

Improvements to the design and description of the Reduceron

Matthew Naylor
mfn@cs.york.ac.uk

Colin Runciman
colin@cs.york.ac.uk

We have developed a reduction machine for lazy functional languages called the Reduceron (paper at IFL 2007). The Reduceron exploits wide, parallel memories to increase evaluation speed, and has been prototyped on an FPGA. We have encouraging results, but the Reduceron is still some five-times slower than a state-of-the-art lazy functional language implementation running on a PC.

A work-in-progress design for a new version of the Reduceron makes improvements at several levels, from the compiler to the machine architecture to the actual FPGA device being used. Improvements include the use of transformations in the compiler to make better utilisation of wide memories, and alternative arrangements of the memories in the machine to improve access times. By the time of the HFL workshop, we expect to have new results to report at one of more of the above levels.

The Reduceron is itself described in the lazy functional language Haskell, using the Lava library. Lava has served us well, but we have encountered some limitations, in particular the lack of support for behavioural description, the lack of static typing on bit-vectors, and the difficulty in interpreting the netlist's critical path in the context of the original Lava description. By the time of the HFL workshop, we expect to have completed our own variant of Lava which addresses these problems.