

High performance solutions for big-data GWAS



Elmar Peise*, Diego Fabregat-Traver, Paolo Bientinesi

AICES, RWTH Aachen, Schinkelstr. 2, 52062 Aachen, Germany

ARTICLE INFO

Article history:

Available online 8 October 2014

Keywords:

Genome-wide association study
 Mixed-models
 Generalized least squares
 Big data
 Distributed memory
 Omics

ABSTRACT

In order to associate complex traits with genetic polymorphisms, genome-wide association studies process huge datasets involving tens of thousands of individuals genotyped for millions of polymorphisms. When handling these datasets, which exceed the main memory of contemporary computers, one faces two distinct challenges: (1) millions of polymorphisms and thousands of phenotypes come at the cost of hundreds of gigabytes of data, which can only be kept in secondary storage; (2) the relatedness of the test population is represented by a relationship matrix, which, for large populations, can only fit in the combined main memory of a distributed architecture. In this paper, by using distributed resources such as Cloud or clusters, we address both challenges: the genotype and phenotype data is streamed from secondary storage using the double-buffering technique, while the relationship matrix is kept across the main memory of a distributed memory system. With the help of these solutions, we develop separate algorithms for studies involving only one or a multitude of traits. We show that these algorithms sustain high-performance and allow the analysis of enormous datasets.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Genome-wide association (GWA) analyses are a powerful statistical tool to identify certain locations of significance in the genome: typically, they aim at determining which *single-nucleotide polymorphisms* (SNPs) influences specific *traits* of interest. Thanks to these studies, hundreds of SNPs for dozens of complex human diseases and quantitative traits have been discovered [1]. In GWA studies (GWAS), one of the most used methods to account for the genetic substructure due to relatedness and population stratification is the variance component approach based on mixed-models [2,3]. While effective, mixed-models based methods are computationally demanding both in terms of data management and computation. The objective of this research is to make large-scale GWA analyses affordable.

Computationally, a mixed-model based GWAS on n individuals, m genetic markers (SNPs), and t traits boils down to the solution of the $m \times t$ generalized least-squares (GLS) problems

$$b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j, \quad \text{with } i = 1, \dots, m \quad \text{and } j = 1, \dots, t, \quad (1)$$

where $X_i \in \mathbb{R}^{n \times p}$ is the design matrix, $M_j \in \mathbb{R}^{n \times n}$ is the covariance matrix, $y_j \in \mathbb{R}^n$ contains the vector of observations, and $b_{ij} \in \mathbb{R}^p$ quantifies the relation between a variation in an SNP (X_i) and a variation in a trait (y_j). Furthermore, M_j is a symmetric positive definite (SPD) matrix, and the full rank matrix X_i can be viewed as composed of two parts: $X_i = (X_L | X_{Ri})$, with

* Corresponding author.

E-mail addresses: peise@aices.rwth-aachen.de (E. Peise), fabregat@aices.rwth-aachen.de (D. Fabregat-Traver), pauldj@aices.rwth-aachen.de (P. Bientinesi).

$X_L \in \mathbb{R}^{n \times (p-1)}$ and $X_{Ri} \in \mathbb{R}^{n \times 1}$, where X_L , which contains $p - 1$ covariates, is fixed, and only X_{Ri} varies with SNP_i . Moreover, the relationship among the individuals is taken into account by the covariance matrix M_j :

$$M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I \right). \quad (2)$$

Here, I is the identity matrix, the *kinship* matrix $\Phi \in \mathbb{R}^{n \times n}$ contains the relationship among all studied individuals, and σ_j^2 and h_j^2 are trait-dependent scalar estimates. Finally, common problem sizes are: $10^3 \leq n \leq 10^5$, $2 \leq p \leq 20$, $10^5 \leq m \leq 10^8$, and t is either 1 (single-trait analysis) or in the range of thousands (multi-trait analysis).

