Best Outstanding Output

Name: Dr. Juan Fumero
Faculty: School of Engineering
Position: Research Fellow
Research area: Parallel Programming, GPU Programming, Compilers, and Virtual Machines.

I am currently a Research Fellow working as part of the Advanced Processor Technology (APT) Research Group. I joined the University of Manchester in October 2017, as a postdoc, after finishing my PhD. Since then, I have been working with the Beehive-lab (part of the APT group) on managed runtime systems, GPU/FPGA computing, and runtime compilation.

I obtained my bachelor's and master's degrees at ULL in Spain. Then, I moved to Switzerland to work as an intern at CERN, in which I collaborated with Intel to study the portability of the existing physics simulators developed at CERN on modern Intel architectures using auto-vectorizing compilers (e.g., CILK+). After that, I decided to do a PhD at the University of Edinburgh, to research on accelerating interpreted programming languages on GPUs through runtime compilation and optimisations.

During my PhD, I also did an internship with Oracle Labs to work on a distributed programming framework for the R programming language and Big Data platforms using GraalVM. This internship also opened the door for more ideas to combine with my PhD, such as the use of partial evaluation and compiler specialisations to transparently accelerate existing R programs on GPUs.

In your own words, please describe your outstanding research whether that be an output, impact, contribution to the environment:

At The University of Manchester, we have been developing a parallel programming framework to facilitate heterogeneous programming for new computer architectures such as GPUs and FPGAs. This framework (called TornadoVM) not only facilitates programming but also complements existing programming languages, such as Java, for transparently using new hardware.

This project has been well received by both, academia, and industry. From the academic perspective, we have been able to publish numerous publications with multiple master's and PhD students. From the industry side, this project has got the attention of companies such as Intel. All of this has led to open new collaborations with the industry, and now we are working together with Intel to improve the technology built at the University of Manchester for the new upcoming Intel architectures through their new programming models.

Furthermore, through this project, we have been invited to numerous industry-focused conferences such as Devoxx, JVMLS, QCon, Joker, and JavaZone as well as Java groups such as the New York Java SIG group to talk about TornadoVM.

Besides, I have been invited to participate in the Level Zero Technical Advisory Board to help shape the new programming models that can be used by high-level language virtual machines (e.g., Java).

As an outcome of all this commitment, I have been accepted into the Intel Software Innovator Program, which is a scheme designed to early evaluate and probe new hardware and software.

What motivated you to do this?

Since very early during my bachelor's studies, I have been interested in compiler implementation, parallel programming, and GPU computing. Those courses fascinated me. This led me to work as a research assistant during my master's degree to help the development of a new compiler for OpenACC (a standard for GPU programming for C/C++ and Fortran programming languages).

Through this project, I really got into the question of how we can make existing programming languages run faster just by using modern hardware as easy as possible. It turns out, that this is a very complicated question to answer, and I have been working on this topic since then.

Apart from having the motivation to solve complex problems, I also want to acknowledge the team I work with, the TornadoVM team. From my point of view, it is very important to be comfortable expressing ideas and debating new directions. I must say that, through many internal discussions, we created new sub-projects that, in many cases, led to new collaborations and publications.

What are you planning to do next?

There are always new things to learn and to do. One of the things we are planning to do is to look for opportunities for commercialization. We believe the technology we are creating at The University of Manchester has value for developers and software engineers, mainly because it offers an alternative to developers that want to use new hardware without having to know all low-level details to make applications work on accelerators. I believe this will be more relevant in the upcoming years, in which many companies are building their hardware accelerators for increasing the performance of custom workloads, such as Deep Learning, graph analytics, etc.