

H2020-FETHPC-2014

Coordination of the HPC strategy



EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D5.4 HPC Training Roadmap

Final

Version: 0.2
Author(s): Catherine Inglis, EPCC
Blaž Rodič, University of Ljubljana
Malcolm Muggeridge - Seagate
Date: 21.02.2017

Project and Deliverable Information Sheet

EXDCI Project	Project Ref. №: FETHPC-671558	
	Project Title: European eXtreme Data and Computing Initiative	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: < D5.4>	
	Deliverable Nature: <DOC_TYPE: Report>	
	Dissemination Level: PU	Contractual Date of Delivery: 28 / February / 2017
		Actual Date of Delivery: 27 / February / 2017
EC Project Officer: Béatrice Marquez-Garrido		

* - 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: HPC Training Roadmap	
	ID: D5.4	
	Version: <0.2>	Status: Final
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): EXDCI-D5.4-v0.2.docx	
Authorship	Written by:	Catherine Inglis, EPCC
	Contributors:	Blaž Rodič, University of Ljubljana; Malcolm Muggeridge - Seagate
	Reviewed by:	François Bodin, IRISA Katia Soto, PRACE
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	09/February/2017	Draft	For internal review
0.2	21/February/2017	Draft	Updated after internal review
0.2	27/February/2017	Final version	Approved by EXDCI Management Board and Technical Board

Document Keywords

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

Copyright notices

© 2017 EXDCI Consortium Partners. All rights reserved. This document is a project document of the EXDCI project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the EXDCI partners, except as mandated by the European Commission contract FETHPC-671558 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

1	Introduction	1
2	Overview of the current status of HPC training provision	2
2.1	Who needs HPC training?	2
2.2	How is HPC training currently delivered?	3
2.3	Overview of short-term, intensive training opportunities.....	3
2.3.1	PRACE	3
2.3.2	Other short training courses and internship-style training	5
2.4	HPC training in higher education.....	7
2.4.1	Postgraduate qualifications in HPC.....	7
2.4.2	HPC training for undergraduates.....	8
2.5	Marie-Curie Innovative Training Networks (ITNs).....	8
2.6	On-line training opportunities.....	9
2.7	The Software Sustainability Institute (SSI).....	10
3	The status quo: a survey of training providers	11
4	A look to the future: perceived training needs of PATC course attendees	16
5	Gaps in the current provision of training	22
5.1	Training undergraduates	22
5.2	Industrial engagement.....	23
6	EuroLab-4-HPC and the HPC Training Curriculum	25
7	Summary of challenges and Recommended Actions	26
8	Conclusion.....	29

List of Figures

Figure 1	Survey response: Institution type	11
Figure 2	Survey response: Courses for externals	12
Figure 3	Survey response: Course duration.....	12
Figure 4	Survey response: Delivery method	13
Figure 5	Survey response: Target group	13
Figure 6	Survey response: Course accreditation	14
Figure 7	Survey response: Course fees	14
Figure 8	Survey response: Course focus	15
Figure 9	Survey response: Emerging topics.....	15
Figure 10	PATC course attendees' future training needs (2014-16)	17
Figure 11	Comparison of perceived future training needs in basic HPC: 2014 vs. 2016.....	18
Figure 12	Comparison of perceived future training needs in advanced HPC: 2014 vs. 2016.....	18
Figure 13	Comparison of perceived future training needs in optimisation: 2014 vs. 2016.....	19
Figure 14	Comparison of perceived future training needs in porting code: 2014 vs. 2016	19
Figure 15	Comparison of perceived future training needs in HPC applications: 2014 vs. 2016.....	20
Figure 16	Comparison of perceived future training needs in community-specific areas: 2014 vs. 2016	20
Figure 17	Comparison of perceived future training needs in visualisation: 2014 vs. 2016	21

List of Tables

Table 1 Challenges and Recommended Actions	28
--	----

References and Applicable Documents

List all external documents referenced in this document

- [1] <http://www.exdci.eu>
- [2] <http://www.prace-project.eu>
- [3] <http://www.etp4hpc.eu>
- [4] <http://www.eesi-project.eu/>
- [5] <http://www.prace-ri.eu/prace-advanced-training-centres/>
- [6] <http://www.prace-ri.eu/prace-seasonal-schools/>
- [7] <https://www.computecanada.ca/>
- [8] <http://www.riken.jp/en/>
- [9] <https://www.xsede.org/>
- [10] <http://www.ihpcss.org/>
- [11] <https://summerofhpc.prace-ri.eu/>
- [12] <http://www.prace-ri.eu/trainings/>
- [13] <https://www.futurelearn.com/courses/supercomputing>
- [14] <https://www.futurelearn.com/courses/big-data-r-hadoop>
- [15] <http://www.prace-ri.eu/prace-codevault/>
- [16] <http://bioexcel.eu/>
- [17] <http://www.hpc-escape.eu/>
- [18] <https://pop-coe.eu/>
- [19] <https://www.hipeac.net/>
- [20] <https://www.hipeac.net/acaces/>
- [21] <https://www.hipeac.net/mobility/internships/>
- [22] <http://www.supercomp.org/>
- [23] <http://isc-hpc.com/>
- [24] <https://www.epcc.ed.ac.uk/msc/programmes/msc-programmes>
- [25] <http://www.ed.ac.uk/studying/postgraduate/degrees/index.php?r=site/view&id=902>
- [26] <https://www.maths.tcd.ie/hpcmsc/>
- [27] <http://www.swansea.ac.uk/postgraduate/taught/science/msc-high-performance-and-scientific-computing/>
- [28] <https://www.liverpool.ac.uk/study/postgraduate/taught/big-data-msc/overview/>
- [29] <http://www.mhpc.it/>
- [30] <http://www.fib.upc.edu/en/masters/miri.html>
- [31] <https://developer.nvidia.com/educators>
- [32] <https://developer.nvidia.com/teaching-kits>
- [33] http://ec.europa.eu/research/mariecurieactions/about/innovative-training-networks_en

- [34] www.anl.gov
- [35] <http://extremecomputingtraining.anl.gov/2016-videos>
- [36] <https://exdci.eu/jobs-and-training/training-portal>
- [37] <http://hpcuniversity.org/>
- [38] <https://www.cyi.ac.cy/>
- [39] <http://supercomputing.cyi.ac.cy/>
- [40] <https://www.xsede.org/training>
- [41] <https://sites.google.com/a/lbl.gov/hpc-training-best-practices/home>
- [42] <https://tess.elixir-europe.org>
- [43] <https://www.software.ac.uk/>
- [44] <https://arxiv.org/abs/1610.02608>
- [45] <https://issuu.com/epccedinburgh/docs/roadmap>
- [46] <http://www.sciencefestival.co.uk/>
- [47] <http://www.bgtb.org.uk/>
- [48] <https://www.thebigbangfair.co.uk/>
- [49] <https://www.epcc.ed.ac.uk/discover-and-learn/resources-and-activities/what-is-a-supercomputer/wee-archie>
- [50] https://www.mathe-im-advent.de/fotos_cray/
- [51] <http://bjc.berkeley.edu/>
- [52] https://www.youtube.com/watch?v=eqbgqSe_hvU
- [53] <http://cordis.europa.eu/fp7/ict/computing/factsheets/planethpc.pdf>
- [54] <http://www.interreg-danube.eu/approved-projects/innohpc>
- [55] <http://www.prace-ri.eu/hpc-access/shape-programme/>
- [56] <https://www.eurolab4hpc.eu/>
- [57] <https://www.eurolab4hpc.eu/static/deliverables/D3-1--final-HPC-curriculum.95d306191a15.pdf>
- [58] https://www.youtube.com/channel/UCfwgjtIQB3puojz_N9ly_Ag/feed

List of Acronyms and Abbreviations

Below is an extensive the List of Acronyms used in previous deliverables. Please add additional ones specific to this deliverable and delete unrelated ones.

ACACES	International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems
AP	Advanced Placement
ATPESC	Argonne Training Program on Extreme-Scale Computing
BoF	Birds-of-a-Feather session
CERN	European Organization for Nuclear Research
CFD	Computational Fluid Dynamics
CoE	Centres of Excellence for Computing Applications
CSE	Computational Science and Engineering
EC	European Commission
EESI	European Exascale Software Initiative
EPCC	Edinburgh Parallel Computing Centre
ERDF	European Regional Development Fund
ESRF	European Synchrotron Radiation Facility
ETP4HPC	European Technology Platform for High Performance Computing
EU	European Union
EXDCI	European Extreme Data and Computing Initiative
FET	Future and Emerging Technologies
FIB	Facultat d'Informàtica de Barcelona – Barcelona School of Informatics
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HiPEAC	European Network on High Performance and Embedded Architecture and Compilation
HLRS	Höchstleistungsrechenzentrum Stuttgart (High Performance Computing Centre in Stuttgart)
HPC	High Performance Computing
ICTP	The Abdus Salam International Centre for Theoretical Physics
ISC	ISC High Performance – the Event for High Performance Computing, Networking and Storage
IT	Information Technology
ITER	International Thermonuclear Experimental Reactor
ITN	Innovative Training Network
MOOC	Massive Open Online Course
MPI	Message Passing Interface
PATC	PRACE Advanced Training Centre
PRACE	Partnership for Advanced Computing in Europe
SC	The International Conference for High Performance Computing, Networking, Storage and Analysis
SHAPE	SME HPC Adoption Programme in Europe

SIGHPC	ACM Special Interest Group on HPC
SISSA	International School for Advanced Studies
SME	Small and Medium Enterprise
SSI	Software Sustainability Institute
STEM	Science, Technology, Engineering and Mathematics
TeSS	Elixir's Training e-Support System
US	United States
xMOOC	extended Massively Open Online Course
XSEDE	Extreme Science and Engineering Discovery Environment

Executive Summary

This deliverable – the HPC Training Roadmap – builds on prior activities, surveys and reports of the EXDCI¹ members – PRACE² and ETP4HPC³ – and their partners, as well as EESI2⁴, in their own specific spheres of expertise. It gives an overview of the current status of HPC training provision, and examines the future training needs of the community and how the gaps can be bridged. It considers how HPC training can be made more accessible, and how it can be scaled up to meet rising demand, including looking at the university sector and provision for undergraduate and graduate students. It recommends short-term and medium- to long-term actions to address the identified challenges.

1 Introduction

In ever more areas of academic research, industry, commerce and government, large-scale computational approaches are used for simulation and increased productivity in research and innovation. However, the shortage of personnel with expertise in HPC skills is still a barrier to increasing the uptake of HPC in academia and industry, and this has been identified as a major risk to European competitiveness. In part, this is due to the insufficient provision of HPC training courses at all levels and across all sectors. EXDCI aims first to generate interest in HPC among young researchers, and also to improve the provision of and access to suitable training for all.

Task 5.3, *Identifying and Meeting Future Training Needs*, aims to review the existing provision of HPC training, identify the future training needs of the HPC community, and produce recommendations based on the resulting gap analysis.

As HPC becomes more pervasive, the potential number of users can only continue to grow – and thus the demand for training will continue to increase. Training must be made easily accessible for all categories of users at all levels, to increase the number of people in the workforce who have suitable skills.

Questions that are considered in this document are:

- What is the current status of HPC training provision in Europe?
- Who currently benefits from HPC training, and who is missing out?
- What is lacking from current provision, and what are the emerging trends to address?
- How can we meet rising demand for training and encourage more integration of HPC training into university undergraduate / postgraduate teaching?