The first reported GWA study dates back to 2005: 146 individuals were genotyped, and about 103,000 SNPs were analyzed with respect to one trait [4]. Since then, as the catalog of published GWA analyses shows [5,6], the number of publications has increased steadily, up to 2404 in 2011 and 3307 in 2012. A similar growth can be observed in both the population size and the number of SNPs: across all the GWAS published in 2012 the studies comprised on average 15,471 individuals, with a maximum of 133,154, and on average 1,252,222 genetic markers, with a maximum of 7,422,970. More recently, advances in technology make it affordable to assess “omics” phenotypes in large populations, resulting in the challenge of analyzing (potentially hundreds) of thousands of traits. From the perspective of Eqs. (1) and (2), these trends present concrete challenges, especially in terms of memory requirements. As $M_j \in \mathbb{R}^{n \times n}$ and the m X_i 's and t y_j 's compete for the main memory, two distinct scenarios arise: (1) if n is small enough for M_j to fit in main memory, the X_i 's and y_j 's are to be streamed from disk; (2) if M_j does not fit in main memory, both data and computation have to be distributed over multiple compute nodes. In this paper, we present efficient strategies for utilizing distributed architectures—such as clusters, Cloud-based systems, and supercomputers—to execute single-trait and multi-trait GWA analyses with arbitrarily large population size, number of SNPs, and traits.

Related work. To perform GWA studies, there exist several freely available libraries. Among them, we highlight GENABEL, a widely spread framework for statistical genomics [7], and FAST-LMM, a high-performance software targeting single-trait analyses [8]. More recently, Fabregat et al. developed OmicABEL—a package for the GENABEL suite—which implements optimized solutions for shared memory architectures [9,10]. However, those algorithms do not support distributed-memory computations, and are only applicable when the kinship matrix fits in the local memory of a single node.

Organization of the paper. The rest of this paper is structured as follows. Section 2 is devoted to single-trait GWAS analyses (i.e., $t = 1$): In Section 2.1, the mathematical algorithm is first introduced; in Section 2.2, we discuss the shared memory implementation GWAS-1D-SMP, an extension to accommodate analyses with an arbitrarily large number of SNPs (large m); then, in Section 2.3, we present GWAS-1D-MPI, a distributed memory extension for analyses with large population size (large n). Section 3 addresses multi-trait studies (i.e., $t > 1$): A discussion of a core algorithm that exploits invariants across multiple traits is given in Section 3.1; to allow the solutions of problems of arbitrary size (in terms of m , n , and t), we apply out-of-core (Section 3.2) and distributed-memory techniques (Section 3.3), thus yielding GWAS-2D-MPI. Conclusions are drawn in Section 4.

2. Single-trait GWAS

We consider Eq. (1) restricted to the study of a single trait y :

$$b_i := \left(X_i^T M^{-1} X_i \right)^{-1} X_i^T M^{-1} y, \quad \text{with } i = 1, \dots, m. \quad (3)$$

2.1. The algorithm

The standard route to solving one such GLS is to reduce it to an ordinary least squares problem (OLS)

$$b_i = \left(\bar{X}_i^T \bar{X}_i \right)^{-1} \bar{X}_i^T \bar{y},$$

through the operations

- 1 $LL^T := M$ (Cholesky factorization)
- 2 $\bar{X}_i := L^{-1} X_i$ (triangular solve)
- 3 $\bar{y} := L^{-1} y$ (triangular solve)

The resulting OLS can then be solved by two alternative approaches, respectively based on the QR decomposition of \bar{X}_i , and the Cholesky decomposition of $\bar{X}_i^T \bar{X}_i$. In general, the QR-based method is numerically more stable; however, in this specific application, since $\bar{X}_i^T \bar{X}_i \in \mathbb{R}^{p \times p}$ is very small and X_i is typically well conditioned, both approaches are equally accurate. In terms of performance, the solution via Cholesky decomposition is slightly more efficient:

- 4 $S_i := \bar{X}_i^T \bar{X}_i$ (symmetric matrix product)
- 5 $\bar{b}_i := \bar{X}_i^T \bar{y}$ (matrix times vector)
- 6 $b_i := S_i^{-1} \bar{b}_i$ (linear system via Cholesky)

In this paper, we only consider this approach.

