

Large-scale genome-wide association studies on a GPU cluster using a CUDA-accelerated PGAS programming model

The International Journal of High Performance Computing Applications
2015, Vol. 29(4) 506–510
© The Author(s) 2015
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1094342015585846
hpc.sagepub.com


Jorge González-Domínguez¹, Jan Christian Kässens², Lars Wienbrandt²
and Bertil Schmidt¹

Abstract

Detecting epistasis, such as 2-SNP interactions, in genome-wide association studies (GWAS) is an important but time consuming operation. Consequently, GPUs have already been used to accelerate these studies, reducing the runtime for moderately-sized datasets to less than 1 hour. However, single-GPU approaches cannot perform large-scale GWAS in reasonable time. In this work we present multiEpistSearch, a tool to detect epistasis that works on GPU clusters. While CUDA is used for parallelization within each GPU, the workload distribution among GPUs is performed with Unified Parallel C++ (UPC++), a novel extension of C++ that follows the Partitioned Global Address Space (PGAS) model. multiEpistSearch is able to analyze large-scale datasets with 5 million SNPs from 10,000 individuals in less than 3 hours using 24 NVIDIA GTX Titans.

Keywords

PGAS, CUDA, GPU, UPC++, bioinformatics

1. Introduction

Modern high-throughput technologies are able to gather information of millions of single nucleotide polymorphisms (SNPs) from thousands of individuals for genome-wide association studies (GWAS). The most common phenotype classification is a binary trait, i.e. the presence (case) or absence (control) of an associated disease. 2-SNPs analyses try to find pairs of SNPs whose joint genotype frequencies show a statistically significant difference between cases and controls which potentially explains the effect of the genetic variation leading to disease (epistasis). Computing epistasis is highly time-consuming due to the large number of pairwise tests to be calculated. Targeting this problem with high-performance computing (HPC) architectures can help to speedup the process. Parallel codes exist to detect epistasis (González-Domínguez et al., 2014; Yung et al., 2011). The best of these approaches are able to reduce the time to analyze a moderately sized dataset from several days on a traditional CPU (Wan et al., 2010) to around 1 hour. Nevertheless, to the best of the authors' knowledge, none of them is able to

perform the analysis of datasets that contain millions of SNPs in an acceptable time.

In this paper we present multiEpistSearch, a tool to detect epistasis on a GPU cluster. Our approach employs a hybrid implementation with CUDA for intra-GPU parallelism and Unified Parallel C++ (UPC++) (Zheng et al., 2014) to distribute data among different GPUs. UPC++ is a novel extension of C++ that combines the advantages of both Partitioned Global Address Space (PGAS) and object-oriented paradigms. To the best of the authors' knowledge, this is the first hybrid implementation of an application using CUDA and UPC++, as most of

¹Parallel and Distributed Architectures Group, Johannes Gutenberg University-Mainz, Mainz, Germany

²Department of Computer Science, Christian-Albrechts-University of Kiel, Kiel, Germany

Corresponding author:

Jorge González-Domínguez, Parallel and Distributed Architectures Group, Johannes Gutenberg University, Staudingerweg 9, 55128 Mainz, Germany.
Email: j.gonzalez@uni-mainz.de

scientific codes for GPU clusters employ MPI for inter-GPU parallelism.

2. Hybrid implementation

Our GWAS tool works with datasets containing information about a large number of biallelic genetic markers from many individuals. For each marker (SNP) there are three genotypes numerically represented as $\{0, 1, 2\}$. Each individual is further characterized as case or control, depending on the presence or absence of an associated disease. Two SNPs present epistasis or interaction if their combination discriminates between cases and controls significantly better than discrimination using each SNP individually. The number of SNPs and individuals is denoted as M and N , respectively.

Following the same approach as in González-Domínguez et al. (2014), three filters (KSASA, KSA and log-linear) are applied to identify which of the $\frac{M(M-1)}{2}$ SNP-pairs present interaction. All of them use a $3 \times 3 \times 2$ contingency table per pair that needs to be calculated from the N individual genotype pair values before the filters are applied. Our parallel implementation creates one UPC++ process per GPU and distributes the workload among them. Each GPU analyzes whether different SNP pairs present epistasis through a CUDA kernel. In order to balance the memory usage per node, the biallelic information of the SNPs is distributed in blocks in a round-robin way, where the user can specify the number of blocks per process. This information is stored in shared memory so it can be directly accessed by all processes. A similar data distribution has been used in our UPC++ implementation for multicore clusters (see Kässens et al., 2014 for more details). When adapting this distribution to the hybrid CUDA&UPC++ implementation the GPUs may need information stored in remote memory to analyze the SNP pairs. In this scenario the UPC++ process associated to that GPU must access the corresponding remote data and load it into the GPU before calling the kernel.