This document is structured as follows:

Chapter 1 – Introduction

Chapter 2 – Overview of the current status of HPC training provision

Chapter 3 – The status quo: a survey of training providers

Chapter 4 – A look to the future: perceived training needs of PATC course attendees

¹ <https://exdci.eu/>

² <http://www.prace-ri.eu/>

³ <http://www.etp4hpc.eu/>

⁴ <http://www.eesi-project.eu/>

Chapter 5 – Gaps in the current provision of training

Chapter 6 – EuroLab-4-HPC and the HPC Training Curriculum

Chapter 7 – Summary of challenges and Recommended Actions

Chapter 8 - Conclusion

2 Overview of the current status of HPC training provision

2.1 Who needs HPC training?

When we think of “the HPC community”, we tend to think of a small but nevertheless considerable population of computational scientists working in academia and using national (or regional) research facilities, and a smaller number of domain experts working in industry. Increasingly, we may also think of Data Scientists, working for commercial companies, processing enormous amounts of data about a company’s customers.

The so-called HPC community is in fact a disparate and fragmented collection of communities, and so training needs to be targeted at different groups of people and at different levels. Courses may be aimed at undergraduates, early-stage researchers, or experienced researchers or developers. They may be pitched at an academic audience, or at a specific commercial sector. Training is mostly aimed at end-users, but can also be needed for developers. A “one size fits all” approach cannot be used, but rather a more tailored programme must be created.

Appropriate HPC training needs to be available for all of the following:

- All career stages
 - Undergraduate students, early-stage researchers (postgraduates and post-docs), experienced researchers
- All levels of HPC experience
 - Existing users – beginners, intermediate and advanced – but also potential users
- Users of different types of facilities
 - Tier-0, Tier-1, Tier-2 facilities...
- Different sectors
 - Academic, industrial / commercial
- Different types of user
 - Developers, programmers, package users, data analysts...

Currently, most easily-accessible HPC training takes the form of short, non-assessed courses (typically lasting no more than a week), that are free of charge to academics. Courses are mostly given by HPC centres, sometimes in association with a hardware manufacturer or software provider. There are usually few pre-requisites for registering, beyond availability of places on the course, and perhaps the funding to travel to the course location. Participants are typically already within the HPC or CSE (Computational Science and Engineering) community, and most attendees already either have some experience of using HPC or have an identified future need to use it.

A small but growing number of European universities offer graduate and postgraduate courses in HPC, and it is also taught as part of some science and engineering courses. In order to create a workforce of graduates from a wide variety of disciplines who are already competent in HPC when they embark on their careers, HPC needs to be a core skill taught to all undergraduate students of CSE subjects.

Ideally, some efforts should also be made at a pre-university level to raise awareness of HPC, and many HPC centres are increasingly involved in outreach activities to engage with school pupils in a more informal way. This is addressed in more detail in D5.3 – Promoting HPC as a Career Choice.

2.2 How is HPC training currently delivered?

HPC training consists of a continuum of learning resources as varied as the community it addresses.

It includes general courses covering the underlying concepts and methodologies of HPC, and courses on HPC-specific technologies (e.g. a specific national service) or programming models and applications (e.g. MPI or OpenMP).

Courses may be short intensive courses (typically from one day to a week long), longer full-time courses (such as a Master's course), or an integral part of a bigger course (e.g. a module within an undergraduate degree).

They may be taught face-to-face, or – increasingly – online, in the form of video tutorials, webinars, or MOOCs (Massive Open Online Courses). Online courses may be synchronous (followed in real-time), or asynchronous (e.g. self-paced tutorials). There is also a large amount of documentation freely available online, including course slides, video recordings, and interactive exercises.

2.3 Overview of short-term, intensive training opportunities

2.3.1 PRACE

PRACE, the Partnership for Advanced Computing in Europe, is the coordinator of EXDCI. PRACE is probably the principal European co-ordinator of intensive HPC training courses targeted mainly at people who need to learn HPC skills for their current studies or work. PRACE training activities range from 1-hour webinars to the 2-month Summer of HPC programme.

In order to widen the reach of the courses as much as possible and provide the best opportunities for everyone, courses are open to all researchers in Europe, are free of charge to academics, and are usually given in English, the de facto lingua franca of the European research community.

- **PATCs – PRACE Advanced Training Centres**

Most PRACE training is delivered by the six PRACE Advanced Training Centres (PATCs)⁵, via short face-to-face training courses (typically lasting 2-5 days). Topics covered range from entry-level introductory courses to more specialised, in-depth topics. The PATCs operate a co-ordinated syllabus, which is offered across the continent throughout the year, to try to ensure a balanced timetable of courses.

PATCs typically offer a combined total of around 70-80 courses per year to a combined audience of more than 1500 attendees. Recent trends have seen a greater emphasis on data science and an increasing number of female participants (growing steadily from 12.9% in the 2012-13 PATC year to 19.6% in 2015-16).

⁵ <http://www.prace-ri.eu/prace-advanced-training-centres/>

While the majority of participants on such courses come from academia, there is growing participation (from 9.9% in 2012-13 to 14.9% in 2015-16) from commercial / industrial companies, ranging from SMEs to large multinational corporations, from a variety of sectors including energy, telecommunications and financial services. The courses that are in demand by attendees from commerce and industry include conventional HPC courses on parallel programming and performance engineering techniques, domain-oriented applications of HPC methods (often engineering simulation methods such as OpenFOAM), and high performance data analytics courses.

- **PTCs (PRACE Training Centres)**

A new complementary initiative will see the establishment of PRACE Training Centres, which will collectively provide a coordinated programme of training courses on HPC and computational science for European researchers. 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.

- **PRACE seasonal schools**

PRACE runs 3-4 seasonal schools⁶ each year. These week-long training courses usually focus on a specific domain or area of HPC programming. They take place all around Europe, but are mainly held in non-PATC countries in order to complement the PATC activities.

- **International HPC Summer School**

PRACE, together with Compute/Calcul Canada⁷, RIKEN⁸, and XSEDE⁹, jointly organises the International HPC Summer School on HPC Challenges in Computational Sciences¹⁰, an annual one-week summer school which takes place alternately in Europe and the USA. Leading scientists from multiple domains and HPC technologies are invited to teach on the summer school, which is aimed at postgraduates and postdoctoral scholars. The objectives are to expand participants' knowledge of HPC and its application in multiple fields of science and engineering, and to foster new friendships and partnerships among the presenters and participants, which is particularly encouraged through a mentoring programme. Demand for places has become highly competitive in recent years.

- **PRACE Summer of HPC**

The PRACE Summer of HPC¹¹ offers two-month student placements at HPC centres across Europe for senior undergraduates and early postgraduates. The programme attracts a slightly different demographic from the other PRACE training activities, as it targets younger researchers who often have not previously used or even considered the use of HPC facilities.

⁶ <http://www.prace-ri.eu/prace-seasonal-schools/>

⁷ <https://www.computecanada.ca/>

⁸ <http://www.riken.jp/en/>

⁹ <https://www.xsede.org/>

¹⁰ <http://www.ihpcss.org/>

¹¹ <https://summerofhpc.prace-ri.eu/>

By fully integrating students into an HPC centre for two months and giving them a real piece of work to do, the programme offers the participants a great insight into HPC work.

- **PRACE online learning**

The PRACE training portal¹² contains an extensive archive of training resources, such as video tutorials and material from courses, including slides as well as video recordings and / or topic-specific exercises where these are available. The number of visits to the training portal increased by 24% in 2015 compared to the previous year, and the number of users increased by 45%, showing that the demand for online training resources is growing significantly.

PRACE aims to intensify online training offerings by providing MOOCs (Massive Open Online Courses), and is launching two MOOCs in March 2017: “Supercomputing: Discover how supercomputers are powering scientific breakthroughs”¹³, and “Managing Big Data with R and Hadoop”¹⁴.

- **PRACE CodeVault**

The PRACE CodeVault¹⁵ is a repository for training codes, open to all students worldwide. It contains various code examples and model solutions of common HPC programming tasks. The code samples are released under open source licences and can be used both in training and as building blocks of real-world computing applications.

- **PRACE members: HPC centres in Europe**

The members of PRACE include many of the top HPC centres in Europe. As providers of many of Europe’s Tier-0 and Tier-1 HPC facilities, they have a remit to provide training for users and potential new users, and all have well-established training schedules with course timetables with a similar range of content to the PATC timetable, but these courses may be aimed at a more local audience (e.g. they may be taught in the local language, instead of in English). They provide a mix of face-to-face courses, themed workshops, online webinars, and archived training material. Face-to-face courses are open to all academics free of charge, and the online material is freely available to all.

2.3.2 *Other short training courses and internship-style training*

- **FETHPC projects / Centres of Excellence**

Some of the EC-funded FETHPC projects and Centres of Excellence (CoE) offer training in their specific areas of research as part of their core activities. These have a more specialised focus. Examples include BioExcel¹⁶, which provides training in HPC for Life Science research, ESCAPE¹⁷, which provides training in Numerical Weather Prediction topics as well as modern computational physics and energy efficient computing, and POP¹⁸, the Centre of Excellence in

¹² <http://www.prace-ri.eu/trainings/>

¹³ <https://www.futurelearn.com/courses/supercomputing>

¹⁴ <https://www.futurelearn.com/courses/big-data-r-hadoop>

¹⁵ <http://www.prace-ri.eu/prace-codevault/>

¹⁶ <http://bioexcel.eu/>

¹⁷ <http://www.hpc-escape.eu/>

¹⁸ <https://pop-coe.eu/>

Performance Optimisation and Productivity, which provides training on parallel programming in general as well as on specific performance tools, analysis methods and optimisation techniques.

- **HiPEAC Network of Excellence**

The mission of HiPEAC¹⁹, the European Network on High Performance and Embedded Architecture and Compilation, is to steer and increase European research in the area of HPC and embedded computing systems, and stimulate co-operation between academia and industry, as well as between computer architects and tool builders.

HiPEAC runs the ACACES (Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems) summer school²⁰. This one-week summer school aims to disseminate advanced scientific knowledge and promote international contact among scientists from academia and industry. A distinguishing feature of this Summer School is its broad scope ranging from low-level technological issues to advanced compilation techniques.

HiPEAC also organises an Industrial PhD Internship Programme²¹ whose primary aim is to enable European companies to identify highly skilled and exceptionally motivated research talent. Many members of the HiPEAC network use these internships as a cost-effective recruiting tool.

- **Events at major international supercomputing conferences (e.g. SC, ISC)**

The US-held SC conference²² (formerly Supercomputing) and the European-based ISC²³ (the International Supercomputing Conference) are the two most significant events in the year for European HPC stakeholders. Both events attract attendees from around the world, bringing together hardware and software vendors, staff from HPC centres, application developers, scientists and more. In addition to the exhibition aspect of the events, there are tutorials, workshops and Birds-of-a-Feather sessions (BoFs). Due to the large number of attendees at these events, HPC stakeholders have an unequalled opportunity to provide training sessions for a varied audience with whom they may not otherwise come into contact

- **Transnational Access research visits (HPC-Europa)**

From as early as 1993, some HPC centres offered Transnational Access visits to European researchers. From 2004 until 2012, there was one co-ordinated Transnational Access programme, HPC-Europa, which ran at six (and later seven) European HPC centres.

The Transnational Access programmes offered researchers of postgraduate level and above the opportunity to carry out a collaborative visit to a host research group in another country, working in a similar field, while benefitting from the HPC facilities at one of the centres.

Some of the participants already had considerable HPC experience, but many were novice HPC users or complete beginners. The Transnational Access programmes therefore offered a unique opportunity for many new HPC users to gain hands-on experience of some of Europe's most powerful HPC facilities, with the added benefit of having a named contact to provide HPC mentoring. Further, participants had easy access to training courses during their visits, and were

¹⁹ <https://www.hipeac.net/>

²⁰ <https://www.hipeac.net/acaces/>

²¹ <https://www.hipeac.net/mobility/internships/>

²² <http://www.supercomp.org/>

²³ <http://isc-hpc.com/>

often advised to schedule their visit to allow them to attend any relevant courses. Even the most experienced users would usually learn new skills during their visit, whether through using new architecture or trying more advanced options with their code.