2.1.1. Multiple SNPs

When the six steps for the solution of one GLS are applied to the specific case of Eq. (3), by taking advantage of the structure of X_i , it is possible to avoid redundant computation.

Plugging $X_i = (X_L | X_{Ri})$ into $\bar{X}_i := L^{-1}X_i$ (line 2), we obtain

$$(\bar{X}_L | \bar{X}_{Ri}) := (L^{-1}X_L | L^{-1}X_{Ri}),$$

that is, $\bar{X}_L := L^{-1}X_L$, and $\bar{X}_{Ri} := L^{-1}X_{Ri}$. These assignments indicate that the quantity \bar{X}_L can be computed once and reused across all SNPs.

Similarly, for $S_i := \bar{X}_i^T \bar{X}_i$ (line 4), we have¹

$$\left(\begin{array}{c|c} S_{TL} & * \\ \hline S_{BLi} & S_{BRi} \end{array} \right) := \left(\begin{array}{c|c} \bar{X}_L^T \bar{X}_L & * \\ \hline \bar{X}_{Ri}^T \bar{X}_L & \bar{X}_{Ri}^T \bar{X}_{Ri} \end{array} \right),$$

from which

$$\begin{aligned} S_{TL} &:= \bar{X}_L^T \bar{X}_L \in \mathbb{R}^{(p-1) \times (p-1)}, \\ S_{BLi} &:= \bar{X}_{Ri}^T \bar{X}_L \in \mathbb{R}^{1 \times (p-1)}, \text{ and} \\ S_{BRi} &:= \bar{X}_{Ri}^T \bar{X}_{Ri} \in \mathbb{R}. \end{aligned}$$

This indicates that S_{TL} , the top left portion of S_i , is independent of i and needs to be computed only once.² Finally, the same idea also applies to \bar{b}_i (line 5), yielding the assignments $\bar{b}_T := \bar{X}_L^T \bar{y}$ and $\bar{b}_{Bi} := \bar{X}_{Ri}^T \bar{y}$.

Algorithm 1. Optimized algorithm for single-trait studies

```

1   $LL^T := M$ 
2   $\bar{X}_L := L^{-1}X_L, \bar{y} := L^{-1}y$ 
3   $S_{TL} := \bar{X}_L^T \bar{X}_L, \bar{b}_T := \bar{X}_L^T \bar{y}$ 
4  for  $i$  in  $\{1, \dots, m\}$ 
5     $\bar{X}_{Ri} := L^{-1}X_{Ri}$ 
6     $S_{BLi} := \bar{X}_{Ri}^T \bar{X}_L$ 
7     $S_{BRi} := \bar{X}_{Ri}^T \bar{X}_{Ri}$ 
8     $\bar{b}_{Bi} := \bar{X}_{Ri}^T \bar{y}$ 
9    set  $S_i := \left( \begin{array}{c|c} S_{TL} & * \\ \hline S_{BLi} & S_{BRi} \end{array} \right), \bar{b}_i := \left( \begin{array}{c} \bar{b}_T \\ \hline \bar{b}_{Bi} \end{array} \right)$ 
10    $b_i := S_i^{-1} \bar{b}_i$ 
11 end

```

The computation for the whole Eq. (3) is given in Algorithm 1. There, by moving all the operations independent of i outside the loop, the overall complexity is lowered from $O(n^3 + mn^2p)$ down to $O(n^3 + mn^2)$.³ This algorithm constitutes the basis for the large-scale versions presented in the next two sections.

2.2. Out-of-core

GWA studies often operate on and generate datasets that exceed the main memory capacity of current computers. For instance, a study with $n = 20,000$ individuals, $m = 10,000,000$ SNPs, and $p = 4$, requires 1.49 TB to store the input data (M and X_i 's), and generates 305 MB of output.⁴ To make large analyses feasible, regardless of the number of SNPs, Fabregat et al. proposed an algorithm that uses asynchronous I/O operations to stream X_{Ri} and b_i from and to secondary storage [9]. This extension of Algorithm 1 is described in the following.