The flowchart in Figure 1 summarizes the behavior of each UPC++ process. The term “metablock” describes each block of SNP pairs that must be computed, i.e. all possible combinations with the blocks of SNPs. In an example with 3 processes where the biallelic data is distributed using 2 blocks per process (6 blocks in total) there are 21 “metablocks” of SNP pairs. For each associated metablock, if necessary, the process initially reads the biallelic information from remote memory and loads it into the GPU memory. Then, it launches the CUDA kernel so that the GPU searches for epistasis in all of the SNP pairs within the metablock. Thanks to the shared global memory space and the one-sided communications available in PGAS, remote copies can be performed without

synchronization with the owner. Moreover, the execution of the CUDA kernel and the copy of the information needed to compute the next metablock are overlapped using asynchronous communication.

We have developed two approaches to distribute the metablocks among the processes and GPUs. First, a static distribution where the metablocks are initially associated to processes. The workload is balanced (by analyzing similar number of metablocks per GPU). The number of copies from remote memory is also minimized. Our distribution guarantees that metablocks are computed by processes that already have at least one of the necessary information blocks in their local memory. Second, an on-demand distribution has been implemented, where metablocks are not initially assigned to certain UPC++ processes. A table of UPC++ locks, with one open lock per metablock, is created in global memory (accessible by all processes). Every time one kernel call finishes, the process accesses this shared table to know which metablocks have not been or are not being computed at the same time by other GPUs, i.e. those metablocks whose lock is still open. Once one idle process finds a metablock to compute, it closes the associated lock.

To increase the efficiency of the UPC++ code, a set of optimization techniques has been applied, mostly focused on minimizing the communication cost. The utilized CUDA kernel is similar to the one explained in González-Domínguez et al. (2014), but it includes two optimizations: (1) a reorder of the biallelic information before loading it to the GPU to improve coalescence; (2) exploitation of fast shared memory by the threads of the same thread block.

3. Experimental evaluation

Two different test systems have been employed in the experimental evaluation. (1) Eight nodes of the Mogan supercomputer, installed at the JGU Mainz, connected through a QDR InfiniBand network. Each node contains 3 NVIDIA GTX Titan GPUs (i.e. 24 GPUs in total). (2) Twelve nodes of Pluton, installed at the University of A Coruña, connected through a Gigabit Ethernet network. Eight of the nodes contain one NVIDIA Tesla K20m and four nodes contain two NVIDIA Tesla 2050.

We have used a real-world dataset obtained from the Wellcome Trust Case-Control Consortium (WTCCC) (The Wellcome Trust Case Control Consortium, 2007) consisting of 3004 controls and 2005 cases genotyped at 500,568 SNPs. In order to provide a fair comparison, the influence of the number of blocks per process is removed by always showing the result for the best block size for each scenario. multiEpistSearch provides an autotuning option in order to automatically identify the best value for this parameter. The autotuning process

