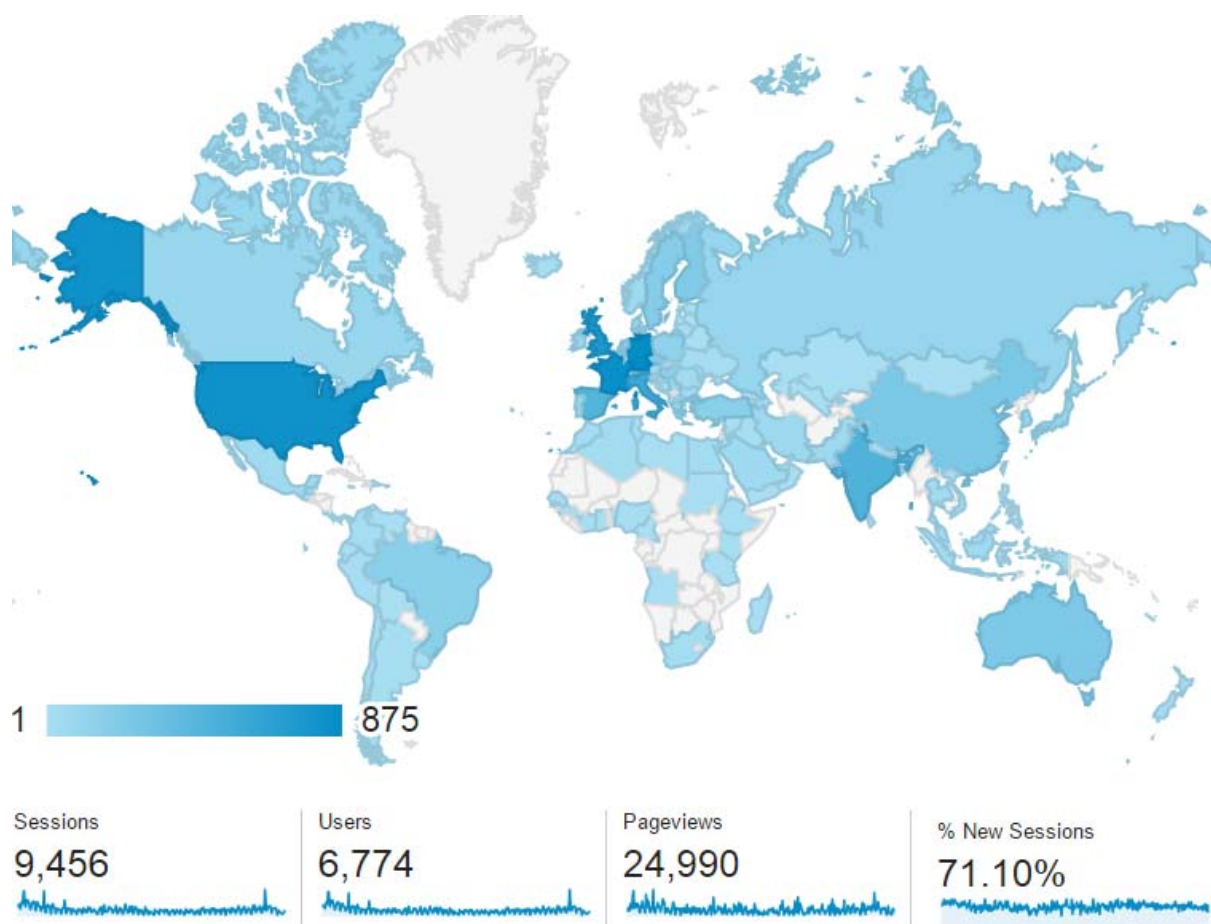


Figure 1. Visitor statistics for the Training Portal during the full year of 2014.

