

Enabling GPU Computing in the R Statistical Environment

Josh Buckner¹, Manhong Dai¹, Brian Athey^{1,2}, Stanley Watson¹ and Fan Meng^{1,2}

¹Molecular & Behavioral Neuroscience Institute and Psychiatry Department

²National Center for Integrative Biomedical Informatics
University of Michigan, Ann Arbor, MI 48109, US
bucknerj@umich.edu


Project URL: <http://brainarray.mbni.med.umich.edu/Brainarray/Rgpgpu/>

Code URL: http://brainarray.mbni.med.umich.edu/Brainarray/Rgpgpu/gputools_0.1-0.tar.gz

License: GPLv3

R is the most popular open source statistical environment in the biomedical research community. However, most of the popular R function implementations involve no parallelism and they can only be executed as separate instances on multicore or cluster hardware for large data-parallel analysis tasks. The arrival of modern graphic processing units (GPUs) with user friendly programming tools, such as nVidia's CUDA toolkit (<http://www.nvidia.com/cuda>), provides a possibility of increasing the computational efficiency of many common tasks by more than one order of magnitude (<http://gpgpu.org/>). However, most R users are not trained to program a GPU, a key obstacle for the widespread adoption of GPUs in biomedical research.

To overcome this obstacle, we decided to devote efforts for moving frequently used R functions in our work to the GPU using CUDA. In the ideal solution, if a CUDA compatible GPU and driver is present on a user's machine, the user may only need to prefix "gpu" to the original function name to take advantage of the GPU implementation of the corresponding R function. We take achieving this ideal as one of our primary goals so that any biomedical researcher can harness the computational power of a GPU using a familiar tool. Since our code is open source, researchers may customize the R interfaces to their particular needs. In addition, because CUDA uses shared libraries and unobtrusive extensions to the C programming language, any experienced C programmer can easily customize the underlying code.



Using the CUDA extension to C and the shared linear algebra library CUBLAS, we have implemented a variety of statistical analysis functions with R interfaces that execute with different degrees of parallelism on a Graphics Processing Unit (GPU). If an algorithm is comprised of common vector or matrix operations each performed once, we involve the GPU by implementing those operations with calls to CUBLAS. If an algorithm involves computing the elements of a large matrix, we can often merely assign each thread executing on the GPU a portion of a row and/or column. Algorithms for which we have implemented GPU enabled versions include the calculations of distances between sets of points (R dist function), hierarchical clustering (R hclust function). Pearson and Kendall correlation coefficients (similar to R cor function), and the Granger test (granger.test in the R MSBVAR package).

Figure 1 provides performance comparison between original R functions assuming a four thread data parallel solution on Intel Core i7 920 and our GPU enabled R functions for a GTX 295 GPU. The speedup test consisted of testing each of three algorithms with five randomly generated data sets. The Granger causality algorithm was tested with a lag of 2 for 200, 400, 600, 800, and 1000 random variables with 10 observations each. Complete hierarchical clustering was tested with 1000, 2000, 4000, 6000, and 8000 points. Calculation of Kendall's correlation coefficient was tested with 20, 30, 40, 50, and 60 random variables with 10000 observations each.

We are committed to implement more R GPU functions, and we hope to contribute packages to the open source community via our project's website. We hope that others can contribute to the R-GPGPU effort and encourage any comments or suggestions.

Acknowledgements: J. Buckner, M. Dai, F. Meng, and S. J. Watson are members of the Pritzker Neuropsychiatric Disorders Research Consortium, which is supported by the Pritzker Neuropsychiatric Disorders Research Fund L.L.C. This work is also partly supported by the National Center for Integrated Biomedical Informatics through NIH grant 1U54DA021519-01A1 to the University of Michigan.