The programme offered a unique combination of both training and research opportunities, with relatively low barriers to participation. As the visitors were removed from their usual daily distractions in their home institute, they were able to focus almost exclusively on their research work and the HPC aspects of it, and almost all found the period spent abroad to be exceptionally fruitful. While training was not a primary objective of the Transnational Access visits, it was clearly one of the major outcomes of most visits, and for some it represented a very useful first step into the world of HPC.

Over the 4 years since the programme ended, it has been greatly missed, demonstrated by the many people in the HPC community who have asked whether it will run again. The consortium is currently waiting for confirmation of funding and hopes that the programme will indeed begin again soon.

2.4 HPC training in higher education

2.4.1 *Postgraduate qualifications in HPC*

An increasing number of Master's courses in HPC are now offered at European universities. For example, the University of Edinburgh alone offers an MSc in High Performance Computing²⁴, an MSc in Data Science²⁵, and a combined MSc in High Performance Computing with Data Science.

Among others, Trinity College Dublin also offers an MSc in HPC²⁶, Swansea University has an MSc in High Performance and Scientific Computing²⁷, and Liverpool University runs an MSc in Big Data and High Performance Computing²⁸. While most of the postgraduate courses in these areas are offered by UK universities, there are others elsewhere, such as the Master in High Performance Computing offered jointly by SISSA (the International School for Advanced Studies in Trieste) and ICTP (Abdus Salam International Centre for Theoretical Physics)²⁹.

All of these courses are fairly new (the EPCC course was the first of its type in the UK, starting in 2001), and respond to an emerging need for graduates from diverse CSE backgrounds with a grounding in HPC and/or Data Science.

Most courses have both a set of compulsory courses, to give a broad-based foundation in HPC and data science, and a range of optional courses, which give students a chance to tailor their studies towards their particular interests. Further, some courses allow students to take a number of credits from other relevant courses offered by different departments of their institute.

The range of Master's projects available for students gives a feel for the variety of work to which HPC can be applied. Industry-based projects can be a great opportunity to have hands-on experience of a real-world project and to gain a variety of practical, transferable skills from the workplace environment. Students who complete successful projects with external organisations have gone on to be offered jobs afterwards; and even if this does not happen,

²⁴ <https://www.epcc.ed.ac.uk/msc/programmes/msc-programmes>

²⁵ <http://www.ed.ac.uk/studying/postgraduate/degrees/index.php?r=site/view&id=902>

²⁶ <https://www.maths.tcd.ie/hpcmsc/>

²⁷ <http://www.swansea.ac.uk/postgraduate/taught/science/msc-high-performance-and-scientific-computing/>

²⁸ <https://www.liverpool.ac.uk/study/postgraduate/taught/big-data-msc/overview/>

²⁹ <http://www.mhpc.it/>

having the experience of working within a company can also give students a clear competitive advantage when job-seeking.

2.4.2 *HPC training for undergraduates*

There is currently a general lack of HPC training in undergraduate curricula, even when HPC centres are located within universities. Often those designing and teaching the undergraduate curricula do not have contact with or awareness of the HPC centres, whose local users are more focused on research. There are some notable exceptions, such as the incorporation of the BSC PATC courses as accredited seminars in a compulsory module on the Barcelona School of Informatics (FIB) Master Programme in Innovation and Research in Informatics³⁰.

Another highly successful initiative is the NVIDIA GPU Educators Program³¹, which provides teaching materials and GPU resources for use in university classrooms and labs, with the aim of equipping students with the skills in deep learning and accelerated computing that they are likely to need in the future. The flagship GPU Teaching Kits³² are co-developed with academia for use in any academic discipline that benefits from accelerated computing. The three comprehensive packages (Deep Learning, Accelerated Computing, and Robotics) contain everything an instructor needs to teach a full-term curriculum course with GPUs, including lecture slides, lecture videos, hands-on labs and coding projects, and source code solutions. The global community of academics who use the Teaching Kits is supported by an online Teaching and Curriculum Forum. The Educators Program shows how industry can work with academia for the benefit of both, and offers a good example of how technology providers can help to foster in the next generation of workers the skills that are needed to exploit their technology.

Integrating HPC courses into undergraduate curricula should be a major focus of efforts to expand HPC training. It has been argued that students in all CSE domains should receive HPC training, as this would not only equip many more students with basic HPC skills, which would be of great benefit to employers, but it might also inspire more students to enrol for dedicated postgraduate courses or actively pursue careers in HPC-related areas.

2.5 **Marie-Curie Innovative Training Networks (ITNs)**

Marie Curie Innovative Training Networks (ITNs)³³ are intended as networks of institutions and enterprises that exchange students and/or researchers working in a given scientific domain.

The networks bring together universities, research centres and companies from different countries, and targets well-identified multi- and interdisciplinary needs in scientific and technological research areas.

Potentials for improving interdisciplinary HPC training and work are huge. Specifically, these networks provide an opportunity for Higher Education institutions with existing courses and competences to share their knowledge and resources with researchers working in various scientific disciplines, enabling them to conduct their research using HPC tools.

Interdisciplinary workshops, summer schools, and similar training events that focus on applied use of HPC in given scientific domains can be organised within such Marie Curie ITNs, paving the way for development of formal study courses or programmes with an HPC orientation, as

³⁰ <http://www.fib.upc.edu/en/masters/miri.html>

³¹ <https://developer.nvidia.com/educators>

³² <https://developer.nvidia.com/teaching-kits>

³³ http://ec.europa.eu/research/mariecurieactions/about/innovative-training-networks_en

well as for new research opportunities that stem from use of supercomputing in essentially all scientific disciplines ranging from natural and life sciences to social sciences and humanities.

Over a 3-year contract, participating researchers are exposed to different sectors, and gain a comprehensive set of transferable skills such as entrepreneurship and communication.

There are three types of Innovative Training Networks:

European Training Networks: Joint research training, where the researcher experiences different sectors and develops their transferable skills. These research positions are offered by at least three partners from at least three different European countries, from inside and outside academia, and do not necessarily lead to a degree award.

European Industrial Doctorates: Joint doctoral training, where the researcher is jointly supervised by at least one academic and at least one non-academic partner from at least two European countries, and spends at least 50% of their time with the non-academic supervisor.

European Joint Doctorates: Joint, double or multiple degrees, offered by a network formed by a minimum of three academic organisations.

Innovative Training Networks have a strong emphasis on mobility, requiring participating researchers to split their time between different sites. The networks are international, interdisciplinary and intersectoral and aim to equip participants with transferable skills that can be applied to both the public and private sectors.

2.6 On-line training opportunities

Although the opportunities for training in HPC-related subjects are always increasing, there are many people for whom attending a year-long Master's course, or even a face-to-face course of a few days' duration, is not an option due to other commitments.

However, at the same time as HPC is becoming more prevalent and demand for training is increasing, so too is the availability and capability of technology to provide on-line training. This is an area which is rapidly increasing, but still has much scope for expansion.

On-line training can take many forms: video tutorials or webinars broadcast live online, or recordings of these which can be watched at a later date; accredited distance learning courses delivered using a Learning Management System, MOOCs, online interactive exercises for self-paced learning, or simply sets of course documents (slides, course notes/handbooks, etc) which can be downloaded for self-study. On-line courses may be synchronous or asynchronous, and may or may not lead to accreditation.

A good example of how training can be made online for people who have not been able to attend in person is given by Argonne National Laboratory³⁴, which runs the Argonne Training Program on Extreme-Scale Computing (ATPESC), an annual two-week programme for early-career researchers. There is great demand for the 65 available places on the programme, but the lectures are made openly available on the Argonne YouTube Training Channel³⁵ so that anyone can, in their own time, watch any of the 76 hours of lectures from some of the world's foremost experts and pioneers in extreme-scale computing.

A number of initiatives have collated HPC training content that is freely available, to make it easier to find.

³⁴ www.anl.gov

³⁵ <http://extremecomputingtraining.anl.gov/2016-videos>

- The **EXDCI Training Portal**³⁶ lists training courses provided by EXDCI members and associated organisations, such as FETHPC projects and CoEs. It also provides links to other sites which provide HPC training resources.
- The **HPC University**³⁷ is a virtual organisation of mainly US-based partners, but which includes PRACE and the Cyprus Institute³⁸. It seeks to provide a cohesive, persistent, and sustainable on-line environment for training materials in all HPC computing environments.
- The **Supercomputing Training Portal**³⁹ was a European-funded project which collected and curated HPC training content and provided access to an educational HPC cluster for the purpose of engaging in hands-on HPC education and training.
- **XSEDE**⁴⁰, the US-based Extreme Science and Engineering Discovery Environment, has links to some online tutorials, and maintains a list of live, online and asynchronous courses, ranging from 1-hour courses to several days.
- The **International HPC Training Consortium**⁴¹ maintains a list of training resources with links to more training providers, tutorials and documentation.

There are also some domain-focused organisations which offer training, including HPC-related courses, such as **TeSS, the Training e-Support System for Elixir**⁴², a distributed infrastructure for life-science information. TeSS offers training events in computing, statistics and bioinformatics, as well as a repository of training materials.

2.7 The Software Sustainability Institute (SSI)

The Software Sustainability Institute (SSI)⁴³ in the UK should also be mentioned, although their focus is more on software development and sustainability in general rather than HPC specifically. SSI recognises that there is a growing requirement for researchers to write and maintain software, but that they receive little training in software development. SSI advises on training curricula for computational skills, analysing training needs and then identifying the right kind of training.

SSI also runs Software Carpentry and Data Carpentry workshops, which teach software development and data handling skills respectively, to enable researchers to be more productive and to make their research robust and reproducible. The demand for these workshops greatly exceeds SSI's capacity to deliver them, so a two-day intensive instructor training workshop was introduced, to increase the pool of certified workshop instructors. In addition to scheduled runs of this workshop, it can be run on demand for a particular group, and an on-line version also exists.

³⁶ <https://exdci.eu/jobs-and-training/training-portal>

³⁷ <http://hpcuniversity.org/>

³⁸ <https://www.cyi.ac.cy/>

³⁹ <http://supercomputing.cyi.ac.cy/>

⁴⁰ <https://www.xsede.org/training1>

⁴¹ <https://sites.google.com/a/lbl.gov/hpc-training-best-practices/home>

⁴² <https://tess.elixir-europe.org>

⁴³ <https://www.software.ac.uk/>

3 The status quo: a survey of training providers

In order to gain a deeper insight into the training activities of the PRACE and ETP4HPC partners, EXDCI sent an online survey to 165 individuals from 86 academic and industrial organisations around Europe, including IT providers and other commercial companies, higher education institutions, research institutes, non-profit technology transfer companies, and national HPC centres.

Twenty-nine responses were received, a slightly larger sample size compared to the previous HPC training provision survey carried out by ETP4HPC, which received 24 responses and was nevertheless considered sufficiently representative. The ETP4HPC survey used two questionnaires, one investigating the needs of the technology provider community, and the other tailored to existing providers of HPC education. These contained questions on the education and training provided to employees, both in-house and from external providers. The results informed the “Report on HPC Education and Training” published in 2014.

Based on the recent “Research and Education in Computational Science and Engineering” report⁴⁴, the EXDCI survey aimed, among other things, to verify the extent to which current training provision covers nine HPC topics that are increasingly in demand. These topics are: Real-time Supercomputing, Embedded Supercomputing, Simulation-based Training and Education, Quantitative Performance Analysis of Algorithms and Software, Performance Engineering and Co-design, Ultrascaleability and Asynchronous Algorithms, Power Management (Power Wall Issues), Fault Tolerance and Resilience, and Big Data.

