H2020-FETHPC-2014

Coordination of the HPC strategy

EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D5.1
HPC Career Case Studies

Final

Version: 1.0
Author(s): Catherine Inglis, UEDIN
Date: 25.01.2016

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: < **D5.1**>	
	Deliverable Nature: <Report>	
	Dissemination Level: PU	**Contractual Date of Delivery:** 31 / Jan / 2016
		Actual Date of Delivery: 29 / January / 2016
	EC Project Officer: Panagiotis Tsarchopoulos	

* - 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 Career Case Studies**	
	ID: D5.1	
	Version: <1.0>	**Status:** ***Final***
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): D5.1.docx	
Authorship	**Written by:**	Catherine Inglis, UEDIN
	Contributors:	Alison Kennedy, UEDIN
	Reviewed by:	M. Bode, RWTH Aachen, I. Yenes, PMO
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	14/January/2016	Draft	Initial draft
0.2	15/January/2016	Draft	Revised after comments and corrections from Alison Kennedy
0.3	25/January/2016	Draft	Revised after internal reviewers' feedback
1.0	29/January/2016	Final version	

Document Keywords

Keywords:	PRACE, Research Infrastructure, HPC careers, case studies

Table of Contents

List of Acronyms and Abbreviations

HPC	High Performance Computing
IT	Information Technology
M	Month
MSc	Master of Science
R&D	Research and Development
R&I	Research and Innovation
WP	Work Package

Executive Summary

This deliverable comprises an initial set of HPC Career Case Studies, which are designed to promote HPC as an interesting and exciting career to the potential HPC specialists of the future.

The case studies are targeted at young people, and aim to promote the many opportunities open to those with HPC skills, both within the HPC community and also in associated professions where modelling skills coupled with domain expertise are of key importance.

A further set of HPC Career Case Studies will be produced as D5.5, to be delivered at Month 23 of the project.

As the initial set of case studies are the key part of this deliverable, this associated deliverable report is brief. The case studies are included as an annex.

1 Introduction

This deliverable comes under WP5, Talent Generation and Training for the Future.

This document has the following sections:

- **Section 1: Introduction** - an overview of WP5 and task T5.1.
- **Section 2: Production of the first set of HPC Career Case Studies** – overview, format of the case studies, selection of first candidates, list of case studies produced.
- **Section 3: Future plans** – planned use of the first set of case studies, and plans for producing further case studies
- **Section 4: Annex**, comprising the case studies themselves, which constitute the main part of the deliverable.

One of the acknowledged barriers to increasing the uptake of HPC in academia and industry is the shortage of skilled personnel. In addition to growing the overall numbers of trained persons, more effective cross-training is needed so that collaborators with differing skills and backgrounds can work together effectively.

WP5 aims to support young talent as a crucial element of the development of European HPC by promoting HPC and HPC skills to young people as exciting and rewarding areas for study and employment. It considers the whole HPC ecosystem, from technology suppliers to systems operators.

Task 5.1, Supporting Talent Generation, aims to examine and document ways to promote the importance of HPC and HPC skills to students in higher education. The task will look at attitudes towards HPC within the target age groups, and make available material that can be downloaded and presented at events such as local science festivals and school or university visits. The task will focus on the benefits of HPC in a gender-neutral way, to encourage more participation in HPC by girls. It will also highlight the many job opportunities open to those with HPC skills, both within the HPC community and in associated professions where modelling skills coupled with domain expertise are of key importance, such as in the field of engineering. In support of these aims, WP5 will produce a series of HPC career case studies targeted at young people to promote the different opportunities that a career in HPC can offer.

This document describes the production of the first set of HPC Career Case Studies, produced for Project Month 5. A further set of HPC Career Case Studies will be produced as D5.5, to be delivered at Month 23 of the project.

2 Production of the first set of HPC Career Case Studies

2.1 Overview of purpose and design of the case studies

The HPC Career Case Studies aim to increase the profile of HPC as an important skill and exciting career option.

The case studies are designed as single-page, double-sided flyers, which can be printed for distribution at relevant events, such as visits to schools and universities, science fairs and careers fairs, and which will be available for download from the EXDCI webpage.

2.2 Format of the case studies

The case studies are targeted at young people, so while they should appear professional and have a consistent design, it is important not to make them too formal or "dry" and technical, to make them appealing and increase the chance of someone stopping to read them. The target readers should feel that they can identify with the people featured in the case study.

For this reason, the following decisions were made regarding the format:

- The case studies have a "question and answer" format, with the answers written in the first person, to keep an informal, conversational feel to them. This allows for personal opinions to be given in a more natural way, and creates a sense of these being frank and honest testimonies.
- All case studies include a photo of the person featured, ideally an informal photo to increase the friendly, open feel of the testimony, and help the potential reader identify with the person.
- Candidates are asked to supply one or two other images, if they have anything relevant.