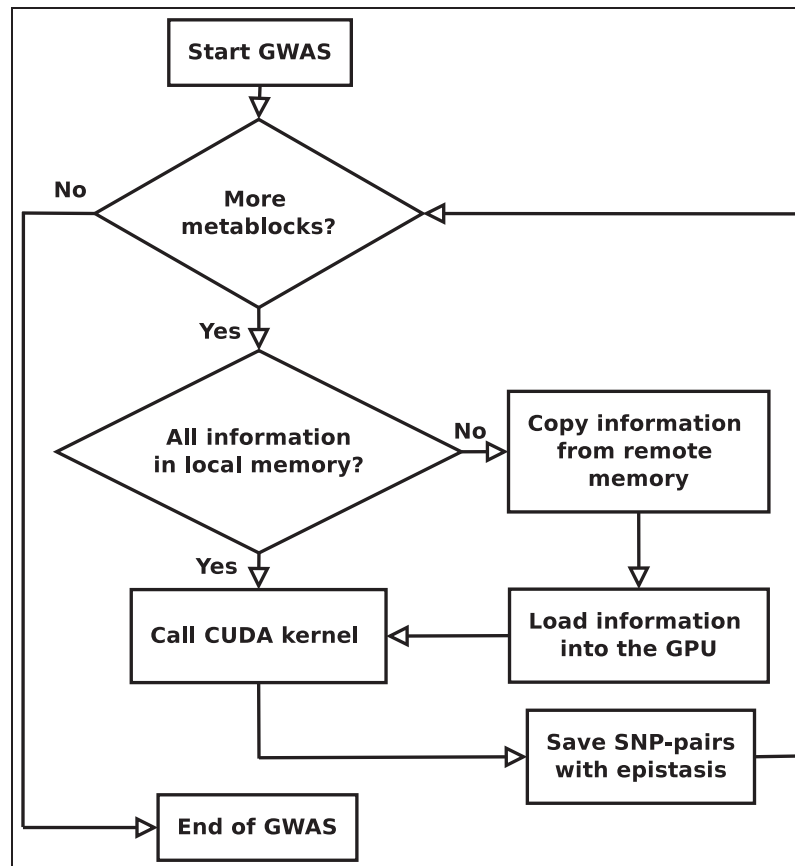


Figure 1. Procedure within each UPC ++ process.

uses information about the runtime of previous executions.

Figure 2 compares the runtime of the distributions described in Section 2 for a varying number of GPUs. The top graph shows that static distribution is better for clusters with the same type of GPUs, with parallel efficiency over 95%. When all of the GPUs are similar, the on-demand version tends to equally distribute the workload. Thus, the workload distribution is the same for both approaches but the accesses to the locks present in the on-demand version lead to a performance overhead that makes it less efficient. However, results for the on-demand distribution are better on the cluster with heterogeneous GPUs, even though it includes the overhead of working with locks; e.g. it is 15% better when using all of the GPUs. While in the static distribution all GPUs analyze a similar number of SNP pairs, the on-demand approach distributes the workload so that the faster GPUs (Tesla K20m) compute more pairs than the slower ones (Tesla 2050).

Table 1 gathers the runtime of multiEpistSearch and other approaches when analyzing the WTCC dataset. Regarding the related work, we have executed the fastest publicly available tool (GBOOST (Yung et al., 2011)) utilizing only one GPU, as it does not support

computation for multiple GPUs. Finally, we have estimated the execution time for BOOST (Wan et al., 2010) on an Intel Core i7 by analyzing a smaller simulated dataset and assuming quadratic increase of time with the number of SNPs. The experimental results show that our CUDA implementation is 2.78 times faster than GBOOST using the same hardware (one GTX Titan). An additional advantage of our tool is the multi-GPU support, which allows us to obtain a speedup of 54.93 using 24 GTX Titans. According to the estimation of the sequential CPU-based BOOST, multiEpistSearch obtains speedups of more than 373 and 8500 against a 3 GHz CPU using 1 and 24 GTX Titans, respectively. Finally, thanks to the on-demand approach, the runtime on *Pluton* is reduced from 5 minutes and 40 seconds using only 8 Tesla K20m to 4 minutes and 20 seconds on the whole cluster, even though the Tesla 2050 is around 1.5 times slower than the Tesla K20m.

A simulated dataset with 5 million SNPs and 10,000 samples has also been analyzed by multiEpistSearch in 2 hours and 45 minutes using 24 GTX Titans. A similar execution using only 1 GPU would need several days, which demonstrates the need of multi-GPU GWAS tools in order to perform large-scale GWAS analyses in




Figure 2. Performance comparison of the two multiEpistSearch workload distribution schemes. The labels indicate the speedup compared to multiEpistSearch executed on a single GTX Titan (left) and on one Tesla K20m (right).

Table 1. Runtime of different designs when looking for epistasis in the WTCCC dataset. *The results for CPU BOOST were estimated from a smaller dataset.

Design	Architecture	Runtime
multiEpistSearch	24 GTX Titan	1 m 11 s
multiEpistSearch	8 Tesla K20m + 8 2050	4 m 20 s
multiEpistSearch	8 Tesla K20m	5 m 40 s
multiEpistSearch	8 Tesla 2050	10 m 12 s
multiEpistSearch	1 GTX Titan	27 m
GBOOST	1 GTX Titan	1 h 15 m
BOOST*	Intel Core i7	7 d

reasonable time. Such a large dataset could not be analyzed by GBOOST due to out-of-bounds problems in the internal arrays.