The responses to the EXDCI survey are summarised in the tables below.

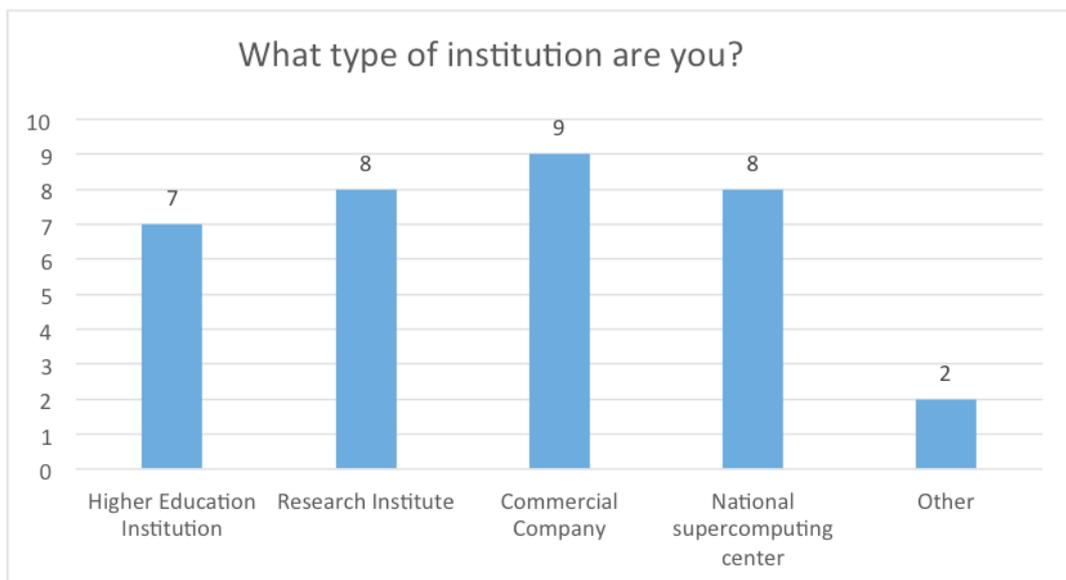


Figure 1 Survey response: Institution type

⁴⁴ <https://arxiv.org/abs/1610.02608>

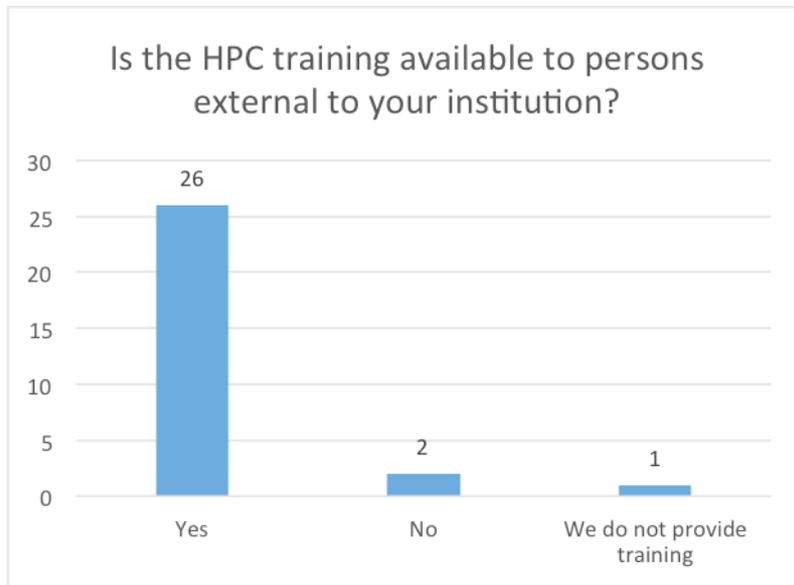


Figure 2 Survey response: Courses for externals

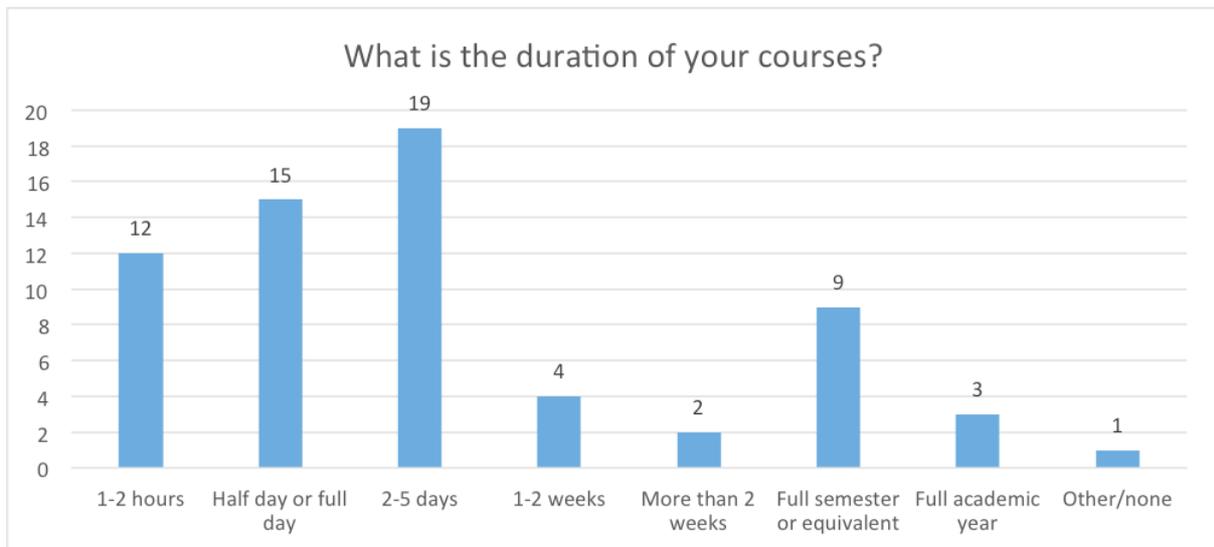


Figure 3 Survey response: Course duration

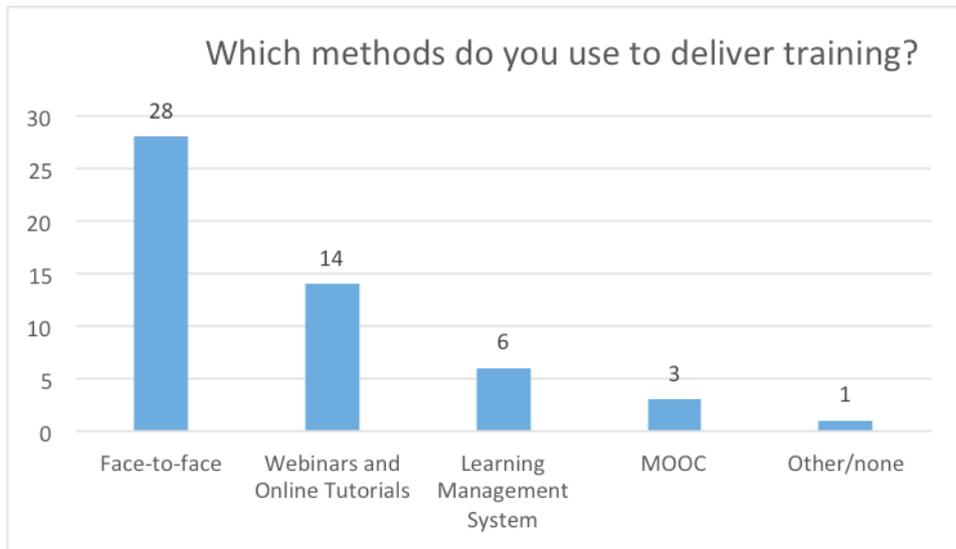


Figure 4 Survey response: Delivery method

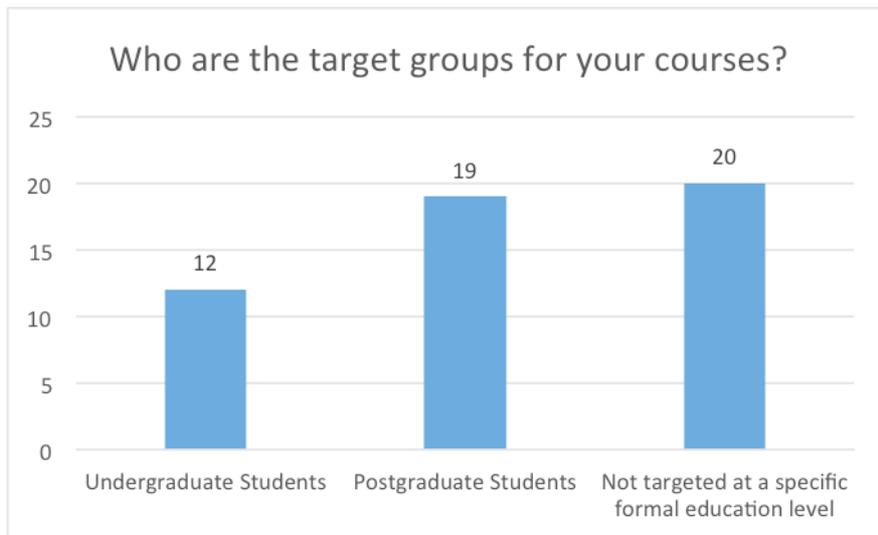


Figure 5 Survey response: Target group

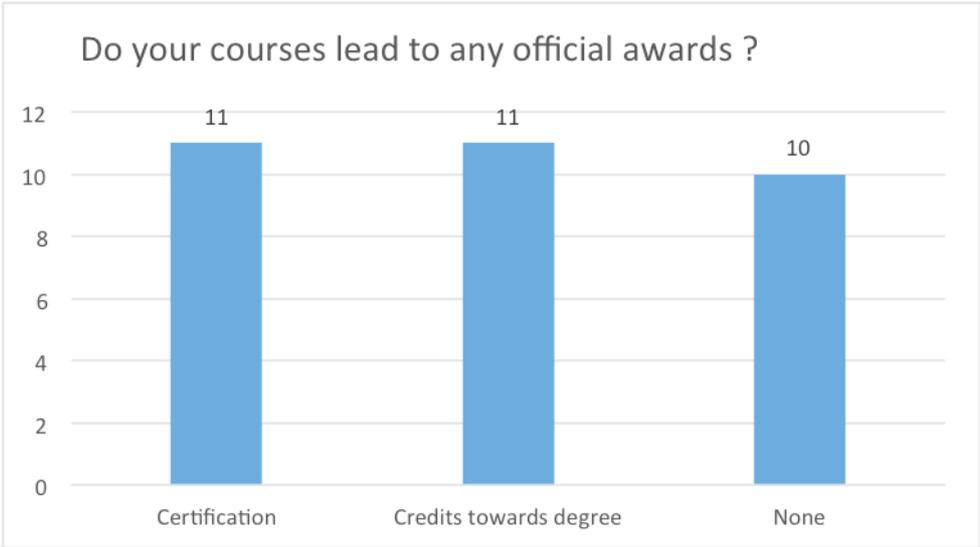


Figure 6 Survey response: Course accreditation

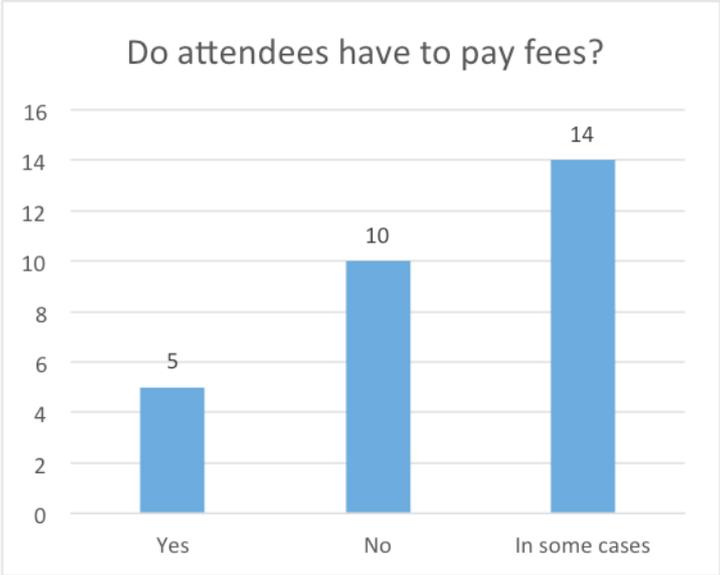


Figure 7 Survey response: Course fees

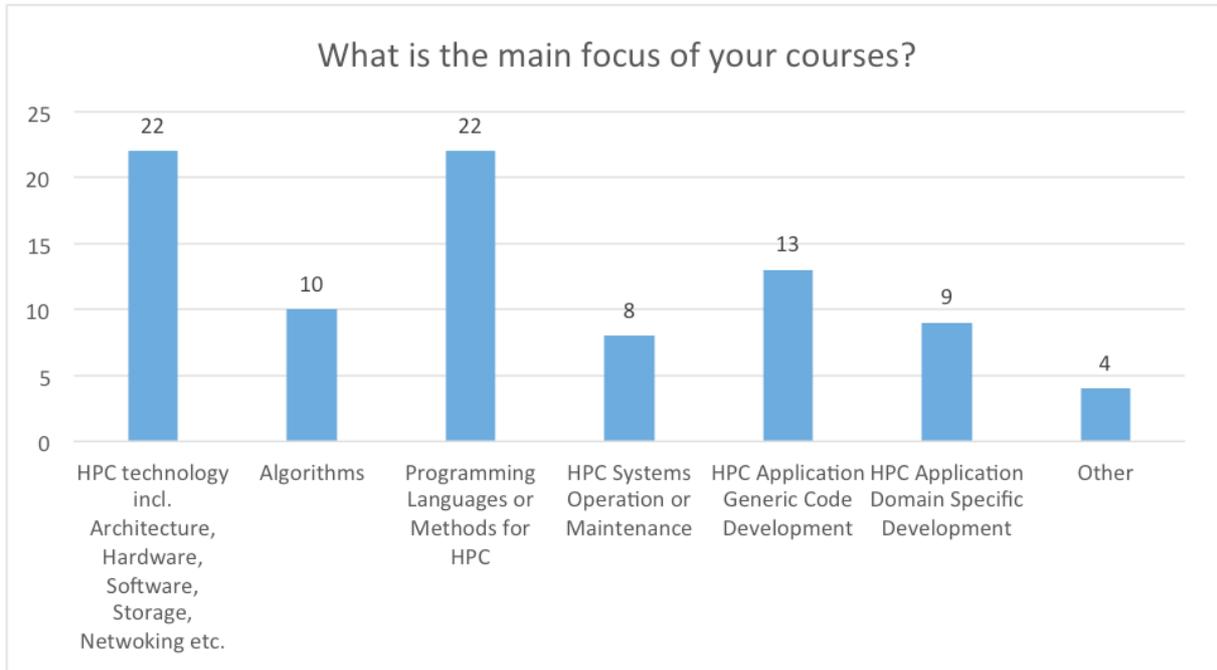


Figure 8 Survey response: Course focus

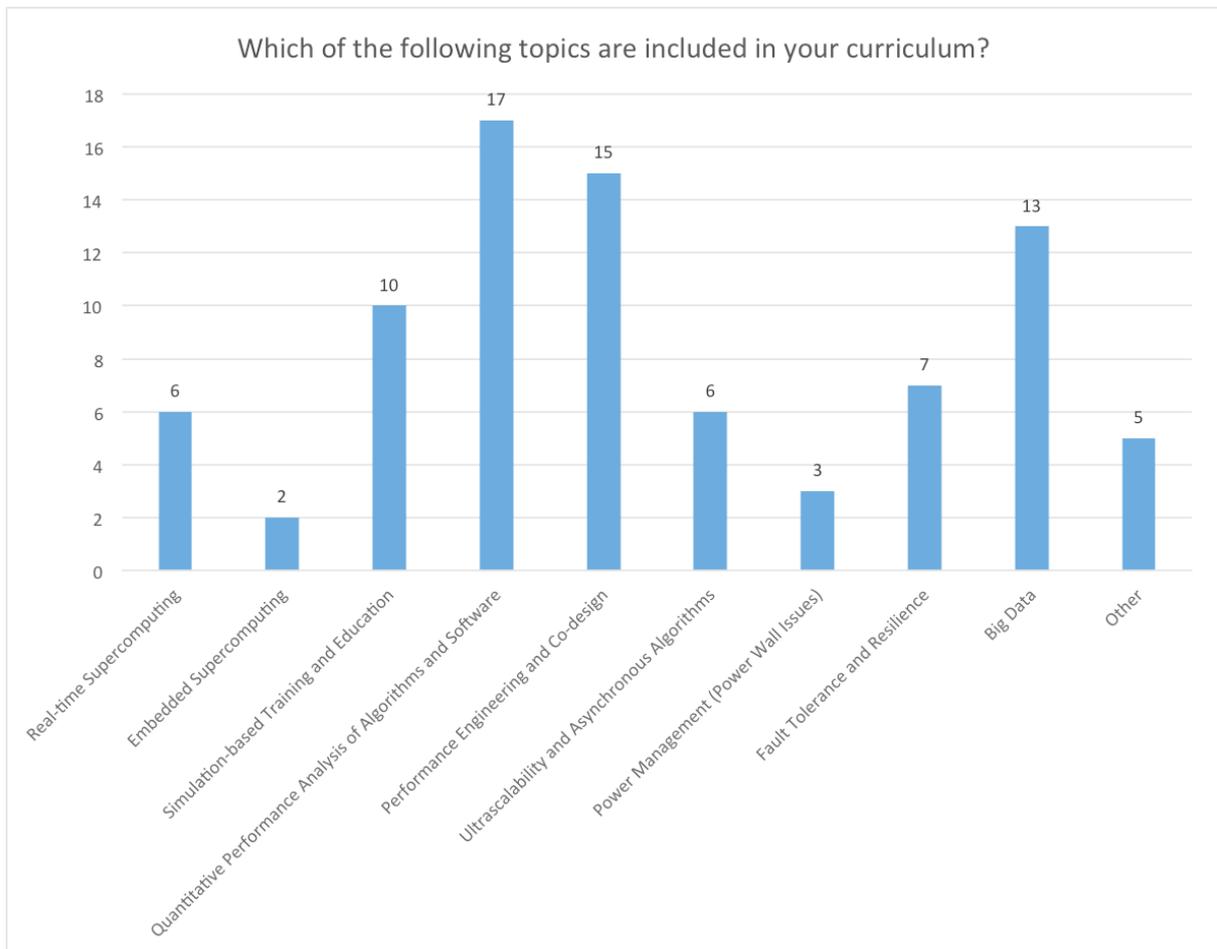


Figure 9 Survey response: Emerging topics

The survey shows that most respondents provide training to external people – only 2 provided training exclusively to internal people, and one did not currently provide any training.

Most courses were short courses of a week or less, with a few courses lasting a semester or academic year.

Most training is still given face-to-face, with webinars and on-line tutorials in second place. The number of courses taught using Learning Management Systems or MOOCs is small, but this could be expected to grow in the coming years as the technology becomes more widespread.

The majority of courses were targeted at postgraduate level, or at no specific level, with fewer being targeted at undergraduate level – although the number (12 respondents – over 40%) was perhaps higher than might have been expected.

There was an almost equal split between courses leading to certification, credits towards a degree, or no formal award.

For most courses, fees were payable in some cases, but in only 5 cases were fees always charged.

The most common topics are HPC Technology (including architecture, hardware, software, storage, networking) and Programming Languages or Methods for HPC, but there was a wide spread of topics.