A word cloud is included on the front page, with the idea that one of the keywords might jump out at a potential reader and catch their attention more easily. This might help the reader to quickly decide which of the case studies is most relevant to them.

A general outline was given to the candidates, to help them to write their contribution, and to keep the content consistent across all the case studies and relevant to the aims of the task. This was as follows:

General introduction, including:

- Current job title and place of work;
- What your job involves in 1-2 sentences;
- A bit of background/biography, e.g. where you are from, briefly (1-2 sentences) what your career has been so far.

Current job:

- What is your current job, and what are your main responsibilities?
- What is the HPC element of the job, or how does HPC relate to what you do?

Career path to current job:

- What was the career path that led you to this point?
- What was your first experience of HPC? Did that change the course of your career?

Why work in HPC?

- What are the exciting aspects of working in an HPC-related career?
- Are there any challenging aspects?

Looking to the future:

- Where do you see your career leading you next?
- What is your outlook on the use of HPC in your field?

As everyone's case is different, not everyone would have something to contribute to each of these questions or sections. Therefore, the candidates were advised that they did not have to cover all of these points if they did not have something relevant to say.

2.3 Selection of potential candidates

The candidates to be featured need to be young enough to seem relevant to the target audience, which might include final-year secondary school pupils, undergraduate and postgraduate students, or researchers in academia or industry – anyone who might be considering which path to follow next in education or employment. At the same time, the candidates need to have enough experience to have an interesting career history to tell.

The featured candidates need to have HPC as a common strand in their stories, but to be as varied as possible in other respects. The following characteristics were considered in selecting potential candidates:

- Male / female;
- Country in which currently working;
- Nationality;
- Academic / industrial;
- Moved into a career in HPC from another background (especially if not informatics), or moved from HPC into a different area where HPC skills are valuable.

Ideal candidates would have had a varied career to date, perhaps have lived in more than one country during their career (showing that HPC is relevant across Europe and offers opportunities for mobility), and ideally – but not necessarily – would have moved into HPC from another discipline, or have secured an interesting job elsewhere thanks to their HPC skills.

The WP5 project partners are BSC, EPCC (UEDIN), University of Ljubljana and Seagate, which gives a mix of academic and industrial partners in different countries with many international contacts and therefore a large pool of potential candidates.

Due to the short timescale involved in creating the first batch of case studies, some of those who were contacted declined, although several of them expressed willingness to be involved in the next set – and indeed some of them are already preparing case studies for D5.5. In one case, the employer does not allow employees to take part in this sort of activity. These issues led to only 5 case studies being produced for D5.1, instead of the 6 case studies originally planned.

Some of those who were contacted but who did not take part in the first series of case studies include:

- Two people with chemistry backgrounds (including one female) who have now moved into HPC / Big Data, and have lived and worked in at least two countries;
- A computer scientist who moved from academia to an HPC centre and then on to a commercial software company, with work experience in two countries;

- An academic who came from a scientific discipline to study and then work in HPC, and then returned to academia;
- A computer scientist who worked in HPC for some years before moving first to a commercial software company and then to a bank;
- A computer scientist who studied and worked in HPC before moving into a varied and international career as a software developer and big data engineer.

The current case studies feature only candidates from the UK and France, although people in other countries were contacted. The next set of case studies will feature people working in a variety of other countries, and some of these are already in preparation.

As most of the current career case studies feature people working in mainstream HPC environments, the next batch will also include a greater variety of careers, including, for example, people working in scientific research who use HPC codes for field-specific research.

2.4 List of initial HPC Career Case Studies (forming D5.1)

Five case studies have been completed for D5.1:

Howard Price, Computer Games Developer (male, UK)

Howard was a physicist who took an MSc in HPC, which led him to a job with Sony Computer Entertainment Europe, developing PlayStation games.

Eilidh Troup, Applications Consultant at EPCC (female, UK)

Eilidh was a biologist who now works at EPCC, where she has been able to work on projects which combine her love of both biology and programming.

Henri Doreau, Research Engineer at CEA (male, France)

Henri studied electronics and computer engineering and first joined CEA as an intern. He now specialises in high performance storage solutions, working both in systems development and as a systems adminstrator.

Romain Dolbeau, HPC Expert at Bull (male, France)

Romain co-founded the CAPS compiler company during his PhD in computer architecture. The focus of this company evolved from general-purpose and embedded compilers to HPC, and this led to Romain developing the skills to become an HPC expert at Bull.