From Figure 1 it is apparent that the PRACE Training Portal continues to reach a global audience. It is interesting to note the increased amount of traffic emanating from the USA, which is of similar (and sometimes higher) levels compared to PRACE member countries. One should also take note of the relative high amounts of traffic from India. Overall, it is clear that the PRACE Training Portal attracts visitors from around the world and has built up an extensive user base making it a significant training resource for the global HPC community.

6 Conclusions

PRACE-3IP WP4 (Training) has continued to offer user training in the form of courses, workshops, and seasonal schools as well as on-line resources for independent study during the extension period of PRACE-3IP. Including the extension period, PRACE-3IP has delivered 307 training days as part of 107 training events that have reached an audience of 2286 participants. In the academic year of 2014-2015 (until January 2015), PRACE has offered 26 courses organised by the PRACE Advanced Training Centres and two seasonal schools in the form of PRACE Autumn School 2014 and PRACE Winter School 2015.

The PRACE Training Portal has also continued to attract visitors from around the world and in collaboration with international partners the series of International HPC Summer Schools has continued to prosper. Overall, PRACE-3IP 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.

7 Annex

Detailed reports of the two PRACE Seasonal Schools organised in Autumn 2014 and Winter 2015 are included below with full programme, statistics and analysis of the overall impact of each school.

7.1 Report on PRACE Autumn School 2014

7.1.1 Basic information about the event

Name: PRACE Autumn School in HPC Programming Techniques - Athens, Greece

Dates: 25-28 November 2014

Location: Athens, Greece

Organizing sites: GRNET

Mission: The training event addresses existing and potential users of High Performance Computing systems in Europe. Researchers and students from Europe will receive advanced training in programming models and optimization techniques, MPI/OpenMP and hybrid programming, profiling and benchmarking. A description of PRACE, its infrastructure and best practices for accessing the top-level European systems will be provided. The PRACE Autumn School 2014 will focus on in-depth presentations for thorough understanding of these HPC topics, followed by hands-on training.

Event URL: <https://events.prace-ri.eu/conferenceDisplay.py?confId=328>

7.1.2 Organizational details

Local organizing committee

Name	Role
Ioannis Liabotis	Chair – Student Selection
Dimitra Kotsokali	Administration
Paschalis Korosoglou	Generic Assistance – Student Selection
Alexandra Charalabidou	Generic Assistance – Student Selection
Dimitris Dellis	Generic Assistance – Student Selection
Vangelis Floros	Generic Assistance – Student Selection

Venue

OTE Academy (<http://oteacademy.com>). OTEAcademy, a member of the OTE Group of companies, focuses on the development of human resources by offering innovative and exclusive training services. Its goal is that the professional, through continuing and systematic learning, is becoming better and living better. The same venue has been used for the successful implementation of the two previous PRACE related training events that GRNET has organized in Greece: PRACE Winter School 2009 and HP-SEE, LinkSceem, PRACE Training Summer 2011.

The premises of the autumn school are being selected for the following reasons:

- Professional training environment
- Extended training facilities
- Good location in the Northern suburbs of Athens

- Professional technical and administrative support
- Adequate areas for lunch and coffee breaks next to the training room
- Organization and provision of catering service

7.1.3 Program & content

Program committee

Name	Role
Ioannis Liabotis	Chair
Paschalis Korosoglou	Member
Alexandra Charalabidou	Member
Dimitris Dellis	Member
Vangelis Floros	Member
Martti Louhivuori	Member

Final program

Program Summary	Time
Technical Presentations about PRACE, HPC Systems and Programming models ²	Day 1 Afternoon
Parallel Programming with MPI (including Hands on sessions) ²	Day 2 Morning
Parallel Programming with OpenMP (including Hands on sessions) ²	Day 2 Afternoon
Hybrid (MPI + Open MP) programming (including Hands on sessions) ²	Day 3 Morning
Hybrid (MPI + Open MP) programming (including Hands on sessions): Case Study: Application of Mixed Mode Programming (MPI+OpenMP) in a Real-World Scientific Code ²	Day 3 Morning
Scientific Libraries Usage (including Hands on sessions) ²	Day 3 Morning
Profiling (including Hands on sessions)	Day 3 Afternoon
Optimization and benchmarking (including hands on)	Day 4 Morning