¹ The subscript letters L, R, T , and B , respectively stand for Left, Right, Top, and Bottom.

² Since S_i is symmetric, its top-right and bottom-left quadrants are the transpose of each other; we mark the top-right quadrant with a * to indicate that it is never accessed nor computed.

³ Since in most analyses $m \gg n$, the complexity reduces by a factor of p , from $O(mn^2p)$ down to $O(mn^2)$.

⁴ In practice the size of the output is even larger, because along with each b_i , a symmetric $p \times p$ matrix is generated.

Algorithm 2. Out-of-core algorithm for single-trait studies. The X_{Ri} and b_i are streamed from and to disk in blocks. Asynchronous I/O operations are highlighted in green

```

1   $LL^T := M$ 
2   $\bar{X}_L := L^{-1}X_L, \bar{y} := L^{-1}y$ 
3   $S_{TL} := \bar{X}_L^T \bar{X}_L, \bar{b}_T := \bar{X}_L^T \bar{y}$ 
4  load_start first  $X_{blk}$ 
5  for each  $blk$ 
6    load_wait current  $X_{blk}$ 
7    if not last  $blk$ : load_start next  $X_{blk}$ 
8     $\bar{X}_{blk} := L^{-1}X_{blk}$ 
9    for  $i$  in  $\{1, \dots, m_{blk}\}$ 
10     set  $\bar{X}_{Ri} := \bar{X}_{blk}[i]$ 
11      $S_{BLi} := \bar{X}_{Ri}^T \bar{X}_L, S_{BRi} := \bar{X}_{Ri}^T \bar{X}_{Ri}$ 
12      $\bar{b}_{Bi} := \bar{X}_{Ri}^T \bar{y}$ 
13     set  $S_i := \left( \begin{array}{c|c} S_{TL} & * \\ \hline S_{BLi} & S_{BRi} \end{array} \right), \bar{b}_i := \left( \begin{array}{c} \bar{b}_T \\ \hline \bar{b}_{Bi} \end{array} \right)$ 
14      $b_i := S_i^{-1} \bar{b}_i$ 
15     set  $b_{blk}[i] := b_i$ 
16   end
17   if not first  $blk$ : store_wait previous  $b_{blk}$ 
18   store_start current  $b_{blk}$ 
19 end
20 store_wait last  $b_{blk}$ 

```

In order to avoid any overhead, the vectors X_{Ri} (and b_i) are grouped into blocks X_{blk} (and b_{blk}) of size m_{blk} , and read (written) asynchronously using double-buffering. This well-known technique logically splits the main memory in two equal regions: while one region is devoted to the block of data that is currently processed, the other is used to both store the output from the previous block and load the input for the next one. Once the computation on the current block is completed, the roles of the two regions are swapped. The algorithm commences by loading the first block of SNPs X_{blk} from disk into memory; then, while the GLS's corresponding to this block are solved, the next block of SNPs is loaded asynchronously in the second memory region. (Analogously, the previous b_{blk} is stored, while the current one is computed.)

When dealing with large analyses, an important optimization comes from, whenever possible, processing multiple SNPs at once: Algorithm 2 shows how combining slow vector operations on X_{Ri} together originates efficient matrix operations on $X_{blk} \in \mathbb{R}^{n \times m_{blk}}$ (line 8).

2.2.1. Performance results: GWAS-1D-SMP

The shared memory implementation of Algorithm 2, here called GWAS-1D-SMP, makes use of parallelism in two different ways [9]. The operations in lines 1 through 8 are dominated by Level 3 Basic Linear Algebra Subroutines (BLAS) and take full advantage of a multithreaded implementation of BLAS and the Linear Algebra PACKage (LAPACK). By contrast, for the operations within the innermost loop (lines 11 through 14), which only involve very small or thin matrices, BLAS and especially multithreaded BLAS are less efficient. Therefore, they are scheduled in parallel using OPENMP in combination with single-threaded BLAS and LAPACK.