Elodie Ardoin, High Performance Networks Architect at CEA (female, France)

Elodie studied Electronic and Signal Processing Engineering and initially joined the seismic research division of CEA, spending 2 years at the Bolivian Seismic Laboratory, but has since specialised in computer networking in the HPC division of CEA.

3 Future plans

Following the internal project review, these case studies will be made publicly available on the EXDCI webpage. Printed copies will be made available at all relevant suitable events attended by EXDCI or its partners.

A second set of case studies is due in Project Month 23. These will be prepared over the course of the next 18 months, as suitable and willing candidates are identified. The existing case studies may also be refreshed as part of D5.5.

Feedback on the format and the content will be sought from both project partners and members of the target audience, in order to identify any improvements to be made and inform the production of the subsequent case studies.

The internal project reviewers have already made some useful suggestions for changes in both the format and the content of the case studies, and these will be taken into account before the next case studies are produced.

In order to allow the impact of the case studies to be monitored, a URL pointing directly to the case studies webpage has been given in the EXDCI information box on the bottom left of the front page. This page will be created before the case studies are made public. Using a utility such as Google Analytics (options still to be investigated), this would allow tracking of visits to the site that arise directly from people who have read the case studies, rather than people who have clicked through from the homepage. Downloads of the case studies could also be tracked.

While the case studies have been produced in English in order to reach the widest possible audience, the template along with the text and photos will be made available for project partners in the event that they want to translate the case studies into their own language for even wider impact.

4 Annex

The case studies are included here in the annex.

EXDCI Case Studies: HPC talent generation

Howard Price: Computer Games Developer

Howard Price is a game programmer who worked for Sony Computer Entertainment Europe for 8 years, developing titles for the PlayStation 3, PlayStation Vita and PlayStation 4 consoles. Before working with Sony, he gained an MSc in High Performance Computing at EPCC (Edinburgh Parallel Computing Centre) at the University of Edinburgh.

EXDCI is a European-funded project led by PRACE and ETP4HPC, the two most significant HPC bodies in Europe. It aims to co-ordinate the development and implementation of a common strategy for the European HPC ecosystem by supporting road-mapping, strategy-making and performance-monitoring activities.

This is one of a series of case studies designed to demonstrate the range of interesting careers in High Performance Computing (HPC). More case studies are available at http://www.exdci.eu/CaseStudies/

Howard, tell us a bit about yourself.

I'm originally from Northumberland but moved to Leeds to study physics, then, following an MSc in CFD (computational fluid dynamics), I worked for British Nuclear Fuels. After later completing an MSc in High Performance Computing and working briefly for EPCC, I was lured away by a job offer from Sony Studio Liverpool to work on WipEout for the PlayStation 3. Since leaving Sony, I have been developing games independently and have recently moved into contract programming.

How did your career path lead you to where you are now?

I was a keen programmer in my teens, but chose to study physics rather than computing to give myself a broader base. When I later joined BNFL on the Modelling & Simulation team, I soon realised that I enjoyed the computing elements of the job far more than the science and engineering aspects, so I started looking for a route into computing. The MSc in HPC then gave me the skills to move into the games industry.

What was your first experience of HPC? Did that change the course of your career?

I was really bitten by the HPC bug during the HPC MSc in Edinburgh. I particularly remember two pieces of coursework: one asked us to take a small scientific simulation and optimise to run efficiently on a single CPU core; another asked us to take some single-threaded image processing code and parallelise it to run on a multicore machine. It was thrilling applying the successive optimisation strategies and watching the code run faster and faster. After that I knew I wanted to work in high performance computing.

The EXDCI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 671558.

So what does your job involve?

Recently I have been specialising in graphics programming, which I love as it combines my favourite disciplines: maths, physics, performance programming – and of course computer games! It involves understanding how to program computer hardware to produce realistic images in real time.

On the last big game I worked on, the PlayStation 4 title DriveClub, I designed and implemented the particle system. Particle systems are used to create special visual effects, e.g. smoke, sparks, explosions, snow and rain, by simulating and drawing many individual small particles.

I worked with the vehicle team to attach particle emitters to the vehicle, to produce tyre smoke and kick-up from surfaces. I also wrote a bespoke precipitation system to simulate and render snow and rain, which featured motion blur and dynamic lighting - think of racing through a storm at night with the world lit only by the car headlights!

Particle effects in PlayStation4 game DRIVECLUB

How does HPC relate to what you do?

A computer game sends an image to the display 30 or 60 times per second. Game code needs to be highly efficient; if the total time taken to execute exceeds the frame period then the game will "drop frames", which can destroy the sense of immersion and ruin the experience.

