

Computational and Simulation Sciences: Applications of Heterogeneous Computing - speeding up Australian science

Tomasz Bednarz, Luke Domanski, Tim E Gureyev, Emma Huang, Pascal Vallotton, Dadong Wang, John A Taylor

Introduction

As the size and complexity of scientific problems and datasets grow, scientists from a broad range of discipline areas are relying more and more on computational methods and simulations to help solve their problems. This work presents a summary of heterogeneous algorithms and applications that have been developed by a large research organization (CSIRO) for solving practical and challenging science problems faster than is possible with conventional multi-core CPUs alone. The problem domains include: computational fluid dynamics, imaging and visualisation, advanced materials modelling, computational biology, chemistry, geosciences and climate and weather research. The algorithms utilize NVIDIA GPUs and multi-core CPUs on a scale ranging from single workstation installations through to large GPU clusters.

Neurite Detection in High-Content Analysis

Analysing the structure of neurons and neurites (thin tree structures mediating neuron communication) in microscope images is an important task in the study of inhibitors of degenerative brain disease or of the regenerative capabilities of brain cells after trauma. Automated techniques for performing this analysis are vital, as manual annotation and measurement by humans is time consuming, tedious, and error prone. We have used NVIDIA GPUs and CUDA to accelerate the neurite detection algorithm [1] included in our HCA-Vision software [2]. Tests showed speed-ups of up to 13x for accelerated algorithm components on a GeForce GTX260 over a single thread running on a 2.3GHz Intel Xeon CPU.

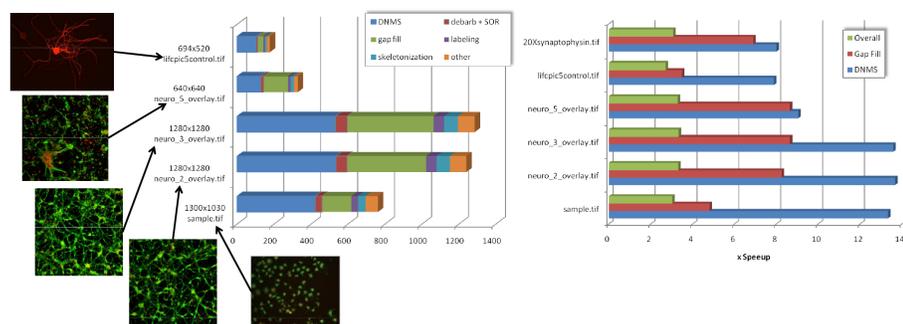


Figure 1: Neurite detection – CPU performance (left plot) and GPU speed up (right plot).

Genetic Analysis in Agriculture

In recent years, the size and complexity of genetic studies have increased dramatically due to improvements in genotyping technology. In agriculture, the CSIRO is breeding the world's first Multiparent Advanced Generation InterCrosses (MAGIC) in wheat [3]. These novel populations include thousands of wheat lines genotyped at tens of thousands of markers. The complexity of the design and the wealth of data require computational solutions which are both efficient and powerful. We are developing GPU implementations of analyses such as genetic map construction and quantitative trait locus mapping. For a small-scale problem with 1000 individuals genotyped at 1000 genetic markers, we have reduced the time for analysis from 2176 seconds on a Xeon E5462 compute node in the CSIRO GPU cluster to 25 seconds on a Tesla S2050. For a cross with 5000 individuals and 10000 markers (currently being genotyped) this analysis would require nearly two weeks on the Xeon, in contrast to only a few hours on the GPU. This reduction in compute time is crucial for practical analysis.

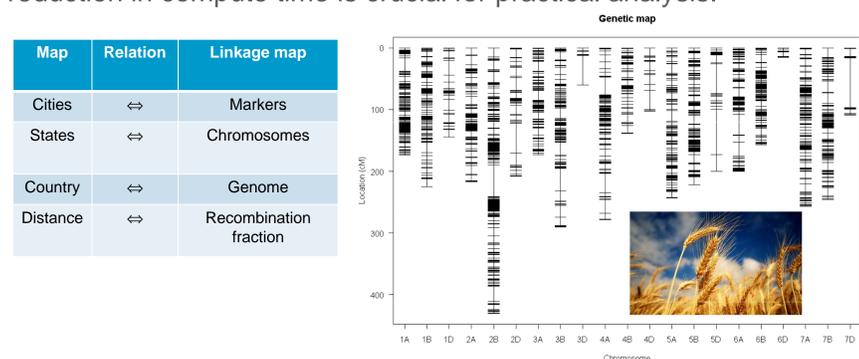


Figure 2: Left: map vs linkage map, Right: sample genetic map.