We compile GWAS-1D-SMP, written in C, with the GNU C compiler (GCC version 4.4.5) and link to Intel's Math Kernel Library (MKL version 10.3). All tests are executed on a system consisting of two six-core Intel X5675 processors, running at 3.06 GHz, equipped with 32 GB of RAM, and connected to a 1 TB hard disk.

Preliminary measurements showed that changing $p \in \{1, \dots, 20\}$ results in performance variations on the order of system fluctuations (below 1%). Therefore $p = 4$, a value encountered in several GWA studies, is considered throughout all our experiments.

In the first experiment, we compare the efficiency of GWAS-1D-SMP with GWAS-BASE, an equivalent in-core version. Fixing $n = 10,000$, $p = 4$, and we let m vary between 10^3 and 10^7 . For the out-of-core version, the SNPs are grouped in blocks of size $m_{blk} = 5000$. As Fig. 1 shows, GWAS-1D-SMP scales linearly in the number of SNPs well beyond the maximum problem size imposed by the 32 GB of RAM. Furthermore, the fact that the lines for the in-core and out-of-core algorithms overlap perfectly confirms that the I/O operation from and to disk are entirely hidden by computation.

In the second experiment, Fig. 2, we compare the performance of GWAS-1D-SMP to that of two other solvers: FAST-LMM, a program designed for GWAS on large datasets [8], and GENABEL, a widely spread library for genome studies [7]. Again, fixing $n = 10,000$ and $p = 4$, m varies between 10^6 and $3.6 \cdot 10^7$. The fairly constant observed speedups of GWAS-1D-SMP over FAST-LMM and GENABEL at $m = 3.6 \cdot 10^7$ are, respectively, 6.3 and 56.8.

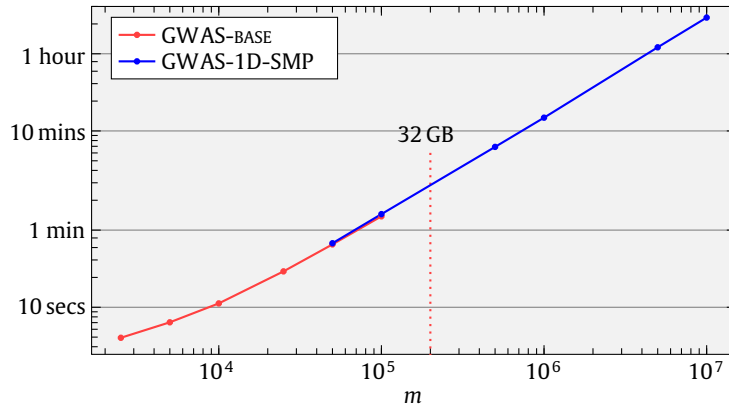


Fig. 1. Performance of single-trait solvers GWAS-BASE and GWAS-1D-SMP as a function of m . $n = 10,000, p = 4$, and m ranges from 10^3 to 10^7 . The vertical line indicates the limit for the in-core solver GWAS-BASE imposed by the RAM size.

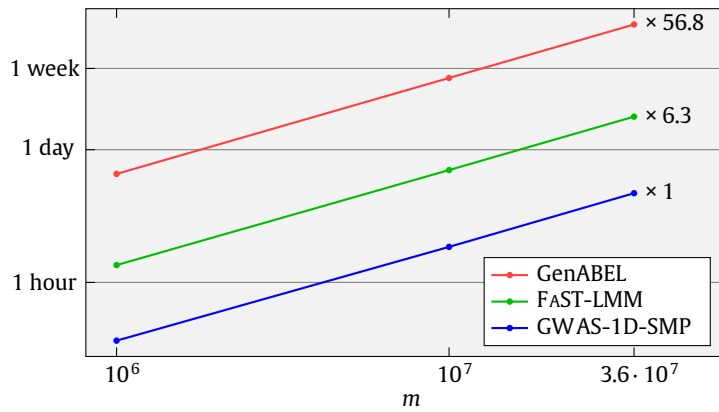


Fig. 2. Performance of the single-trait solver GWAS-1D-SMP compared to GENABEL and FAST-LMM. $n = 10,000, p = 4$, and m ranges from 10^6 to $3.6 \cdot 10^7$.