Modern games consoles are highly parallel systems consisting of CPUs and the GPU connected by high-speed memory buses. Current-generation CPUs have around 8 cores, while GPUs have thousands of cores. GPUs were originally designed to accelerate graphics operations but have evolved into very powerful general-purpose processors. Modern game engines are designed to take full advantage of all of the processing power available, but this is only possible if the coders know how to write efficient, performant code that takes advantage of the machine's processor and memory architecture. These are skills I developed during my time in Edinburgh.

What are the exciting aspects of working in an HPC-related career?

The two best things are the people and the machines. At EPCC and Sony I was surrounded by a mix of young, extremely bright colleagues, and ridiculously experienced industry veterans. It's fantastic having such a wealth of enthusiasm and knowledge at your fingertips on a daily basis.

One benefit of working for a hardware platform holder like Sony is early access to new hardware and the opportunity to discuss with a network of industry-leading experts how best to use it. It's amazing when the game ships and you see your code being reviewed by journalists and players!

Are there any challenging aspects?

The codebase for a commercial game code can be absolutely enormous and highly complex, so it can take time to get up to speed when moving onto a project. Debugging parallel code and analysing profiles is tricky, but good tools, documents and colleagues are usually available. There's never the time in a project to implement every feature you'd like to, so you need to prioritise tasks judiciously. It can be heart-breaking to realise there isn't enough time to implement your favourite feature.

Where do you see your career leading you next?

I'm currently enjoying working on smaller projects again, with more control over the final product as a whole. However, I'd really like to get back into engine programming at some point. Graphics programming involves interacting with hardware at a very low level, which I find hugely satisfying.

What is your outlook on the use of HPC in your field?

As the trend in gaming devices continues towards more, rather than faster, cores, understanding parallel programming will be increasingly relevant. As graphics cards are fast becoming low-cost, high-performance general-purpose processors, there will be big performance wins from moving CPU-intensive code to GPU, where possible.

Much as music and video are moving to streamed services, in future we'll see more people playing games streamed over the internet, with the game stored and executed on remote machines, not on a box in the user's home. The computing power will migrate from the home to commercial sites.

EXDCI Case Studies: HPC talent generation

Eilidh Troup: Applications Consultant at EPCC

Eilidh Troup is an Applications Consultant at Edinburgh Parallel Computing Centre (EPCC). She has a background in genetics, but was not too keen on lab work, and wanted to do something more interesting than regular programming, so she moved into working in HPC. Outside work, she loves cycling, reading, and spending time with her young daughter.

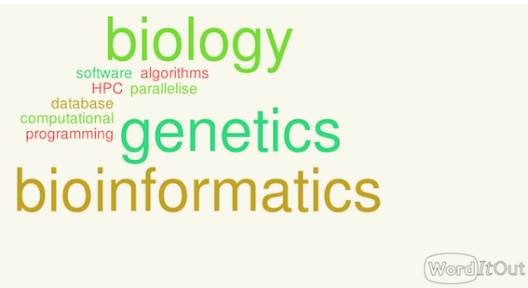

EXDCI is a European-funded project led by PRACE and ETP4HPC, the two most significant HPC bodies in Europe. It aims to co-ordinate the development and implementation of a common strategy for the European HPC ecosystem by supporting road-mapping, strategy-making and performance-monitoring activities.

This is one of a series of case studies designed to demonstrate the range of interesting careers in High Performance Computing (HPC). More case studies are available at http://www.exdci.eu/CaseStudies/

Eilidh, tell us a bit about yourself.

I work as an Applications Consultant at EPCC at the University of Edinburgh. My first degree was in genetics at the University of Glasgow, so I particularly enjoy the opportunities that I get in EPCC to work on projects with a biological focus.

How did your career path lead you to where you are now?

I loved genetics but didn't really enjoy lab work, so in my final year I did a computational project. Having enjoyed this greatly, I then did an MSc in IT, followed by a research MSc in bioinformatics. I worked on a computing project for the Institute of Neuroscience and Psychology for a while, then got a job with a pensions company. I started working in HPC when a friend recommended working at EPCC, and I applied for a chance to do something more interesting than routine programming.

What was your first experience of HPC? Did that change the course of your career?

I worked at EPCC for a few years before I asked to work on an HPC project I liked the sound of – SPRINT (http://r-sprint.org/), a parallel version of R. It was challenging to learn R, brush up on C and figure out how to use the supercomputer all at the same time, but the documentation was good and all of the more experienced HPC programmers were very helpful and supportive.

It was a really nice project because SPRINT makes it easy for people who just know R to be able to use HPC. SPRINT takes R functions and provides a parallel version which you can call in just the same way, without needing to get into tricky parallel stuff as a user.

The EXDCI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 671558.

So what does your job involve?

I'm currently working with the School of Biological Sciences at Edinburgh University, setting up a database for their experimental results. Having a biology background makes it easier to work with biologists and understand what they are doing.

