Geophysical Research Abstracts,
Vol. 11, EGU2009-1464, 2009
EGU General Assembly 2009

Parallelisation of storage cell flood models using OpenMP and accelerator cards

J. Neal, T. Fewtrell, M. Trigg, and P. Bates
School of Geographical Sciences, University of Bristol, University Road, Bristol. BS8 1SS. United Kingdom

Recent developments in computer processors have moved away from increasing clock speed towards multi-core approaches. For computationally intensive flood inundation models this development shift will need to be exploited if simulation runtimes are to be reduced in the near future. This work describes the implementation and benchmarking of a parallel version of the LISFLOOD-FP coupled 1D-2D hydraulic model. The motivation behind the study was that reducing model run time through parallelisation would increase the utility of such models by expanding the domains or resolutions over which they can be practically implemented, allowing previously inaccessible scientific questions to be addressed. Of the many parallelisation methods only two were considered here: The first used the shared memory Open Multi Processor (OpenMP) application programming interface (API) to achieved parallelisation on standard CPU's. The second method used Clearspeed accelerator boards to run the computationally intensive 2D floodplain component of the model, whilst the computationally less intensive 1D channel model runs at a longer time step on a standard CPU.

Parallel speedup with OpenMP was calculated for 13 models distributed over seven study sites and implemented on one, two, four and eight processor cores. Selected cases were then run on one core with a Clearspeed CSX600 accelerator board. A key advantage of OpenMP, with an explicit rather than implicit model, was the ease of implementation and minimal code changes required to run simulations in parallel. Using the Clearspeed accelerator boards required selected functions to be re-written in the Cn parallel programming language (an extension of the C language) and an interface based on Clearspeeds CSPX Accelerator Interface Library for calling these remote functions from the host code where the rest of the model was running. Preliminary results indicate that both approaches can be used to parallelise storage cell flood inundation models and that the explicit scheme used in LISFLOOD-FP is well suited to parallelisation.