List of trainers

Dr. **Costas Bekas**, is managing the Foundations of Cognitive Computing Group at IBM Research - Zurich. Costas received B. Eng., Msc and PhD diplomas, all from the Computer Engineering & Informatics Department, University of Patras, Greece, in 1998, 2001 and 2003 respectively. In 2003-2005, he worked as a postdoctoral associate with Professor Yousef Saad at the Computer Science & Engineering Department, University of Minnesota, USA. Costas's main focus is in HPC systems and their impact in every day life, science and business. Costas is a recipient of the 2013 ACM Gordon Bell Prize, and the 2012 PRACE Award.

Dr. **Alan O'Cais** has worked at the Juelich Supercomputing Centre since 2010 within the Application Support division. His work is primarily focussed within the LinkSCEEM-2 project helping to develop a Virtual Research Community in Computational Science in the Eastern Mediterranean region. He has been active in the field of computational research and

high performance computing since receiving his bachelors degree in Theoretical Physics from Trinity College Dublin in 2001. He received a Masters Degree in High Performance Computing in 2002 and a PhD in Lattice Quantum Chromodynamics in 2005. He has held research positions at Trinity College and the University of Adelaide. In 2008 he joined the Cyprus Institute as Scientific Coordinator of the Computation-based Science and Technology Research Centre (CaSToRC).

Mr. **Paschalis Korosoglou** is the technical coordinator of the Scientific Computing Office at the Aristotle University of Thessaloniki. He holds an MSc in Computational science and has participated in major national and European projects, mainly in user services provisioning and application support activities. In this respect he holds a solid background on scientific applications porting, profiling, optimisation and benchmarking activities and has developed under several occasions code optimisation improvements related to hybrid parallel implementations (MPI/OpenMP) and parallel I/O best practices. Currently, he is leading the activities of the Scientific Computing Office with regard to Grid, Cloud, Big Data and HPC activities. He has been involved in national and international Working Groups, such as CSIRT, EUGridPMA, GOC, SAM and has contributed within white papers published by the e-Infrastructures Policy Group (e-IRG).

Dr. **Dimitris Dellis** holds a B.Sc. in Chemistry (1990) and PhD in Computational Chemistry (1995) from the National and Kapodistrian University of Athens, Greece. He has extensive HPC and grid computing experience. He was using HPC systems in computational chemistry research projects on fz-juelich machines (J90, T3E, Regata, 2003-2005). He received an HPC-Europa grant on BSC (2009). In EGEE/EGI projects he acted as application support and VO software manager for SEE VO, grid sites administrator (HG-02, GR-06), NGI_GRNET support staff (2008-2014). In PRACE 1IP/2IP/3IP he was involved in benchmarking tasks either as group member or as BCO (2010-2014). Currently he holds the position of “Senior HPC Applications Support Engineer” at GRNET S.A. where he is responsible for activities related to user consultations, porting, optimization and running HPC applications at national and international resources.

Mr **Ioannis Liabotis** is a project manager at GRNET S.A. During his 8 years at GRNET he has been involved in a series of FP6 and FP7 eInfrastructures projects related to HPC, Grids and Cloud. He has been the technical coordinator of the SEE-GRID-2, SEE-GRID-SCI and HP-SEE projects. He is currently a member of the PRACE Council and a member of the PRACE projects management boards representing Greece. Ioannis obtained his Diploma in Electrical and Computer Engineering from NTUA, Greece, in 1998; and his Mphil in Computation from UMIST, UK, in 2000.

Mr **Nikolaos Tryfonidis** is a PhD Candidate in the Aristotle University of Thessaloniki, Greece. After graduating with a Physics Degree, he completed a Master's Degree in High Performance Computing. His PhD project involves the computational study of plasma physics, using a variety of HPC methods, from mixed-mode programming to GPGPUs.