The emerging topics most covered were Quantitative Performance Analysis of Algorithms and Software, and Performance Engineering and Co-design, followed by Big Data and Simulation-based Training and Education, but again there was a wide spread of topics covered.

This survey shows the great variety of HPC training topics, teaching delivery methods and course formats and target audiences that make up the HPC training landscape. This variety is essential in order to be able to provide training in an accessible manner to all those who need it.

4 A look to the future: perceived training needs of PATC course attendees

PATC course attendees are asked to complete a questionnaire at the end of their course. In addition to providing feedback on the courses, they are asked to assess their future training needs in a variety of areas, as follows:

- In the future, I will need training in general HPC programming (MPI, OpenMP)...
- In the future, I will need training in advanced HPC programming (Hybrid MPI-OpenMP; next-gen HPC languages e.g. PGAS; GPU computing e.g. CUDA)...
- In the future, I will need training in code optimisation and performance analysis.
- In the future, I will need training in porting of existing codes to HPC architectures.
- In the future, I will need training in specific HPC application(s).
- In the future, I will need training in HPC programming and applications specific to my research community.
- In the future, I will need training in visualisation techniques.
- Are there some other fields of training you feel PRACE should provide training events in?

While the perceived needs of course attendees may not reflect the wider HPC community as a whole, and completion of the questionnaire is not obligatory and so responses may not even

accurately represent all participants, the views are those of people who have already had a need for HPC training and so should be considered as a good indication of future needs.

The responses are shown below in Figure 10.

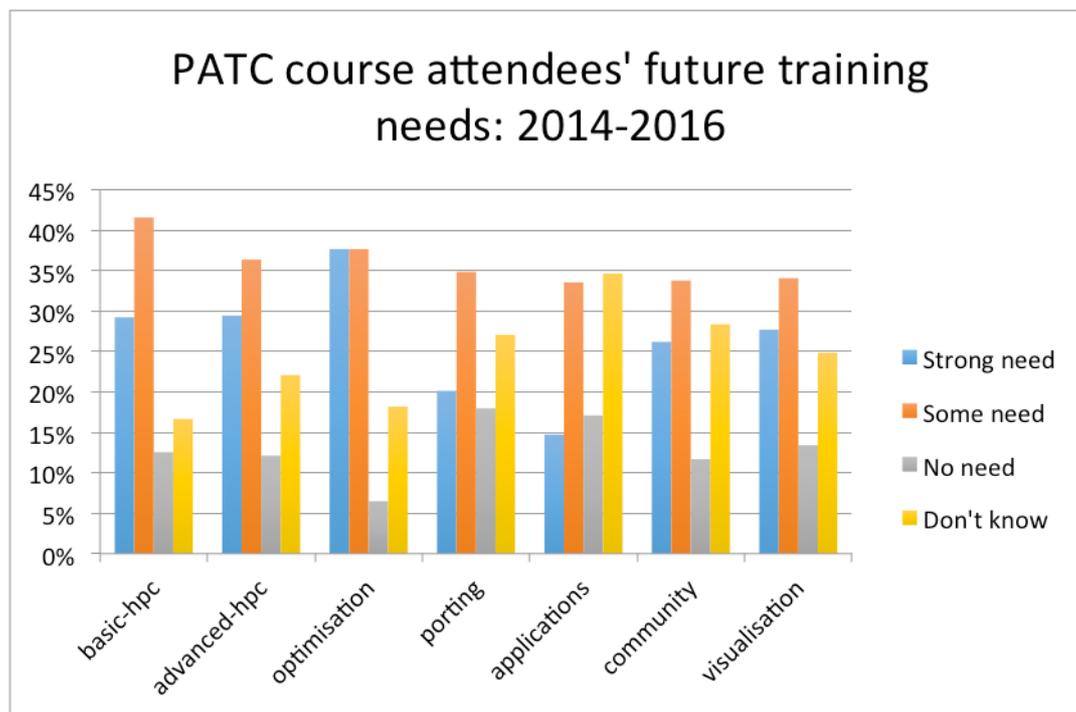


Figure 10 PATC course attendees' future training needs (2014-16)

More than a quarter of respondents felt that they would have a strong need in the future for training in: optimisation (38%), basic HPC and advanced HPC (29%), visualisation (28%), and community-specific training (26%).

