Porting HPCG to Heterogeneous HPC systems

This paper presents certain lessons learnt while porting a widely used HPC, data intensive benchmark to Heterogeneous HPC systems incorporating state-of-the-art FPGAs . Using targeted benchmarks for HPC systems is the most widely accepted path to measure and rank the performance of large-scale machines. The dominating metric for the last decade in the HPC benchmarking field has been the High Performance Linpack (HPL) benchmark. HPL provides computations on dense matrices, and results in a favourable computation-to-memory-access ratio of O(N), where N is the problem size. This results in an optimistic performance prediction that is dominated by floating-point arithmetic performance. More recently, another metric has been proposed, namely the High Performance Conjugate Gradients (HPCG) benchmark, which has gained high acceptance, and a corresponding listing alongside the TOP500 ranking. HPCG is at the opposite extreme, compared with HPL, in terms of its computation-to-memory-access ratio, which is O(1). HPCG is therefore strongly influenced by memory bandwidth, i.e., being internally I/O bound. This important differentiation signifies a highly degraded performance target, which can become even more unfavourable for FPGA-based HPC systems. In effect, reconfigurable logic can benefit when applications adhere to the opposite scenario, i.e., high computation-to-data-transfer ratios. As the latter is the case for many real applications, they can take advantage of an FPGA's inherent effectiveness in implementing many parallel data-flow processing engines. However, even under the stringent requirements shaped by an HPCG-like execution scenario, making the most of FPGAs is a challenging and significant use case for current and future heterogeneous HPC systems.

We first present an analysis of the HPCG-related mathematics and the corresponding structures, designating a score of computational kernels. Then we move on to HPCG implementations, focusing on FPGA platforms and describing some of our earlier as well as our latest efforts, along with alternative implementations and their related concerns and limitations. Then we report on our advances targeting a heterogeneous HPC system incorporating FPGAs. This process involves many interesting steps, which we elaborate on, along with our latest efforts and results. Importantly, and in parallel with these efforts, we are investigating variable- and mixed-precision techniques, in order to alleviate the memory bottleneck stemming from sparse matrix computations. We present results on various techniques which seem rather promising, as measured through extensive simulations. Furthermore, we elaborate on an alternative HPCG data layout scheme, which stands somewhere in the middle, regarding complexity, as well as performance. We finally present a set of measurements from a variety of platforms, including CPUs, GPUs and FPGAs, for various configurations and data sizes, in order to have a clearer view of the possible benefits that FPGA-based platforms are able to offer.