Designing the program

The training addressed existing and potential users of High Performance Computing systems in Europe. Researchers and students from Europe received advanced training in programming models and optimization techniques, MPI/OpenMP and hybrid programming, profiling and benchmarking. A description of PRACE, its infrastructure and best practices for accessing the top level European systems has been provided. The PRACE Autumn School 2014 focused on in-depth presentations for thorough understanding of these HPC topics, followed by hands-on training.

Computer resources

The hands-on sessions of the training were run on the CaSToRC (Cyprus) Training HPC System. The provided resources were in general sufficient for the event. Slight delays in queues were observed due to some trainees' errors. This was actually a good practical example that assisted in the success of the hands-on session.

7.1.4 Participants & feedback

Number of participants by country

Country	Number of Participants
Greece	47
Cyprus	4
Switzerland	1
Poland	1
Czech Republic	1
Total	54

Process for selecting the participants

The registration process contained an extensive questionnaire requiring applicants to express their HPC qualifications both using multiple choice queries or verbal statements. The student selection committee evaluated the students' responses and selected 60 of the 83 applicants. Finally 54 students attended the event.

Statistics and analysis of the feedback survey

The average mark on the generic question: "Overall, how would you rate this school" was 8.0, indicating that the participants were overall very satisfied with the school. Some indicative graphs from the responses are shown in Figure 2.




Figure 2. Overview of feedback results from PRACE Autumn School 2014.

Analysing the results shown in Figure 2 we can conclude that the organization and the selected subjects were of high quality. The responses were mixed in the question related to the pace of the lectures. We believe that this was due to the varying levels of expertise among the participants in the event.

Selected verbal comments received as feedback and which are very helpful for future events follow below.

GENERAL COMMENTS

- Keep your good job!! (sic)
- It was such a great experience!
- In general very well organized, great experience and valuable knowledge. There is space for some improvement on the interactive part and the hands on exercises. Also it is very difficult to keep up for 3 hours, maybe a couple more coffee breaks (not 30 minute breaks though!)

SUGGESTIONS FOR THE FUTURE

- Details for preparation of code for proposal submission. Step by step from serial to scalable code.
- Maybe a more fundamental approach at parallel computing (MOOC?) to open new opportunities and make community grow.
- Advanced HPC Programming (CUDA & PGAS) is a topic of great importance in developing HPC applications. Numerical Algorithms and usage of Scientific Libraries
- CUDA language
- I would like to attend once a more introductory school about HPC programming.
- Establish of small training classes (on demand) within PRACE infrastructure regionally
- Heterogenous Cluster Programming. (sic)

7.1.5 *Conclusions & lessons learned*

In what follows are some statements from the organizing committee that reflects the lessons learned after the completion of the PRACE training.

- Demand for such trainings in countries like Greece is very high.
- The level of expertise of the students is important to match this of the lectures and the pace of lecturers.
- Practical examples during hands on sessions should be provided both in C and FORTRAN where relevant.

7.2 Report on PRACE Winter School 2015

7.2.1 *Basic information about the event*

Name: PRACE Winter School 2015 – HPC Tools for Data Intensive Processing

Dates: 12-15 January 2015

Location: Ostrava, Czech Republic

Organizing sites: VŠB-Technical University of Ostrava and IT4Innovations Centre of Excellence, Czech Republic

Mission: Historically, High Performance Computing (HPC) was primarily about computing-intensive processing, opposed to data-intensive jobs. Nowadays, complex HPC simulations combine these two attributes: Large-scale modelling as a rule performs demanding computations on very large data sets. Moreover, a new subset of HPC jobs occurs, searching for useful information and patterns in the data itself – high-performance data analysis. Data

becomes the focus of HPC. PRACE Winter School 2015, presenting HPC Tools for Data Intensive Processing, reflects this development and offers corresponding training to existing and potential users of HPC systems in Europe.