I have also been working on a really interesting project with a company which breeds animals. We are looking at the genetics pedigree, and that code is not traditional HPC, but is based on a Hadoop cluster with a spark interface on top. We have some interesting ideas for how to put the data into the database with some data structures which are right for that specific problem.

What are the exciting aspects of working in an HPC-related career?

The best thing is when users are happy with what I have done; when it lets them do something that they couldn't before. That's a really nice feeling.

A couple of years ago I did a piece of work for the biologists, who were studying circadian rhythm – the day and night cycle of plants – and different mutations that can give them different clock lengths. There was a bit of Fortran code that would fit the periodic curve to the gene expression, but it was a bit buggy and some bits of data would just crash when it tried to run. I rewrote it in Java and they were really pleased because it gave the same results and it didn't crash with the new data set that they had. It's really nice when you get a complete piece of work that somebody can just use like that.

I have always really liked working on algorithms, and I like the interesting ideas of different data structures, so that made HPC appeal to me because you can do more unusual stuff with the data structures and algorithms.

Are there any challenging aspects?

I think the big challenge for me moving to HPC work was that I had not done much programming in C before. Some things were not exactly HPC issues, but were definitely hurdles to get past, and these can be hard for your confidence at the beginning. I had already picked up working from the command line and I was all ready to get stuck into the conceptual issues of HPC, such as message passing and optimisation, but it was really some of the very early things which were hard for me. For example, I had some problems with the compilers, as things are set up differently on the HPC systems.

Where do you see your career leading you next?

Coming from a biology background, I was drawn to this job as a chance to work with scientists, and I'm happy working on various biology-related computing projects here. The science is always interesting, and there are plenty of new coding challenges as the amount of data generated in biology is increasing so rapidly.

What is your outlook on the use of HPC in your field?

This is an exciting time for HPC and biology. Physicists have been using HPC for years, and now biologists are beginning to need that kind of computing power too.

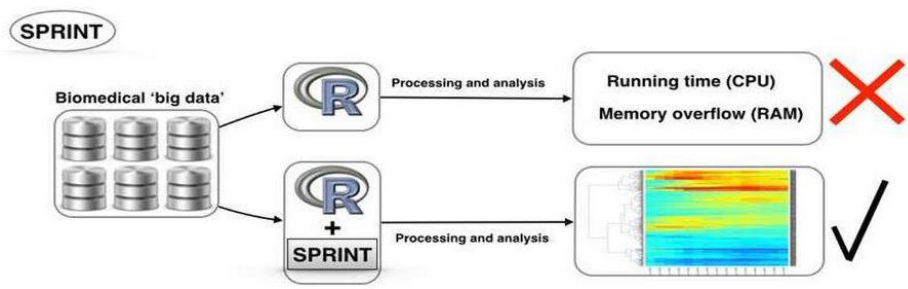

SPRINT helps biomedical researchers use HPC without needing to learn how to parallelise code

EXDCI Case Studies: HPC talent generation

Henri Doreau: CEA Research Engineer

Henri Doreau is as a Research Engineer at CEA, the French Alternative Energies and Atomic Energy Commission. As a software developer, he contributes to numerous projects in high-performance storage and big data domains. Away from the keyboard, Henri nurtures a passion for literature and likes oenology, cooking, running and cycling.

EXDCI is a European-funded project led by PRACE and ETP4HPC, the two most significant HPC bodies in Europe. It aims to co-ordinate the development and implementation of a common strategy for the European HPC ecosystem by supporting road-mapping, strategy-making and performance-monitoring activities.

This is one of a series of case studies designed to demonstrate the range of interesting careers in High Performance Computing (HPC). More case studies are available at http://www.exdci.eu/CaseStudies/

Henri, tell us a bit about yourself and how your career path has led you to where you are now.

I grew up in Angers, in the Loire valley (France), where I also studied electronics and computer sciences at ESEO engineering school.

I have a strong interest in software development, and especially free software. In the past, I used to focus mostly on IT security, and discovered HPC through an internship at CEA[1]. I applied out of curiosity, enjoyed it, and I am still in the same team today! Back then, my goal was to implement scalable command execution techniques within the ClusterShell library[2]. It matched perfectly my inclination to push IT to the limits and design algorithms that "do more".

Free software communities (and the nmap security scanner project in particular) had taught me to adopt a rigorous approach to try to produce elegant and robust code. I realised that this was even more relevant within CEA due to the nature of the concerns behind HPC.

I enjoyed the culture of excellence, learning things every day, working on challenging problems, and the dual research and engineering aspect of the work. Free software is really important to me, and the fact that CEA values it too made it a perfect place for me to work.

