

EXDCI Case Studies: HPC talent generation

Mehdi Amini: Compiler Engineer

Mehdi Amini is a tech enthusiast who works for Apple, and is based in the San Francisco Bay Area. He has a particular interest in parallel and heterogeneous programming, including GPUs, and holds two Master's degrees and a PhD. A cross-disciplinary Master's project, working with an astrophysicist, first sparked his interest in HPC.



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

Originally from Strasbourg, I am now enjoying the California weather in the San Francisco Bay Area, where I work as a compiler engineer at Apple.

My first contact with HPC was during an introductory Master's class on parallel computing. My Master's thesis project was a collaboration between the HPC centre and an astrophysicist who was interested in doing galactic simulations using GPUs. This was right when CUDA was first released and I found it mind-blowing at that time. The cross-disciplinary collaboration with brilliant researchers was also a very interesting part of the job. My interest has been in HPC since that time.

After completing two Master's Degrees in Strasbourg, I was employed at the University as an engineer for three years, working on the HPC resources amongst other things.

I left the university to join SILKAN, a startup specialising in high-performance simulations, which has now built up a particularly strong presence in the aerospace and defence, transport, and energy sectors. There I was working in collaboration with MINES ParisTech on a research project which involved compilers technology and heterogeneous computing with GPUs.

After completing my PhD, I relocated to Silicon Valley to help with the launch of the subsidiary which the company was just opening there. Finally, my last move led me to join the compiler group at Apple, where I'm currently working on the open-source Clang/LLVM project (<https://clang.llvm.org/>). My focus is on trying to improve the optimisations in the project.

Innovation
Quality Performance
Compiler
Silicon Valley
HPC
Efficiency

WordItOut

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/>

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

As a compiler engineer, multiple aspects are important: supporting new hardware, helping to develop new languages or language features, and improving performance.

My main focus is currently on performance: optimisations in the compiler provide a double-digit percent improvement in performance.

This translates directly into cost-efficiency in devices of all sizes, from improved efficiency in the utilisation of supercomputers, all the way down to a slower drain on the battery in your smartphone.

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

The exciting aspect of working in a performance-constrained environment is how challenging it can be to get the next few percents of improved efficiency out of the available components. We constantly have to be innovative, and usually go to great lengths to reach the next step.

Another aspect of being a compiler engineer is the need both to understand the software, as written by the programmer, and have a strong knowledge of the details of the hardware. We often gain significant improvements when collaborating with the hardware designers directly. It is definitely an advantage to work in a company – such as Apple – which controls the end-to-end project.

Are there any challenging aspects?

The mobile market today is also very demanding, presenting its own challenges. The power budget is limited because the battery is a scarce resource. From a compiler point of view, improving the performance of the software usually translates directly into power saving.

Another important aspect is the binary size: building and shipping a full OS on a watch requires great care in terms of the size of the program we produce, and we sometimes have to go the extra mile to get just a few percent reduction.

Where do you see your career leading you next?

Compilers are very large beasts (LLVM is approaching 1.5M of lines of C++, and Clang is over 1M). I am sure that I have plenty of time ahead of me working on various related projects.

The compiler is also a critical component of the developer toolkit, and it will be for some time.

Working on the compiler is a place where I feel I have opportunities to be innovative, while being surrounded by brilliant engineers from whom I am learning every day.

Finally working on an open-source project is definitely a plus, and the LLVM project has a focus on quality and moves at a fast pace that matches what I like.

The prominence of heterogeneous hardware everywhere, from the smartphone to the supercomputer, and the respective challenges they offer, make it a common goal to extract even more performance out of the software.

HPC used to be a niche for supercomputers, but that has been over for some time now: HPC is everywhere.



Golden Gate Bridge, San Francisco (USA) ©iStock.com/Stella Levi