Event URL: <https://events.prace-ri.eu/conferenceDisplay.py?confId=330>

<https://www.prace-ri.eu/PRACE-Winter-School-2015/>

7.2.2 Organizational details

Local organizing committee

Name	Role
Ondřej Jakl	Chair – overall coordination
David Horák	Deputy – technical equipment
David Hrbáč	Collaborator – computing resources
Anna Němcová	Collaborator – accounting, agreements
Karina Pešatová	Collaborator – dissemination, event organization
Vít Vondrák	Head of the local PRACE team

Venue

VŠB - Technical University of Ostrava (VŠB-TUO, <http://www.vsb.cz>) was founded in 1849, and has since grown into a modern institution of higher learning. Nowadays it is the fourth largest university in Czech Republic with over 20,000 students studying in bachelor's, master's and doctoral degree programs in seven faculties. It is the home of the National Supercomputing Centre IT4Innovations, the factual representative of the Czech Republic in PRACE.

VŠB-TUO has in Ostrava-Poruba one of the best-equipped campuses in the Czech Republic, with many new or refurbished buildings, and is conveniently located on bus and tramlines. Its Nová aula (New Aula), the venue of the School, belongs to new architectural symbols of the university and as a building equipped with all the modern amenities and technical features provides ideal conditions for training.

7.2.3 Program & content

Program committee

Name	Role
Ondřej Jakl (IT4Innovations)	Chair
Branislav Jansík (IT4Innovations)	Member
Michel Kern (INRIA & Maison de la Simulation)	Member
Michael Krieger (RISC, Austria)	Member
Tomáš Kozubek (IT4Innovations)	Member
Martin Palkovič (IT4Innovations)	Member
Vít Vondrák (IT4Innovations)	Member
Wolfgang Schreiner (RISC, Austria)	Member

Final program

Day 1 – Mon, January 12, 2015			
<i>Start</i>	<i>Stop</i>	<i>Concept</i>	<i>Speaker</i>
9:00	10:00	High performance I/O discussion	Jeremy Logan
10:00	11:00	Using HDF5	Norbert Podhorszki
11:00	11:30	Break	
11:30	12:30	Using Parallel HDF5	Norbert Podhorszki
12:30	1:30	Lunch	
1:30	2:00	ADIOS overview	Jeremy Logan
2:00	3:00	ADIOS – write files	Jeremy Logan
3:00	3:30	Break	
3:30	4:00	ADIOS tools	Jeremy Logan
4:00	5:00	ADIOS no-xml	Norbert Podhorszki
5:00		HW assignments	
Day 2 – Tue, January 13, 2015			
<i>Start</i>	<i>Stop</i>	<i>Concept</i>	<i>Speaker</i>
9:00	9:30	Go over HW assignment	Norbert Podhorszki
9:30	10:30	ADIOS reading	Jeremy Logan
10:30	11:00	Break	
11:30	12:30	ADIOS Staging	Norbert Podhorszki
12:30	1:30	Lunch	
1:30	2:30	pdbR overview	George Ostrouchov
2:30	3:00	Break	
3:00	5:00	pdbR hands-on	George Ostrouchov
5:00		HW assignments	
Day 3 – Wed, January 14, 2015 - Visit & pdbR session			
<i>Start</i>	<i>Stop</i>	<i>Concept</i>	<i>Speaker</i>
9:00	10:30	Introduction to Visit	Dave Pugmire
10:30	11:00	Break	
11:00	12:15	Hands on with Visit	Dave Pugmire
12:15	1:30	Lunch	
1:30	3:00	pdbR advanced topics	George Ostrouchov
3:00	3:30	Break	
3:30	5:00	pdbR advanced topics	George Ostrouchov
Day 3 – Wed, January 14, 2015 - Hadoop session			
[This parallel session was cancelled, c.f. below]			
<i>Start</i>	<i>Stop</i>	<i>Concept</i>	<i>Speaker</i>
9:00	10:00	An introduction to Big Data	Shadi Ibrahim
		The MapReduce programming	
10:00	10:30	model	Shadi Ibrahim
10:30	11:00	Break	
11:00	12:15	Hadoop ecosystem: An overview	Shadi Ibrahim
12:15	1:30	Lunch	
1:30	3:00	Hands on deploying/using Hadoop	Shadi Ibrahim
3:00	3:30	Break	
3:30	5:00	Configuring Hadoop hands on	Shadi Ibrahim