[1] http://www.cea.fr/english-portal
[2] https://github.com/cea-hpc/clustershell

The EXDCI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 671558.

So what does your job involve?

Today, my work is divided into two parts.

Most of my time is spent on system development of high performance storage solutions. This includes contributions to the POSIX filesystem world, but also more innovative techniques combining various approaches such as object storage or machine learning. I started working on the Lustre filesystem to fix a couple of small things we observed in production, and then I increased the scope of my contributions to bigger and more complex patches.

I am also a system administrator on the storage infrastructure of our Tera (defence) and TGCC (academic and industrial research) computing centres. My role is to ensure that the storage system runs smoothly and with the expected performance.

These two activities are mutually self-exclusive, so I cannot work on them both at the same time, but production experience is incredibly valuable when developing. Having a deep understanding of the critical tools (namely the Lustre filesystem and RobinHood policy engine) makes me a much more efficient administrator. This is usually referred to as "DevOps" in the computer literature.

What are the exciting aspects of working in an HPC-related career?

My activities allow me to meet with many people from very different technical areas. First, there are the users of the supercomputers, who include physicists, mechanical engineers and biologists, among others. HPC has become a key tool for many scientists working on major questions of our time. Activities range from climate modelling to personalised medicine or high energy physics. It is important to understand what the users do and how their code operates in order to provide them with sensible advice for using the existing systems, and to design even more efficient ones.

I also travel to meet with various stakeholders in the storage community, including vendors, other HPC sites, and researchers. This really is a job where you learn things every day. It shapes your way of thinking. Being exposed to such a variety of approaches, smart people, innovative ideas and technologies forces you into considering many more options to solve what you are working on. You must simultaneously stay receptive to new techniques and yet know precisely what you want to achieve, in order to filter out what may currently be trendy, but is not appropriate in your case.

What is your outlook on the use of HPC in your field?

As a team and an organisation we put significant effort into making our system more efficient, smarter and more robust. We are working hard on improving interfaces between layers of our storage stack, the mid-term goal being to prepare it for new models. We are currently integrating machine-learning techniques into our centralised logging infrastructure to re-structure information and assist troubleshooting. I can see that we are making progress despite being completely new to this particular field, and this is very exciting!

Exascale is a major goal, and there are still many open questions before we can achieve it. We are investing significant amounts of time in order to design secure and reliable systems that will scale to this extent. I am fully committed to these collective and ambitious targets.

Coming up this year is LAD'16, an international workshop on Lustre, which we are organising in Paris. I expect that this event will be another opportunity to strengthen the Lustre community, share ideas and start projects between teams.

Where do you see your career leading you next?

I have a very strong interest in solving technical problems. I would like to continue accumulating experience and develop my expertise in order to be able to tackle larger and more complex challenges.

Organisational questions behind HPC are also especially interesting. This field has only recently been identified as a strategic one here in Europe. From my position, I have the feeling that parallel storage is not researched as much as its critical role requires. I have also noticed that many groups are not aware of others' work and so miss opportunities for highly fruitful collaborations. HPC plays a pivotal role between research and industry, which I believe can be strengthened. I hope that I can contribute to that at some point.

CEA's Tera 100 supercomputer

EXDCI Case Studies: HPC talent generation

Romain Dolbeau: HPC Expert at Bull

Romain Dolbeau works as an HPC Expert at the Center of Excellence in Parallel Computing of Bull, an Atos company. In this role, he works with pre-sales and support teams to help Atos customers fully leverage the computing power and energy savings brought by many-core technologies.

EXDCI is a European-funded project led by PRACE and ETP4HPC, the two most significant HPC bodies in Europe. It aims to co-ordinate the development and implementation of a common strategy for the European HPC ecosystem by supporting road-mapping, strategy-making and performance-monitoring activities.

This is one of a series of case studies designed to demonstrate the range of interesting careers in High Performance Computing (HPC). More case studies are available at http://www.exdci.eu/CaseStudies/

Romain, tell us a bit about yourself.

I grew up in a Paris suburb where I discovered computing at a young age from a primary school substitute teacher. I later studied computer science at the Université Paris Sud, before moving to Rennes to work on a PhD in computer architecture under Prof. André Seznec at the ENS Cachan (now the ENS de Bretagne).

How did your career path lead you to where you are now? How did you move into HPC?

It was during my PhD work that I co-founded the compiler company CAPS entreprise with other students, research engineers and Prof. François Bodin. I joined CAPS as an employee, both as a system administrator and also as the computer architecture guy among the many compiler experts.

It was during this decade at CAPS that I really discovered HPC. The focus of the company quickly moved from general-purpose and embedded compilers to HPC, in particular when we introduced the first directives-based programming model for hybrid computing, HMPP, in 2007.

