



Contribution ID: 52

Type: **not specified**

Massive Parallelism for Combinatorial Problems

Massive parallelism is a design paradigm in computer architecture which trades the complexity of sequential processing units for many simpler units operating in parallel [1]. Massively parallel architectures have a higher theoretical throughput than a similar sequential architecture, because many of the transistors which would otherwise be committed to inefficient optimisations are instead available for additional cores; thus, increasing the overall throughput of the device [1].

Many contemporary research projects have investigated the use of massively parallel computer devices to accelerate computation [2, 3, 4]. One such device is the Graphics Processing Unit (GPU.) The GPU has become ubiquitous with compute driven science due to its wide support base for HPC and its low cost. GPU driven HPC applications are found in a number of data driven research fields such as bioinformatics [4].

Our work leverages massively parallel platforms to accelerate combinatorial problems. A central idea of the research is a model. The model is used to describe the structure of algorithms which solve combinatorial problems on massively parallel platforms.

The model emphasises two central ideas:

- Data parallelism
- Collaboration between solvers

These two mechanisms ought to lead to a faster execution of each step of the algorithm and an increase in the rate of convergence of the system as a whole.

We present an overview of the model as well as a review of two contemporary massively parallel platforms – the GPU and the Convey FPGA Hybrid Computer – investigating their architectures and providing a brief review of contemporary research being aided by these platforms.

[1] Asanovic, K. et al. “A View of the Parallel Computing Landscape,” 2009, Communications of the ACM 52, pp 56–67.

[2] Ujaldon, M. “High performance computing and simulations on the GPU using CUDA,” 2012, HPCS, pp 1–7.

[3] Yang, B. et al. “GPU accelerated Monte Carlo simulation of deep penetration neutron transport,” 2012, PDGC, pp 899–904.

[4] Z. Ying et al. “GPU-Accelerated DNA Distance Matrix Computation,” 2011, Chinagrid Conference, pp 47–47.

Primary authors: Mr STEERE, Edward (University of the Witwatersrand); Prof. HAZELHURST, Scott (University of the Witwatersrand)

Presenter: Mr STEERE, Edward (University of the Witwatersrand)