Day 4 – Thu, January 15, 2015 - VisIt & pdbR session			
<i>Start</i>	<i>Stop</i>	<i>Concept</i>	<i>Speaker</i>
9:00	9:30	Go over HW assignment	George Ostrouchov
9:30	10:30	Staging Plugins	└
10:30	11:00	Break	
11:30	12:30	Hands on creating/using plugins	└
Day 4 – Thu, January 15, 2015 - Hadoop session			
[This parallel session was cancelled, c.f. below]			
<i>Start</i>	<i>Stop</i>	<i>Concept</i>	<i>Speaker</i>
9:00	9:30	Go over HW assignment	Shadi Ibrahim
9:30	10:30	Hadoop: Optimizations and open issues	Shadi Ibrahim
10:30	11:00	Break	
11:30	12:30	Hands on writing MapReduce applications	Shadi Ibrahim

List of trainers

Jeremy Logan is a Computational Scientist at the University of Tennessee and works closely with the Scientific Data Group at Oak Ridge National Laboratory. Jeremy's research interests include I/O performance, data and workflow management, and the application of domain specific, generative techniques to High Performance Computing.

Norbert Podhorszki is a Research Scientist in the Scientific Data Group at Oak Ridge National Laboratory. He is the lead developer of ADIOS. He works with application users of the Oak Ridge Leadership Facility to improve their I/O performance using ADIOS. His research interest is in how to enable data processing on-the-fly using memory-to-memory data movements, e.g. for speeding up I/O, coupling simulation codes, and building in-situ workflows.

George Ostrouchov is a Senior Research Scientist in the Scientific Data Group at the Oak Ridge National Laboratory and Joint Faculty Professor of Statistics at the University of Tennessee. His doctoral work was on large sparse least squares computations in data analysis. His research interests have been for many years at the intersection of high performance computing and statistics. He initiated and continues to lead the pbdR project. George is a Fellow of the American Statistical Association.

David Pugmire is a Research Scientist in the Scientific Data Group, in the Computer Science and Mathematics Division at ORNL. His research interests are in visualization of large scientific data.

Shadi Ibrahim is a permanent INRIA research scientist within the KerData research team. He obtained his Ph.D. in Computer Science from Huazhong University of Science and Technology in Wuhan of China in 2011. His research interests are in cloud computing, big data management, data-intensive computing, high performance computing, virtualization technology, and file and storage systems. He has published several research papers in recognized big data and cloud computing research journals and conferences, among which, several papers on optimizing and improving Hadoop MapReduce performance in the cloud and one book chapter on MapReduce framework.

Unfortunately, Shadi Ibrahim had to cancel his lectures only three days before the school started due to a *force majeure*. The organisers received notification of a possible risk of

cancellation on Friday January 9 2015, i.e. three days before the event, with the cancellation then confirmed later by him.

Designing the program

The Programme committee decided to focus the School on “HPC Tools for Data Intensive Processing”, i.e. on big data processing in HPC environment and corresponding tools, in accordance with the scientific interests of the hosting institution. Applications in bioinformatics, life sciences or earth sciences were expected. This topic was quite novel in the repertoire of PRACE Seasonal Schools, which was also one reason for this choice.

The aim, reflected in the programme design, was to acquaint the students (scientists and developers who wish to use high-level tools but at the same time require fine-grain control over data processing) with a suite of high-quality and open-source tools suitable for coping with large sets of scientific data. This suite was “complete” in the sense that it covers a natural sequence of tools from high performance I/O (HDF5 and Adios), via analytics (R, pbdR), to visualization (VisIt). A complementary, distributed approach is offered by the Hadoop platform.