Fig. 3. Default 2D matrix distribution on a 2 × 3 process grid.

2.3. Distributed memory

While GWAS-1D-SMP scales up to an arbitrarily large amount of SNPs m , the main memory is still a limiting factor for the population size n : in fact, the algorithm necessitates the matrix $M \in \mathbb{R}^{n \times n}$ (or equivalently, its Cholesky factor L) to reside fully in memory. Due to the triangular solve (Algorithm 2, line 2), keeping the matrix in secondary storage is not a viable option. Our approach here consists in distributing M, L , and all matrices on which L operates across multiple compute nodes. Thereby, any constraint on their size is lifted.

2.3.1. ELEMENTAL

As a framework for distributed-memory dense linear algebra operations, we use ELEMENTAL [11]. This C++ library, which is based on the Message Passing Interface (MPI), operates on a virtual two-dimensional grid of processes; its name is inspired by the fact that, in general, matrices are cyclically distributed across this grid in an element-wise fashion. This principal distribution⁵ is shown in Fig. 3.

⁵ In ELEMENTAL's notation: $[MC, MR]$.

Algebraic operations on distributed matrices typically involve two stages: data redistribution (communication), and invocation of single-node BLAS or LAPACK routines (computation). Optimal performance is attained by minimizing communication within the redistributions. In most cases, as shown in [11], this is achieved by choosing the process grid to be as close to a perfect square as possible.

While a square process grid is optimal for performance, since all processes only hold non-contiguous portions of the matrix, it complicates loading contiguously stored data from files into a distributed matrix. In the context of GWAS, the algorithm has to load two objects of different nature: the matrix M , and the collections of vectors X_{blk} ; the special nature of the latter determines that the vectors can be loaded and processed in any order.

For loading M , we first read contiguous panels into the local memory of each process via standard file operations, and then, by accumulating these panels, construct the global (distributed) version of M . This is done via ELEMENTAL's axpy-interface, a feature that makes it possible to add node-local matrices to a global one.

For loading X_{blk} instead, a collection of contiguously stored vectors is read into memory through more efficient means than the axpy-interface by exploiting that, as long as consistently handled, the order of the vectors is irrelevant. The trick is to use a matrix that is distributed on a virtual 1D reordering of the grid into a row of processes. As shown in Fig. 4, the process-local data of such a matrix is a set of full columns, which can be loaded from a contiguous data-file. While these local columns are not adjacent in the distributed matrix, ELEMENTAL guarantees that all algebraic operations performed on them maintain their order. For performance reasons, prior to any computation, the matrix on the 1D ordering of this grid needs to be redistributed to conform to the initial 2D process grid (Fig. 3). This redistribution, provided by ELEMENTAL, can internally be performed most efficiently through a single `MPI_Alltoall` if the 1D grid is the concatenation of the rows of the 2D grid.⁶

