Conference Report for the Old Centralians' Trust

SIAM Conference on Parallel Processing for Scientific Computing
Seattle, 11th-15th February 2020

Jan Eichstädt
jan.eichstaedt13@imperial.ac.uk
Imperial College London, Department of Aeronautics
February 21, 2020

This report gives a brief summary of my participation at the *SIAM Conference on Parallel Processing for Scientific Computing* (PP20) in Seattle, Washington. I am very thankful to the *Old Centralians' Trust* for awarding me with a travel grant of 530GBP to support my participation.

This SIAM conference, hosted by the *Society for Industrial and Applied Mathematics*, is one of around a dozen annual SIAM conferences in the United States and worldwide, that each cover a different area in applied mathematics. The SIAM PP20 conference focuses on current and future trends in high-performance computing (HPC), where a recurrent theme is the parallel processing of data. The research in this area lies at the interface of scientific applications, mathematical algorithms, high-performance computing hardware, and their programming paradigms.

Currently, I am in the fourth year of my PhD studies at the *Department of Aeronautics* at *Imperial College London*. My research is based around the question how we can future-proof computational fluid dynamics (CFD) solvers for ever more parallel computing hardware. My work is supervised by Professor Joaquim Peiró and embedded in the *Nektar++* group at the department, which jointly advances CFD algorithms and solvers for scientific and industrial applications.

At the SIAM PP20 conference, I was delivering a presentation with the title *Efficient vectorised CUDA kernels for high-order finite element flow solvers* in the session *Numerical Methods for Flow Simulations*. The talk was based on my work which optimises the performance of element-based compute kernels for flow solvers using the *CUDA* language for general purpose graphics processing

units (GPGPUs). This computing hardware has different characteristics in terms of parallel processing and memory access than classical CPUs, so we had to adapt our legacy CPU implementations accordingly. The algorithms that we considered cover all operators required to solve the Helmholtz equation, which lies at the heart of our more elaborate flow solvers, and is therefore a concise but still representative testcase. We placed special emphasis on the layout of data in memory and compared the performance of a few different options, and found that the optimum depends on the specific hardware at hand and the type of the elements.

My presentation was well received, a few participants were very interested to learn about the details of the performance optimisation that I applied in my work and the methods of performance benchmarking that I employed to quantify these.

Personally, I have benefited from participating in this conference in multiple ways. First of all, the dissemination of my research; through the direct feedback of participants I feel encouraged that my research goes in the right direction, is relevant, and transferable to related research questions with different applications.

Additionally, I took the chance to listen to talks that are closely related to my own research and potentially useful for my career after finishing the PhD this year. I attended very interesting talks on performance benchmarking for HPC, that is becoming increasingly complex as the computing hardware becomes more heterogeneous. In the next years and decades HPC will not only see CPU and GPGPU hardware, but also specialised chips for machine learning algorithms, and probably quantum computers and so called neuromorphic computers, that can dynamically change their circuits to adapt to different algorithms. There have been invited presentations from directors of US national labs, laying out their strategy and latest results towards computations at Exascale, which will enable to conduct new scientific simulations, for example in atmospheric modeling, molecular design for new materials, interactions of wind turbines, or design of more efficient reactors. After this conference, I feel much more informed about these trends and could very well imagine to work with some of these fascinating technologies in the next years.

This conference also allowed me to travel to the United States again and strengthen my ties with friends and colleagues that I worked with over the last summer at MIT in Boston, as well as getting to know Seattle as a city, too, which is home to research centres of companies such as Boeing and Amazon.

Again, I like to thank the *Old Centralians' Trust* and their donors for supporting my attendance at the SIAM Parallel Processing conference 2020, which has been very beneficial for my research, and my development as a researcher.

Jan Eichstädt