The programme was also affected by the availability of trainers, since our effort was to invite outstanding specialists, preferably leaders or members of the tool/library developer teams.

Description of contents

The timeline (see School’s web presentation at <https://www.prace-ri.eu/PRACE-Winter-School-2015/> for details) was as follows. The School started on Monday with a short Opening covering its programme, organization, hosting institution and PRACE. The first training session was called High Performance I/O, teaching serial and parallel HDF5 and later ADIOS, which makes scaling I/O easy, portable, and efficient.

This topic continued also on Tuesday, however the main session of the second day was Analytics, providing the basics of the R language with a focus on its recent high performance extensions enabled by the Programming with Big Data in R (pbdR) project.

Since Wednesday until its end on Thursday after lunch the School was planned to split into two parallel sessions. The Visualization session introduced VisIt is an open source, turnkey application for large scale simulated and experimental data sets, in fact an infrastructure for parallelized, general post-processing of extremely massive data sets. On the other hand the Hadoop session should have served as a first step towards exploring the Hadoop platform and also to provide a short introduction into working with big data in Hadoop, including MapReduce as an important programming model for Big Data processing in the Cloud.

As explained above, the contracted tutor of the Hadoop session, Shadi Ibrahim, had to cancel his participation in the school at the very last moment, and as the Programme committee could not find any other free Hadoop tutor in such a short period of time, the Hadoop parallel session had to be cancelled. That is why all participants attended the Visualization session. However, the ca. 20 participants who chose Hadoop as their preferred parallel session during the registration, obtained Ibrahim’s presentation (slides), prepared for the school, as a small remedy.

Computer resources

The main computing resource for the School was the Anselm supercomputer at the National Supercomputing Centre IT4Innovations, which is a Bull cluster from 2013. After intensive preparatory work of the organizers, Anselm’s administrators and tutors this platform could meet very well the requirements of the school, because of dedicated allocation of nodes and

high throughput queues. For some hands-on sessions, e.g. during the VisIt tutorial, virtual machines prepared beforehand and installed on participants' laptops were more appropriate.

7.2.4 *Participants & feedback*

Number of participants by country

Country	Number of Participants
Czech Republic	37
France	1
Italy	2
Slovakia	1
Austria	1
Belgium	1
Denmark	1
Sweden	1
Total	45

Process for selecting the participants

The registration form contained an extensive questionnaire requiring applicants to express their short autobiography, motivation, HPC qualifications, interest to present poster, etc. Based on this data, the Admission committee, personally the same as the Programme committee, selected 45 participants of the School out of 55 registrants. When necessary, the registrations were assigned scores allowing comparisons. The selection procedure was performed on the fly, in four sets, which gave early registrants the advantage to learn the result promptly (within three weeks). Three applications from Asia were not found eligible.

There were three late cancellations and two no show-ups, but those absences have been replaced by local, conditionally accepted students interested in the school.

Statistics and analysis of the feedback survey

In order to include results in this deliverable, the feedback from the participants was collected earlier than usual, i.e. after seven days of opening the PRACE training survey. Nevertheless, a good sampling was achieved with 31 responses collected in total. The participants were in general satisfied with the school as indicated by an average mark of 7.8 to the generic question "Overall, how would you rate this school". The objections mentioned explicitly were quite diverse, from "it would be better to define prerequisites more precisely before the school start for the next time" to "many tables in auditorium were broken". On the other hand many participants agreed on the fact that especially the Visualization session was great. Surprisingly, there were no complaints about the cancelled Hadoop tutorial.

7.2.5 *Conclusions & lessons learned*

PRACE Seasonal Schools are precious events, which have tremendous knowledge dissemination effect. Another benefit is in shaping the local/national HPC training infrastructure, which is for many countries, including the Czech Republic, a task of the day.