More than a third of respondents also reported some need for training in each of these areas, and only around one sixth of respondents (varying from 12% to 18%) felt that they would have no need for training in each area – with the exception of optimisation, where only 6% felt that they would have no future need for training.

It should be noted that between a sixth and a third of respondents felt unable to assess their future training needs in each area, which may reflect the fast-changing nature of HPC, or could be due to many of the PATC course attendees being relatively early-stage HPC users who do not yet have a clear idea of how they will use HPC and which areas will be important to them in the future.

We can compare the answers given to each question in 2014 to those given in 2016, to try to ascertain whether future training needs are evolving.

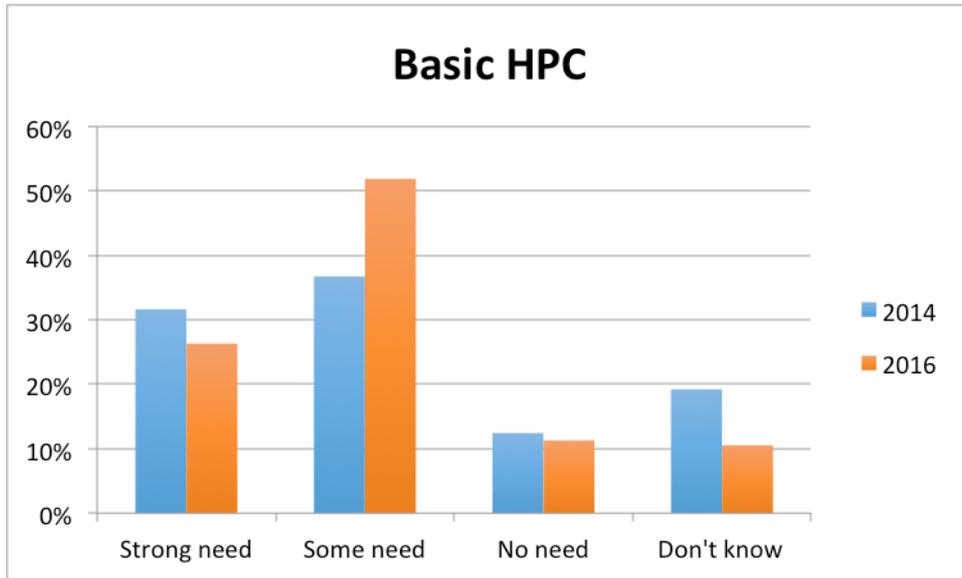


Figure 11 Comparison of perceived future training needs in basic HPC: 2014 vs. 2016

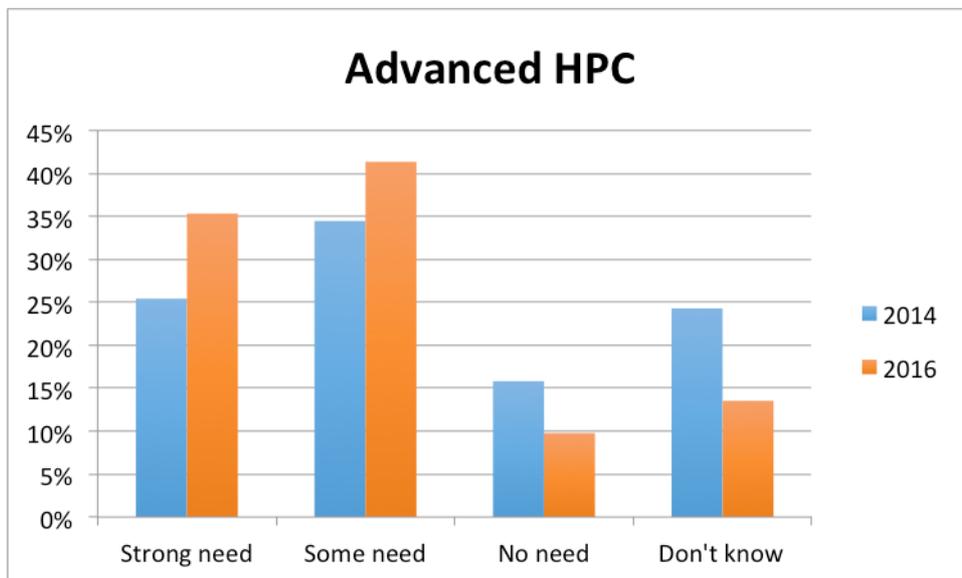


Figure 12 Comparison of perceived future training needs in advanced HPC: 2014 vs. 2016

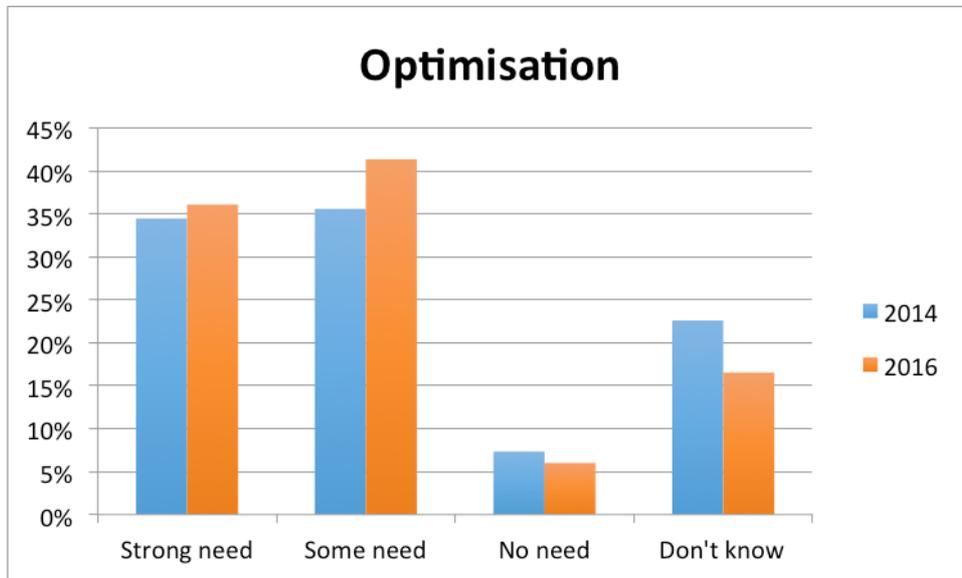


Figure 13 Comparison of perceived future training needs in optimisation: 2014 vs. 2016

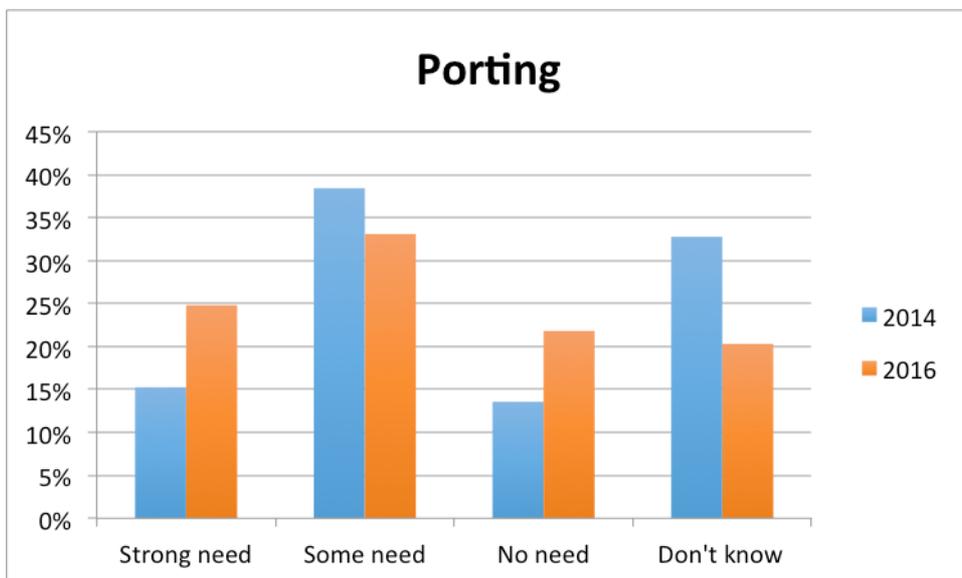


Figure 14 Comparison of perceived future training needs in porting code: 2014 vs. 2016

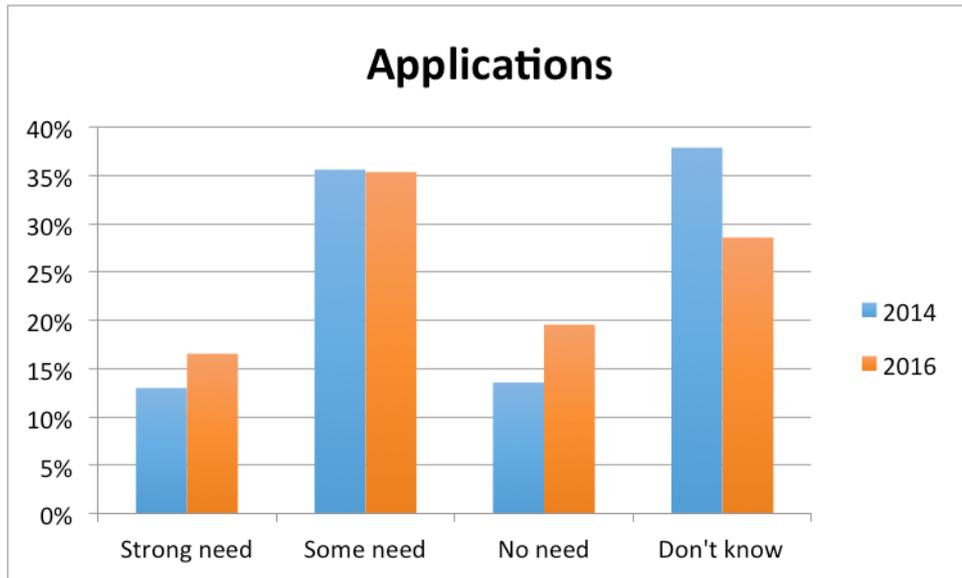


Figure 15 Comparison of perceived future training needs in HPC applications: 2014 vs. 2016

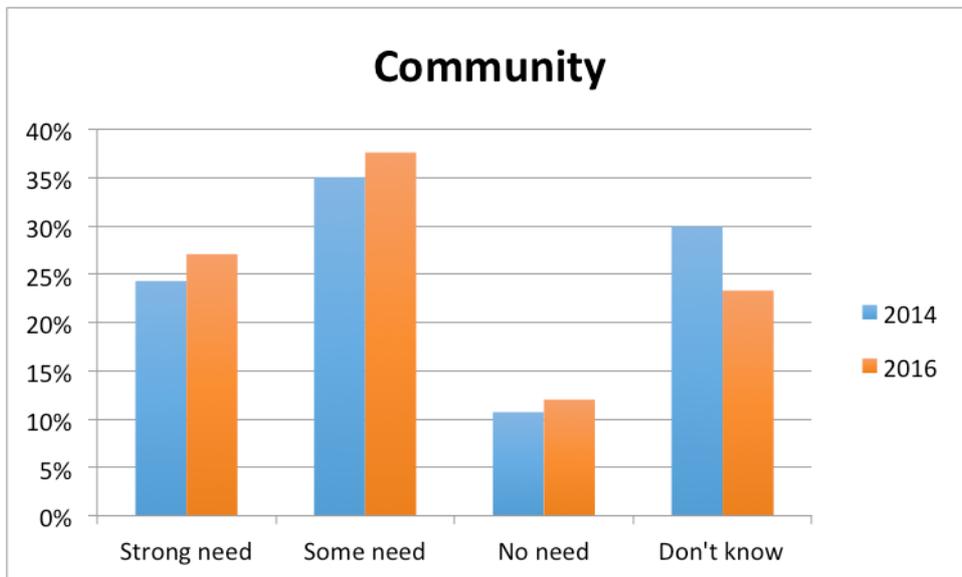


Figure 16 Comparison of perceived future training needs in community-specific areas: 2014 vs. 2016

