

Accelerating QTL mapping with graphics cards

Guillaume CHAPUIS¹, Olivier FILANGI², Jean-Michel ELSÉN³, Dominique LAVENIER¹ and Pascale LE ROY²

¹ IRISA, UMR 6074, Campus de Beaulieu, 263 avenue du Général Leclerc - Bâtiment 12 35 042 Rennes, cedex, France
{Guillaume.Chapuis, Dominique.Lavenier}@irisa.fr

² INRA Agrocampus Ouest, UMR Pegase, 65 rue de Saint-Brieuc, INRA, 35042 Rennes, France
{Olivier.Filangi, Pascale.Leroy}@rennes.inra.fr

³ INRA, UPR 0631, Chemin de Borde-Rouge - Auzeville BP 52627 31326 CASTANET-TOLOSAN, cedex, France
{Jean-Michel.Elsen}@toulouse.inra.fr

Keywords QTL mapping, GPGPU, GPU, linkage analysis, linkage disequilibrium, multiQTL analysis, QTLMap, Double-Precision, Cuda

1 Context

Most of the traits characterizing individuals are influenced by heredity. Geneticists are interested in detecting, localizing and identifying genes, the polymorphism of which explains a part of observed trait variability. Such genes are often called QTLs (for "Quantitative Trait Locus"), the term locus pointing to a physical position on the genome.

QTL detection procedures consist in a series of statistical hypotheses tests at successive putative locations on the genome. We focus here on regression approaches performed on sets of large families. These approaches were developed for exploiting both the linkage and the gametic association between loci observed on a per family basis and / or at the population level. QTLMap implements three types of QTL analyses: Linkage analysis (LA), Linkage Disequilibrium Analysis (LDA), and Linkage Disequilibrium Linkage Analysis (LDLA).

QTL mapping analyses are computationally intensive. They often take weeks to run on modern computers and run times increase linearly with the density of available genetic markers.

Despite the computational burden that QTL mapping represents, few parallel tools exist. The first attempt was made by [3] with gridQTL. This tool is derived from QTLexpress ([2]), a popular web based tool for QTL analyses, and harnesses the power of computational grids to try and reduce run times. The previous version of QTLMap, developed by [1], takes advantage of modern CPUs by using all their cores simultaneously.

We have developed a new version of QTLMap, which takes advantage of Graphics Processing Units (GPUs). The empirical approach used in QTLMap makes it an ideal candidate for GPU computations. Statistical tests computed at each genome position and for each simulation are almost identical in terms of instructions and also completely independent. This type of data parallelism is a perfect fit for the single instruction multiple data (SIMD) architecture of GPUs. The new GPU implementation of QTLMap performs about 70 times faster than the previous multicore implementation, while maintaining the same level of precision.

2 Experiments and results

Tests were run on machines with two quadcore Intel® Xeon® E5420 (12M Cache, 2.50 GHz, 1333 MHz FSB) processors. Multicore cpu tests were run on the Genotoul platform (<http://www.genotoul.fr>). GPU tests were run on a machine equipped with an Nvidia® C2050 card. Each test consists of an LDLA analysis over a simulated dataset from the 2011 QTL-MAS workshop (<https://colloque.inra.fr/qtmmas>).

Fig. 1 shows the evolution of the speedup with respect to the number of studied genome positions between the previous multicore implementation and the GPU implementation in Double-Precision (DP). The openMP



version of QTLMap is not optimized for low numbers of genome positions ; therefore we observe a high speedup, when the number of genome positions is low. Fig. 1 shows a speedup of about 70 for the GPU version in Double-Precision over the openMP version.

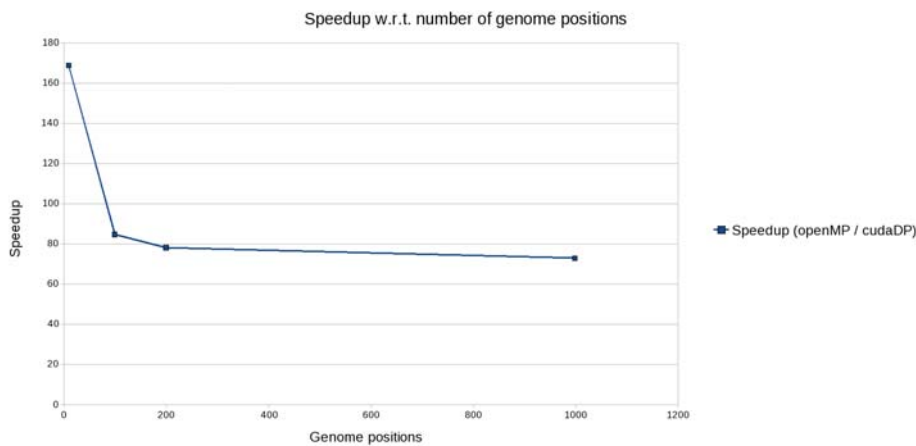


Figure 1. Speedup of GPU-QTLMap / OpenMP-QTLMap (8 cores) over the number of genome positions

Reduced runtimes allow geneticists to consider more precise and time consuming analyses by increasing the number of simulations or the number of studied genome positions. All versions of QTLMap are available under CeCILL licences at <http://www.inra.fr/qtldmap/>.

Future work include the promotion and use of parallel computing in statistical genetics, focusing on two applications of the Single Nucleotide Polymorphism (SNP) chip technology:

- Dissection of the genetic architecture of characters through Genome Wide Association Studies (GWAS);
- Genomic Selection (GS).

SNP chip technology now makes possible the genotyping on millions of SNPs of tens or hundreds of thousands of individuals, thus increasing the demand for much faster computations. Faster computations are needed both for implementing more precise genetic models in research of trait genetic determinants, and for the industrial exploitation of genomic data, with production of statistical information at regular time intervals. Our aim is to produce, when needed, new algorithms better suited for parallel architectures (GPUs and/or clusters of computers).

Acknowledgments

This work is part of the AAP INRA/INRIA and is also supported by the region of Brittany, France.

References

- [1] O. Filangi, C. Moreno, H. Gilbert, A. Legarra, P. Le Roy, and J. Elsen. Qtlmap, a software for qtl detection in outbred populations. In *Proceedings of the 9th World Congress on Genetics Applied to Livestock Production: 1-6 August: Leipzig*, number 787, 2010.
- [2] G. Seaton, C. S. Haley, S. A. Knott, M. Kearsey, and P. M. Visscher. Qtl express: mapping quantitative trait loci in simple and complex pedigrees. *Bioinformatics applications note*, 18(2):339–340, 2002.
- [3] G. Seaton, J. Hernandez, J. Grunchev, I. White, J. Allen, D. De Koning, W. Wei, D. Berry, C. Haley, and S. Knott. Gridqtl: a grid portal for qtl mapping of compute intensive datasets. In *Proceedings of the 8th World Congress on Genetics Applied to Livestock Production*, pages 13–18, 2006.