Most of the early adopters of the technologies (hybrid computing and HMPP) were large-scale HPC users, and we had to understand their needs to best serve them and improve our brand-new technology.

Ultimately, I became a pure HPC expert, working with our customers and partners to get the best out of their hardware and software investment.

The EXDCI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 671558.

So what does your job involve?

I joined Bull, now an Atos company, as an HPC Expert after the end of CAPS. As such, I again work with our customers, partners and other teams at Atos to fully leverage their hardware for HPC. There are different aspects to this, depending on the specific project.

Some clients are in need of training. They are experts in their own fields – which might be quantum physics, hydrodynamics, or climate prediction – but not necessarily in computer science. Therefore I need to train them to a varying degree of expertise in the fields that are relevant to HPC: hardware, compilers, parallelism, etc. Some might only need an entry-level course to avoid beginner's mistakes, while others may apply for multiple days of intensive training in vectorisation, cache exploitation, and so on. The goal is for them to be able to create their new numerical codes in a way better suited to the underlying hardware, and to be able to do their science faster.

Other clients are in need of expertise or consulting work. The code is there, they are knowledgeable in the required subjects, but they want even more from their hardware. I work with them to study their codes, and adapt them to the current – or future – hardware. Sometimes minimal changes are required, sometimes entire parts of the code need to be updated to fit the machine. This requires up-to-date knowledge of the hardware and the software tools and can be quite challenging since some of the customers are already experts themselves.

What are the exciting aspects of working in an HPC-related career?

I need to constantly study the state of the art in my field to keep up with innovations and methodologies, so that I can have up-to-date answers for our customers. And since those customers come from every background, there is always something new to learn from them and new challenges born from different requirements.

Where do you see your career leading you next, and what is your outlook on the use of HPC in your field?

HPC has only become bigger and more relevant since I started in the field. It is not only numerical simulations which have become ubiquitous in every science and every industrial process, but also all the data-mining algorithms, deep-learning algorithms and so on from the social networking and internet-of-things era. All of them require massive amounts of computation.

So at least in the near future, I see my job as being broadly similar to what it is now. The technologies are constantly evolving, so that yesterday's advice might be tomorrow's mistake – or vice versa. Hybrid computing is common in the high-end now but was non-existent less than a decade ago. GPUs are ubiquitous, but alternatives, such as the Xeon Phi or PEZY-SC, are challenging them. The debate between proponents of many small cores and supporters of a few large cores is as intense at it ever was. Non-volatile memory is rising fast and threatening assumptions common in many user codes and perhaps even operating systems. Optical communication is shaking things up by moving from fibre optics to silicon photonics. The leadership of Intel in conventional CPUs is under threat by the manufacturers of mobile devices.

So even with the same job description, the ever-changing landscape of HPC is what keeps it interesting for me.

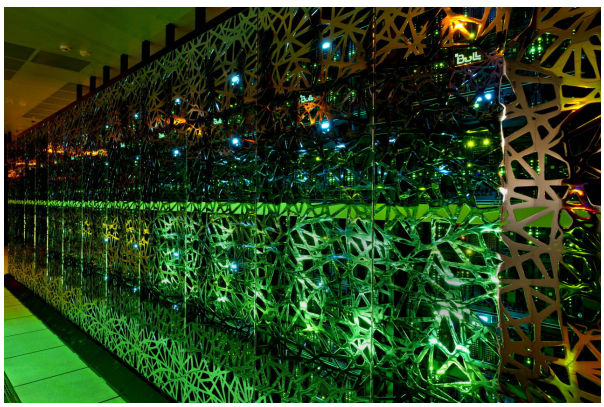

CURIE, a Bull supercomputer at the Très Grand Centre de Calcul (TGCC), near Paris. Copyright: Cadam/CEA

EXDCI Case Studies: HPC talent generation

Elodie Ardoin: High Performance Networks Architect

Elodie Ardoin joined the seismic research division of CEA, the French Alternative Energies and Atomic Energy Commission, in 2002. In 2009, she moved to the HPC division to work on high-performance networks, participating in storage network design and operations for the CRRT and TGCC supercomputing infrastructures. She has two children aged 5 and 9.