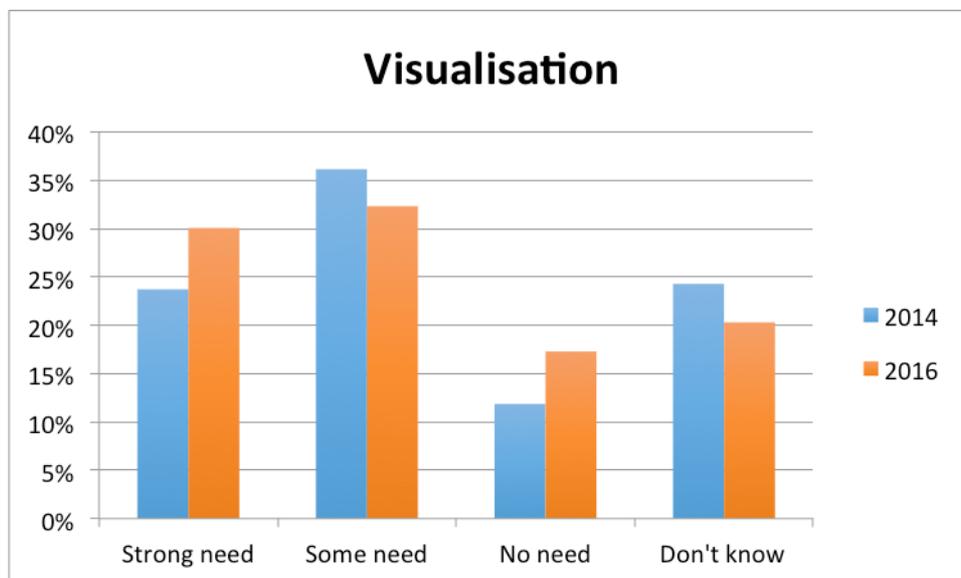


Figure 17 Comparison of perceived future training needs in visualisation: 2014 vs. 2016

From these comparative tables, we can see that in 2016, more people reported some need for basic training in HPC, while fewer reported a strong need for this. This may be due to more respondents having already undertaken some basic training.

There has also been an increase in the number of people reporting some need or a strong need for advanced HPC training, which again may reflect an increased number of people having acquired some basic HPC knowledge, and needing to extend their knowledge.

In some areas, such as optimisation, applications and community-focused training, although overall demand has increased, there is little difference between 2014 and 2016. Meanwhile, porting and visualisation have seen a shift from some need to a strong need.

Although there is little major difference between the figures for the two years, there is a consistent increase in some need or strong need for future training in every area. Although the number of people responding “don’t know” remains relatively high, there is a consistent, and often considerable, decrease between the number of people answering “don’t know” and those answering “no need” in each year, suggesting that people are becoming more informed about their future needs, which in turn suggests that people have a better general understanding of what HPC can do and how it can help them with their research.

Among the responses to the question, “Are there some other fields of training you feel PRACE should provide training events in?”, the most common answers included:

- Big Data / Data management / Data analysis
- Bioinformatics (linked to Big Data)
- Advanced MPI
- CFD / Application of MPI to CFD / OpenFoam and C++
- Techniques for handling IO of large data in HPC
- Introduction to Python (run more frequently)
- Good practice in coding
- File I/O and profiling tools
- Parallel algorithms
- Visualisation

5 Gaps in the current provision of training

5.1 Training undergraduates

Although HPC is still a relatively small sector, we can see from the above that much is currently being done to provide a wide variety of training opportunities in as flexible a fashion as possible, in order to make it accessible to the widest possible audience. There are important initiatives, driven mainly by PRACE, to co-ordinate the provision of HPC training across Europe, to make it as easy as possible for the greatest number of users and potential users to attend the training that they need, in a reasonably convenient location – whether this is entry-level courses on general HPC concepts, or more specific topics at an advanced level. Academic researchers in particular benefit from these courses, as all courses are open to them, free of charge.

Despite the wide provision of HPC training and the variety of formats in which it is provided (short intensive courses to full MSc programmes, in-person tuition versus on-line learning, etc), the European workforce still suffers from a lack of people with appropriate HPC skills.

In part, this is due to the growing demand for these skills. Multicore systems are now ubiquitous, and the number of processing cores per microprocessor is firmly on an upward trend. There is a wider interest in supercomputing and parallel computing generally, especially in industry, and HPC is being adopted by ever more research domains and commercial enterprises.

However, many of the existing training opportunities are targeted towards the academic sector, and in particular, towards those with an established need for training. The people who attend the courses are likely to be existing users of HPC facilities, or close associates of existing users, such as postgraduates in a computational science research group whose members already use the facilities. In most cases, people who hear about the courses will already be working in a field to which HPC can be applied.

One of the main issues identified with HPC education is that it starts too late, and for many their first opportunity only comes when they reach postgraduate level. By this stage, researchers are primarily led to HPC simply as a result of choices they have already made. However, if HPC was taught more widely at undergraduate level, perhaps the reverse would happen, and an interest in HPC would lead more students to choose to continue pursuing careers in Computational Science and Engineering, either in academia or industry.

In the PRACE Training Survey, only 20% of respondents had received training in parallel programming at university. Of course we cannot be sure how representative this figure is of CSE students as a whole. As HPC evolves and becomes more prevalent, and training opportunities increase in turn, the past undergraduate experience of today's postgraduate students may not accurately reflect the current experience of today's undergraduates - the undergraduate experience of more senior researchers will be even less representative of the picture today. However, there is likely to be a correlation between those who have had some previous exposure to parallel programming, and those who go on to work in a field which requires HPC skills. Therefore, if anything, this figure probably shows a higher proportion of students receiving undergraduate training than is reflected among CSE students as a whole.

The PlanetHPC roadmap ("Challenges facing HPC and the associated R&D priorities: a roadmap for HPC research in Europe"⁴⁵) established that:

⁴⁵ <https://issuu.com/epccedinburgh/docs/roadmap>

“Graduate and post-graduate training in HPC should be further encouraged. The capabilities of HPC should be taught to students of other disciplines; knowledge of this tends to be restricted to students of physical sciences and engineering.

The intention should not be to teach all students to program supercomputers, but rather to raise awareness of the capabilities of HPC as a tool for research.”

Promoting HPC at an even earlier level, to high school pupils and even primary school pupils, is an obvious extension of targeting the future workforce. Several initiatives already target school-age pupils.

- HPC centres attend science festivals aimed at the general public, and often in particular at children. For example, EPCC regularly undertakes outreach activities at events both big and small, such as: the Edinburgh Science Festival⁴⁶, the world’s first public celebration of science and technology and still one of Europe’s largest; Bang Goes the Borders⁴⁷, a free science festival aimed at families; and the Big Bang UK Young Scientists and Engineers Fair⁴⁸, the largest celebration of science, technology, engineering and maths (STEM) for young people in the UK.
- Using Raspberry Pi units or similar, mini supercomputers, such as EPCC’s Wee Archie⁴⁹ and HLRS’s Konni⁵⁰, can be used in outreach visits to schools and other public science outreach events to explain what a supercomputer is. Instructions can be developed so that individuals or classes can build their own similar systems.
- “The Beauty and Joy of Computing”⁵¹ is a non-majors university course developed at UC Berkeley that is also offered as a College Board endorsed AP Computer Science Principles course in over a hundred high schools across the US. Its content includes HPC Big Ideas of concurrency and distributed computing. This is discussed in detail in a SIGHPC webinar by Dr Daniel Garcia (University of California Berkeley): “Bringing HPC Big Ideas to high school students and learners worldwide”⁵²

Raising awareness of HPC at school level is an important objective, but outreach rather than formal training should be the focus of any activities, at least until there is better provision of training at the undergraduate level to equip students with HPC skills to take them on to postgraduate or commercial research.

5.2 Industrial engagement

HPC uptake is still relatively low in industry and commerce, and for small to medium sized enterprises (SMEs) in particular the barriers in terms of expertise and the cost of facilities are difficult to overcome without initiatives such as the PRACE SHAPE programme, described later in this section.

As identified in the ETP4HPC Report on Education and Training, much of the current HPC training provision is focused on the needs of HPC system user communities, rather than on

⁴⁶ <http://www.sciencefestival.co.uk/>

⁴⁷ <http://www.bgtb.org.uk/>

⁴⁸ <https://www.thebigbangfair.co.uk/>

⁴⁹ <https://www.epcc.ed.ac.uk/discover-and-learn/resources-and-activities/what-is-a-supercomputer/wee-archie>

⁵⁰ https://www.mathe-im-advent.de/fotos_cray/

⁵¹ <http://bjc.berkeley.edu/>

⁵² https://www.youtube.com/watch?v=eqbgqSe_hvU

technology suppliers and infrastructure providers. This means that training material and course delivery tends to be oriented towards academic needs and schedules.

Training provided by commercial enterprises (such as hardware and software vendors) tends to be closed to people external to the company, or expensive to attend, unless they are given in association with an academic-focused organisation – such as courses given by Cray or Allinea in association with PATCs or other HPC centres. Commercial organisations may run training sessions at events such as SC and ISC, but there are significant costs associated with attending these events in the first place (travel costs, conference fees, and perhaps more significantly, staff time).

PATC courses, and many others offered by HPC centres, are free to academics, who make up the vast majority of attendees. However, the courses are usually open on the same terms to non-academics (ie the same prerequisites and priority for places apply), the only difference being that they may be required to pay a registration fee. Some courses are targeted specifically at industry, and this helps to foster collaboration with academics. Most HPC centres are also open to scheduling their courses on demand for non-academic institutions, where demand warrants it. However, overall, non-academic participation in courses remains low.

The commercial sector is harder to target as they are generally not users of the PRACE services or national facilities. PRACE course announcements are sent to a list of over 2000 people, but these are mainly people from the academic sector who are already HPC users and / or have previously attended other PRACE training courses.

Further, the requirements of the commercial sector are generally very different. Firstly, they may require basic training on the benefits of HPC – something which is largely taken for granted within the academic community. They are more likely to use commercial packages and have domain-specific requirements.

Employees of commercial organisations who do not have in-house HPC training may also find that either the non-academic fees charged or the delivery of the courses in blocks of 2-3 days present barriers to their attendance. Most manufacturing companies in Europe are SMEs and as such they generally have no research or advanced computing capability. Moreover, they need a direct and short path to achieve a return on investment.

Research has shown that 97% of industrial companies that employ HPC consider it indispensable for their ability to innovate, compete and survive. However, PlanetHPC⁵³ ran an intensive campaign to find out HPC's potential in a wide spectrum of European industries by contacting trade and research associations, and the results showed that the capabilities of HPC are not well enough understood by many industries, with many thinking that HPC is simply not important for them.

Low levels of HPC expertise, and the perceived difficulty of accessing such expertise from external sources, are contributing factors to the poor uptake of HPC in SMEs in particular. To address this, some HPC service providers are now expanding their business model to an “HPC as a full service” approach, providing not just access to HPC capabilities, but also to higher level services. One example of this is the InnoHPC programme (High-Performance Computing for Effective Innovation in the Danube Region)⁵⁴, funded through the European Regional Development Fund (ERDF)'s Interreg Danube Transnational Programme. This aims to promote HPC use among SMEs by creating a transnational HPC laboratory to pool regional HPC infrastructure and competencies.

