A Reconfigurable Vector Instruction Processor for Accelerating a Convection Parametrization Model on FPGAs

Syed Waqar Nabi (Univ. Glasgow), Saji Hameed (Univ. Aizu) and Wim Vanderbauwhede (Univ. Glasgow)

High Performance Computing (HPC) platforms allow scientists to model computationally intensive algorithms. HPC clusters increasingly use General-Purpose Graphics Processing Units (GPGPUs) as accelerators; FPGAs provide an attractive alternative to GPGPUs for use as co-processors, but they are still far from being mainstream due to a number of challenges faced when using FPGA-based platforms. Our research aims to make FPGA-based high performance computing more accessible to the scientific community. In this work we present the results of investigating the acceleration of a particular atmospheric model, Flexpart, on FPGAs. We focus on accelerating the most computationally intensive kernel from this model. The key contribution of our work is the architectural exploration we undertook to arrive at a solution that best exploits the parallelism available in the legacy code, and is also convenient to program, so that eventually the compilation of high-level legacy code to our architecture can be fully automated. We present the three different types of architecture, comparing their resource utilization and performance, and propose that an architecture where there are a number of computational cores, each built along the lines of a vector instruction processor, works best in this particular scenario, and is a promising candidate for a generic FPGA-based platform for scientific computation.

We also present the results of experiments done with various configuration parameters of the proposed architecture, to show its utility in adapting to a range of scientific applications.