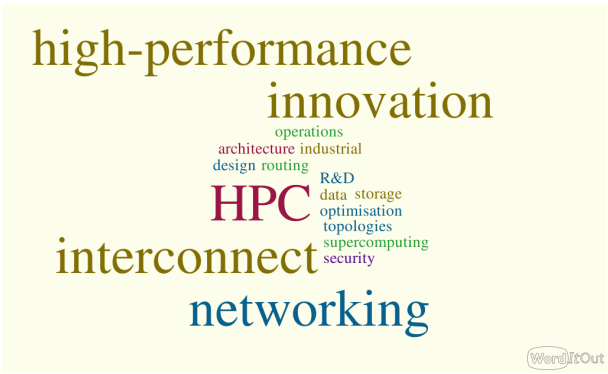

EXDCI is a European-funded project led by PRACE and ETP4HPC, the two most significant HPC bodies in Europe. It aims to co-ordinate the development and implementation of a common strategy for the European HPC ecosystem by supporting road-mapping, strategy-making and performance-monitoring activities.

This is one of a series of case studies designed to demonstrate the range of interesting careers in High Performance Computing (HPC). More case studies are available at http://www.exdci.eu/CaseStudies/

Elodie, tell us a bit about yourself.

I graduated as an Electronic and Signal Processing Engineer from the ENSEEIHT engineering school in 2001, and joined CEA[1], the French Alternative Energies and Atomic Energy Commission, in 2002. I worked in the seismic research division as an electronic and signal processing engineer for 7 years before moving over to the HPC division.

How did your career path lead you to where you are now? How did you move into HPC?

My career at CEA began with a 2-year period abroad in a Bolivian seismic laboratory, under a partnership convention, where I performed on-site troubleshooting on the seismic and infrasonic monitoring networks of the CTBTO, the Comprehensive Test-Ban Treaty Organization. On my return to France, I continued to work on transmission, networking and real-time data acquisition systems. I then worked with the R&D team on new systems validation and international on-site deployments for 5 more years.

I got the opportunity to specialise in computer networking in the HPC division at CEA in 2009, working on the design of network architectures for projects of all kinds (security, routing and data centres).

In 2013, I had a break during which I took postgraduate courses in marketing at the ESAN business school in Peru, returning to CEA in 2014.

[1] http://www.cea.fr/english-portal

The EXDCI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 671558.

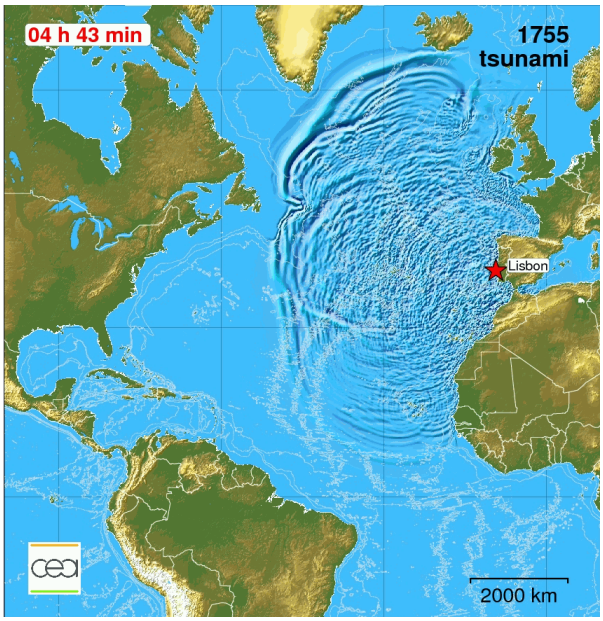

Seismic research: computer simulation of tsunami

So what does your job involve? How does HPC relate to what you do?

I work in the division in charge of large systems architecture design and operation. I participate in R&D activities led by our industrial partner on new interconnect technologies for supercomputing clusters. This includes investigating interconnect topologies, studying the impact of routing and placement, and identifying traffic patterns to optimise communications and to get the best overall performance. I am also in charge of the operational follow-up and design of the Très Grand Centre de Calcul (TGCC) high performance networks.

The ever-growing size of HPC clusters makes interconnect a crucial subsystem in terms of costs and performance. Resource optimisation, security, and message-passing rate performance are key aspects that need to be addressed. My operational activities allow me to be aware of the needs and limitations of the current technologies and help me to develop my expertise.

What are the exciting aspects of working in an HPC-related career?

Innovation: HPC and simulation play a significant role in innovation for society (eg research in genetics, geophysics, health and energy, related big data processing, etc). The same is true of systems architecture. It is innovation at a high pace, and a great place to learn.

Are there any challenging aspects?

Without an extensive background in computing science, getting good technical expertise on quite specific high performance networks is challenging for me, but so enriching: understanding the ins and outs, and getting the global view in order to understand technical trade-offs and put them in place. Connecting and understanding all actors with different backgrounds and needs (eg computer architects, storage architects, users) is not always easy – but that's another area where networking makes sense too!

Where do you see your career leading you next?

More of the same…I still have so much to learn!

The complex world of interconnect topology

Elodie with members of the Women in HPC network (www.womeninhpc.org)