Computational Fluid Dynamics

The main aim of the Computational Fluid Dynamics (CFD) simulations is to reconstruct the reality of fluid motion and behaviour as accurately as possible in order to better understand natural phenomena under specified conditions. Ideally, general solutions can also cover different scales and geometric configurations. With the advent of high-performance GPUs with massively-parallel multi-threaded architectures, basic CFD algorithms can now be implemented to give results in near real-time. For instance our OpenCL based implementation of HSMAC code [4] showed speed ups of ~200x, when compared S2050 GPU execution times against Intel's i7 2.66Ghz.

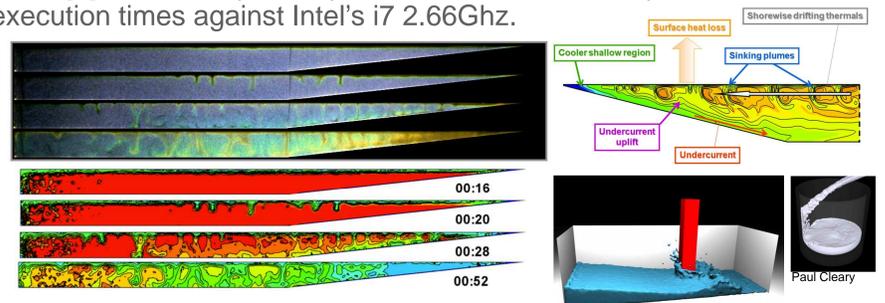


Figure 3: Examples problems: exchange flow in reservoir model and smoothed particle dynamics simulations.

Remote Computed Tomography Reconstruction Services

A software package for X-ray image analysis, processing and simulation called "XLI" [5] has been developed in CSIRO. We developed implementations for some of the most computationally intensive algorithms in XLI that allow them to run on GPUs, which can be utilised both at the client and server ends (our GPU Cluster). A typical CT data set obtained with a CCD-based X-ray detector with $(4k)^2$ pixels over multiple view angles is of the order of 128 GB. If a conventional algorithm is used to reconstruct the 3D internal structure of the sample from such CT data, it typically takes several days of computations using a modern single core CPU. In the case of GPUs, we have demonstrated speed-up factors of more than 200x in the back-projection step on a GeForce GTX280 GPU compared to a single Xeon E5420 CPU core. On a large mixed CPU/GPU cluster the system can currently reconstruct a volume with $\sim(4k)^3$ voxels in approximately 10 minutes using 17 nodes with 204 Xeon CPU cores and 34 NVIDIA Tesla M2070 GPUs.



Figure 4: Reconstructed Images, from left: fossil bryozoan, oil/water interface in sand, acacia plant, loblolly pine.

3-D Deconvolution

Deconvolution is an important operation in many areas of science, including astronomy, microscopy, and medical imaging. It reduces the effects of blurring introduced during image capture, revealing objects and details that may not have been visible in the raw image.

The Richardson-Lucy algorithm [6] is used to find the true object in the presence of Poisson image noise by iteratively re-applying the imaging model in the discrete space towards an improved estimate. Fig.5 shows the performance of the algorithm on up to 64 GPU workers. Result are also shown for CPU workers using Intel's MKL FFT library. One GPU from a Tesla S1070 or S2050 can process a sub-tile about 3.6x or 5.4x faster than a single CPU core.

The performance for the CPUs and S1070 GPUs scale almost linearly, however, we notice the S2050 results do not scale well past 16 nodes, and that they converge with the S1070 results. This communication overheads dominate past 16 workers.

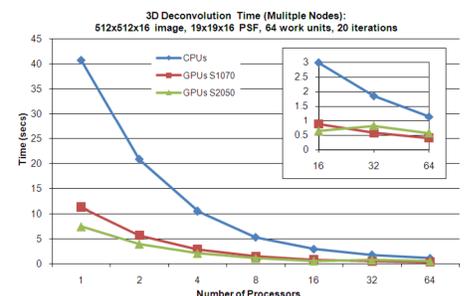


Figure 5: 3D deconvolution on the GPU Cluster.



References

- Domanski L. et al., DOI:10.1109/DICTA.2010.112
- CSIRO CMIS, HCA-Vision website: <http://www.hca-vision.com/>
- Cavanagh C. et al., <http://bit.ly/wPidvq>
- Bednarz T., <http://www.mssanz.org.au/modsim2011/A8/bednarz.pdf>
- Website: <http://www-ts.imaging.net/Services/AppInfo/X-TRACT.aspx>
- Richardson WH, <http://alturl.com/hk3c3>

Acknowledgements

We would like to thank CSIRO ASC staff for their tireless efforts in maintaining and supporting the CSIRO GPU Cluster.

Other people involved

Ya.I.Nesterets, D.Thompson, E.Jones, C.Cavanagh, C.Sun, N.Margvelashvili, A.Khassapov, A.Stevenson, S.Mayo, M.Morell and A.W. George

Further information

contact: Dr Tomasz Bednarz
phone: +61-2-9325-3213
email: tomasz.bednarz@csiro.au
web: www.csiro.au/org/CMIS

www.csiro.au