4 Conclusions

We have presented multiEpistSearch, a tool that exhaustively measures the interaction of all SNP pairs of a GWAS dataset on GPU clusters using CUDA and UPC ++, a new PGAS extension of C ++. Static and on-demand approaches have been implemented for the workload distribution among GPUs. They have been tested on two different GPU clusters. The static distribution, that initially assigns metablocks to GPUs, obtains the best performance on Mogon, a cluster with 24 GTX Titans. The on-demand distribution, where the workload is dynamically assigned to the GPUs during execution time, is more suitable for clusters such as Pluton, where there are different types of GPUs. Our evaluation shows that multiEpistSearch is able to exhaustively look for epistasis in large-scale SNP datasets in feasible time. For instance, it takes less than 3 hours to analyze a dataset with 5 million SNPs and 10,000 individuals on Mogon.

Acknowledgements

This study makes use of data generated by the Wellcome Trust Case-Control Consortium. A full list of the

investigators who contributed to the generation of the data is available from <http://www.wtccc.org.uk>. We gratefully thank the Computer Architecture Group at the University of A Coruña for providing access to Pluton.

Funding

This work was supported by the Wellcome Trust (award numbers 076113 and 085475).

References

- González-Domínguez J, Schmidt B, Wienbrandt L, et al. (2014) Hybrid CPU/GPU acceleration of detection of 2-SNP epistatic interactions in GWAS. In: *Proceedings 15th international European conference on parallel and distributed computing (Euro-Par'14)*.
- Kässens JC, González-Domínguez J, Wienbrandt L and Schmidt B (2014) UPC ++ for bioinformatics: a case study using genome-wide association studies. In: *Proceedings 15th IEEE international conference on cluster computing (Cluster'14)*.
- The Wellcome Trust Case Control Consortium (2007) Genome-wide association study of 14,000 cases of seven common diseases and 3,000 shared controls. *Nature* 447(7145): 661–678.
- Wan X, Yang C, Yang Q, et al. (2010) BOOST: a fast approach to detecting gene–gene interactions in genome-wide case-control studies. *American Journal of Human Genetics* 87(3): 325–340.
- Yung LS, Yang C, Wan X, et al. (2011) GBOOST: a GPU-based tool for detecting gene–gene interactions in genome-wide case control studies. *Bioinformatics* 27(9): 1309–1310.
- Zheng Y, Kamil A, Driscoll M, et al. (2014) UPC ++: a PGAS extension for C ++. In: *Proceedings 28th IEEE international parallel and distributed processing symposium (IPDPS'14)*.

Author biographies

Jorge González-Domínguez received the BSc, MSc and PhD degrees in Computer Science from the University of A Coruña, Spain, in 2008, 2010 and 2013, respectively. He is currently a postdoctoral researcher in the Parallel and Distributed Architectures Group at the

Johannes Gutenberg University Mainz, Germany. His main research interests are in the areas of high-performance computing for bioinformatics and PGAS programming languages.

Jan Christian Kässens received his MSc in Computer Science from the Christian-Albrechts-University of Kiel, Germany, in 2012. He is currently a research assistant at the Technical Computer Science group at the CAU and working on his PhD. His research focus lies on hardware/software co-development, hardware-assisted parallelization and FPGA technology.

Lars Wienbrandt received his MSc (Dipl.-Inf.) in Computer Science from the Christian-Albrechts-University of Kiel, Germany, in 2009. He is currently a research assistant at the Technical Computer Science group at the CAU and working on his PhD. His research area includes the parallelization and

implementation of bioinformatics algorithms on FPGA architectures.

Bertil Schmidt (M'04–SM'07) is tenured Full Professor and Chair for Parallel and Distributed Architectures at the University of Mainz, Germany. Prior to that he was a faculty member at Nanyang Technological University (Singapore) and at University of New South Wales (UNSW). His research group has designed a variety of algorithms and tools for bioinformatics mainly focusing on the analysis of large-scale sequence and short read datasets. For his research work, he has received a CUDA Research Center award, a CUDA Academic Partnership award, a CUDA Professor Partnership award and the Best Paper Award at IEEE ASAP 2009. Furthermore, he serves as the champion for bioinformatics and computational biology on <http://www.gpucomputing.net>.