⁵³ <http://cordis.europa.eu/fp7/ict/computing/factsheets/planethpc.pdf>

⁵⁴ <http://www.interreg-danube.eu/approved-projects/innohpc>

More industrial outreach programmes to promote HPC need to be established to encourage the uptake of HPC in industry, and existing connections should be better exploited. For example, the PRACE SHAPE programme⁵⁵ (the SME HPC Adoption Programme in Europe) represents an ideal route for understanding the training needs of SMEs. SHAPE aims to raise awareness and equip European SMEs with the expertise necessary to take advantage of the opportunities for innovation that HPC brings to increase competitiveness. SHAPE aims not only to provide access to HPC facilities, but to define workable and sustainable HPC solutions adapted to the resources and staff skills of the enterprise.

MOOCs should be publicised as widely as possible to the commercial sector – the first PRACE MOOCs require just 3 hours per week over 5 weeks, and the attendee simply joins from their own desk.

Access could also be widened by co-locating training with other events. This is already achieved successfully at major events such as SC and ISC, with tutorials and workshops given by HPC technology providers and service providers alike, and a genuine effort to bring together all the stakeholders where possible. However, these events are to a large extent still “preaching to the converted”, and efforts should be made to identify suitable occasions to promote the benefits of HPC and the availability of training to new audiences, potentially by proactively seeking opportunities to present training slots at domain-focused conferences and summer schools.

6 EuroLab-4-HPC and the HPC Training Curriculum

EuroLab-4-HPC⁵⁶ aims to establish an HPC systems curriculum⁵⁷ for technology leaders in HPC systems, and to develop training practices based on MOOC technologies to train future technology leaders. The curriculum combines courses that can be offered both in traditional form and on-line.

The purpose was not to propose a detailed, specific curriculum, but rather to provide a basis, or framework, for the development of such curricula. The structure is flexible, with a suggested core set of courses and a list of optional courses to address specific educational targets. A suggested course description is provided, with a one-paragraph overview of the topics that should be covered on each course.

The curriculum is designed to be flexible enough to be adapted for different course durations, background of attendees and specific educational goals. The curriculum describes three example scenarios, which encompass 1-year and 2-year MSc programmes.

The curriculum is designed to be suitable for online delivery as well as face-to-face teaching, and includes a section on Best Practices for Online Education, focusing on extended MOOCs (xMOOCs). While these were identified as the most practical online education methodology, they are also characterised by very low completion rates, and the curriculum seeks to address this. Guidelines for implementing a MOOC are also included.

⁵⁵ <http://www.prace-ri.eu/hpc-access/shape-programme/>

⁵⁶ <https://www.eurolab4hpc.eu/>

⁵⁷ <https://www.eurolab4hpc.eu/static/deliverables/D3-1--final-HPC-curriculum.95d306191a15.pdf>

7 Summary of challenges and Recommended Actions

Challenge 1: Coping with rising demand

- **Increased use of online learning**, especially MOOCs, is the clearest way to increase the reach of courses. However, MOOCs require significant investment in terms of effort; a quicker and simpler way to extend the reach of training is to record all tutorials and courses whenever possible, and make them available online, e.g. via a dedicated YouTube channel such as that of the Argonne National Laboratory⁵⁸.
- **More summer schools** should be developed, as demand outstrips the capacity of existing summer schools. Summer is a good time for students (including postgraduates), as with term-time demands on hold, they have time to investigate subjects which may interest them but are less immediately essential. Summer schools are ideal as most students enjoy both the travel opportunities and the community spirit and are enthusiastic about attending.
- **Collaboration** between training providers can lead to the sharing of materials, joint development of new materials, contribution to shared repositories, and the sharing of best practices. This can result in more material being produced and being made accessible, for less effort.
- **“Train the trainers” approach:** Following the model of the Software Sustainability Institute, there could be benefits in developing one or more packaged courses covering general interest material, and an associated certificated instructor course. By increasing the training capacity, this can allow better penetration into specific communities which need to increase their HPC uptake and which have good community networks in place (e.g. in a specific domain, or for users of a specific package). The expectation is that the course then only needs to be delivered once to a new community, and is self-sustaining within the community thereafter.
- **Transnational Access:** reinstating a European Transnational Access programme for High Performance Computing would bring back a unique learning opportunity which is particularly valuable for new HPC users.

Challenge 2: Addressing the widening target audience

- **Target users of specific community codes** with growing HPC use. This could mean publicising existing courses or developing new targeted course material.
- **Increase co-location of courses with other events** – e.g. provide workshops and tutorials at domain-focused conferences. This could be one method of targeting users of community codes.
- **More industrial outreach programmes to promote HPC** should be established, and existing connections (e.g. PRACE SHAPE) better exploited. The possibilities of securing funding from the ERDF, which aims to reduce disparities between EU regions, should be explored. Outreach events can be used to promote existing training, but also

⁵⁸ https://www.youtube.com/channel/UCfwgjtIQB3puojz_N9ly_Ag/feed

to better understand the needs of industrial users. There may be a need for new targeted material, or simply better targeted publicity of training. Training providers could be more proactive about offering training on demand for individual companies. While this would require additional staff effort, it would also generate revenue from course fees.

- **Gain leverage from good relationships between technology providers and HPC service providers** with the development of more collaborative courses, and increased participation of technology providers as lecturers on courses. These could be general courses on specific HPC tools, or in domain-specific areas such as CFD or structural mechanics.
- **Improve dissemination** to ensure training opportunities are publicised beyond the existing contacts, who are for the most part existing HPC users or people who have already attended courses. Develop a strategy for reaching new communities, particularly where these are composed of potential or beginner HPC users.

Challenge 3: Integrating HPC training into university undergraduate / postgraduate curricula

- **Undergraduate curriculum:** Courses on basic HPC concepts such as parallel computing should be integrated into the undergraduate curriculum for **all** CSE courses, and taught to all CSE postgraduates in their first semester. Elements of existing Master's courses could be repackaged to form a general course. While this would involve significant initial effort, the course could then be delivered to large numbers of students at the same time, and/or made available through online learning systems. An online learning approach could allow participation of students from other universities which lack HPC facilities and expertise. This would have a more general effect of increasing awareness of the benefits of HPC, and of the facilities and training that are available.
- **Short presentations** could be integrated into existing teaching – a one-hour overview of HPC could be given to all new undergraduate and postgraduate students, covering very briefly the basic concepts of HPC, examples of real-world applications, an overview of local/national facilities that are available to the students, and any training opportunities (courses, summer schools, etc.) that are relevant. This could be developed quickly and could reach a wide number of students.

Challenge 4: Encouraging interdisciplinary working

- **ITNs** are a potential source of collaborative partners, as they are typically involved in training PhD students across different domains. They offer huge potential for improving interdisciplinary HPC training and work by providing an opportunity for higher education institutions already involved in HPC training to share their knowledge and resources with researchers working in diverse, but relevant, scientific disciplines, giving them the basis to be able to conduct their research using HPC tools. Interdisciplinary workshops, summer schools, and other events focusing on applied uses of HPC can be organised, paving the way for development of formal study courses or programmes with an HPC orientation, as well as for new research opportunities that stem from use of supercomputing in essentially all scientific disciplines ranging from natural and life sciences to social sciences and humanities.

Challenge Addressed	Short-Term Actions	Medium / Long-term Actions
1. Rising demand	<ul style="list-style-type: none"> • Record more courses which are already being delivered. <ul style="list-style-type: none"> ○ Archive systematically in accessible place (e.g. a dedicated YouTube channel). ○ Publicise well. • Develop community of training providers to encourage collaboration between different institutes to increase efficiency. 	<ul style="list-style-type: none"> • Develop MOOCs to address specific target audiences. • Increase number of summer schools. • Develop “Train the Trainers” package of training material and associated certificated instructor course. • Identify ITNs working in appropriate areas and investigate ways to work with them.
2. Widening target audience	<ul style="list-style-type: none"> • Increase co-location of courses with other events. • Improve dissemination beyond network of established contacts. • Target users of specific community codes. <ul style="list-style-type: none"> ○ Targeted publicity. ○ New targeted course material. ○ Co-locate courses with relevant events. • More industrial outreach and better exploitation of existing programmes. 	<ul style="list-style-type: none"> • Increase technology providers’ involvement in training provision and course development.
3. Integrating HPC into undergraduate / postgraduate curriculum	<ul style="list-style-type: none"> • Build links with relevant CSE departments / faculties to discuss provision of HPC training in UG/PG curriculum. • Develop a one-hour presentation on basic HPC concepts, uses, and opportunities for new UG/PG students. 	<ul style="list-style-type: none"> • Develop UG/PG courses (from existing MSc material?) and liaise with universities to include a module (or equivalent) in UG/PG curriculum. • Investigate online training opportunities to open these courses beyond home institute.
4. Encouraging interdisciplinary working		<ul style="list-style-type: none"> • Identify ITNs working in appropriate areas and investigate ways to work with them.

Table 1 Challenges and Recommended Actions

8 Conclusion

This report has looked at the current provision of HPC training in Europe, the gaps that exist in that provision and perceived training needs of the community. As demand for HPC training increases to keep up with the growing pervasiveness of HPC and the heightened demand for HPC skills in the workforce, it examines how HPC training can be scaled up to meet this growing demand.

A survey of HPC stakeholders in Europe highlighted the great variety of HPC training topics, teaching delivery methods, course formats and target audiences that make up the HPC training landscape. As HPC training now needs to be targeted at all career stages from undergraduate level up, and increasingly needs to meet the needs of the non-academic sector as well as the academic sector, it is essential to embrace this variety in order to provide training in an accessible manner to all who need it. A “one-size fits all” approach is not appropriate, but rather HPC training must consist of a continuum of learning resources as varied as the community it addresses.

Currently, the academic HPC community is reasonably well served, especially when it comes to existing or likely future users of HPC services. In particular, PRACE is leading the way in this arena, notably through the PATCs, complemented by seasonal schools, the International Summer School and the Summer of HPC. As well as face-to-face courses, there is a growing focus on online training, with the first PRACE MOOCs due to start in March 2017.

The use of online training is growing generally, but still needs to be adopted more fully. While there are considerable overheads to setting up online courses, long-term benefits derive from the greater potential reach of each course. However, non-completion rates are relatively high, and the EuroLab-4-HPC curriculum provides guidelines on how to design courses to boost student engagement.

The number of universities offering HPC-focused postgraduate courses is growing, but there continues to be a shortfall in HPC training provision at the undergraduate level. HPC needs to be a core skill taught to all undergraduate students of CSE subjects, and HPC centres need to work more closely with university teaching units to increase integration of HPC into the curricula.

To encourage interdisciplinary learning and working, opportunities should be sought to collaborate in Innovative Training Networks. These networks are interdisciplinary, intersectoral and international and are designed for young researchers not only to develop their research skills, but also to acquire a range of transferable skills that can be applied in their future career, either in the public or private sector. They provide an opportunity for providers of HPC training to share their knowledge and resources with educational institutions and private companies, focusing on the applied use of HPC, in order to firmly establish this as a core skill in the future workforce.

Finally, it is clear that HPC training continues to evolve as rapidly as HPC itself. The subject matter naturally gets out of date quickly, meaning that the content of specific courses and of the curriculum as a whole need to be reviewed as often as every year. The methods of training provision are rapidly evolving, with a growing emphasis on online training. HPC training is an area which never stands still, and thus emerging needs must be addressed quickly, and the global training strategy must also be reviewed every 4-5 years to identify and tackle new trends.