2.3.2. The parallel algorithm

Algorithm 3. Distributed memory algorithm for single-trait studies. Asynchronous I/O operations are depicted in green, distributed matrices and operations in blue, and quantities that differ across processes in red

```

1  load_start first  $X_{blk}$ 
2   $LL^T := M$ 
3   $\bar{X}_L := L^{-1}X_L, \bar{y} := L^{-1}y$ 
4  copy  $\bar{X}_L := \bar{X}_L, \bar{y} := \bar{y}$ 
5   $S_{TL} := \bar{X}_L^T \bar{X}_L, \bar{b}_T := \bar{X}_L^T \bar{y}$ 
6  for each  $blk$ 
7    load_wait current  $X_{blk}$ 
8    if not last  $blk$ : load_start next  $X_{blk}$ 
9    set  $X_{blk} := \text{combine}(X_{blk})$ 
10    $\bar{X}_{blk} := L^{-1}X_{blk}$ 
11   set  $\bar{X}_{blk} := \text{localpart}(\bar{X}_{blk})$ 
12    $S_{blk} := \bar{X}_{blk}^T \bar{X}_L$ 
13   for  $i$  in  $\{1, \dots, \frac{m_{blk}}{n}\}_p$ 
14     set  $\bar{X}_{Ri} := \bar{X}_{blk}[i], S_{BLi} := S_{blk}[i]$ 
15      $S_{BRi} := \bar{X}_{Ri}^T \bar{X}_{Ri}$ 
16      $\bar{b}_{Bi} := \bar{X}_{Ri}^T \bar{y}$ 
17     set  $S_i := \left( \begin{array}{c|c} S_{TL} & * \\ \hline S_{BLi} & S_{BRi} \end{array} \right), \bar{b}_i := \left( \begin{array}{c} \bar{b}_T \\ \hline \bar{b}_{Bi} \end{array} \right)$ 
18      $b_i := S_i^{-1} \bar{b}_i$ 
19     set  $b_{blk}[i] := b_i$ 
20   end
21   if not first  $blk$ : store_wait previous  $b_{blk}$ 
22   store_start current  $b_{blk}$ 
23 end
24 store_wait last  $b_{blk}$ 

```

In Algorithm 3, we present the distributed-memory version of Algorithm 2 for np processes; the matrices distributed among the processes and the corresponding operations are highlighted in blue; the quantities that differ from one process to another are instead in red.

⁶ In ELEMENTAL: $[*, VR]$.




Fig. 4. 1D matrix distribution on a 1×6 process grid.

The algorithm begins (line 1) by loading the first $\frac{m_{blk}}{np}$ vectors X_{Ri} into a local block X_{blk} on each process asynchronously. Then, from the initially distributed M, X_L , and y , it computes L, \bar{X}_L , and \bar{y} (lines 2–3). Next, X_L and y , respectively local copies of X_L and y , are created on each process (line 4). Since small local computations are significantly more efficient than the distributed counterparts, all processes compute S_{TL} and b_T redundantly (line 5).

In order to compute $\bar{X}_{blk} := L^{-1}X_{blk}$, Elemental requires all involved operands to be distributed across its 2D process grid. However, the process-local X_{blk} is stored as contiguous columns. These matrices, which are seen as a cyclically distributed matrix on a 1D grid (see Section 2.3.1), are therefore redistributed to X_{blk} on the 2D grid (line 9). After the computation in line 10 completes, the resulting X_{blk} is distributed back: each process receives those contiguous columns \bar{X}_{blk} of \bar{X}_{blk} that correspond to X_{blk} .

In addition to blocking X_{Ri} and b_{Bi} , by stacking the m_{blk} row vectors S_{BLi} 's belonging to the current block into S_{blk} , their computation is combined into a single matrix product (line 12). In line 14, S_{BLi} and X_{Ri} are selected from, respectively, X_{blk} and S_{blk} for the innermost loop. This loop then computes the local b_{blk} independently on each process. Finally, while b_{blk} (whose columns b_i corresponds to the initially loaded vectors X_{Ri} within X_{blk}) is stored asynchronously, the next iteration commences.

In addition to Elemental's distributed memory parallelism, we exploit node-local shared memory parallelism in two different ways: since the innermost loop (lines 13 through 20) works on very small quantities, it is parallelized with OpenMP; all other operations involve larger matrices and make use of multithreaded BLAS libraries.

2.3.3. Performance results: GWAS-1D-MPI

We compile GWAS-1D-MPI, the C++-implementation of Algorithm 3, with the GNU C++ compiler (GCC version 4.8.1) with OpenMPI (version 1.6.4), use Elemental (version 0.82-p1) and link to Intel's Math Kernel Library (MKL version 11.0). In our tests, we use a compute cluster with 40 nodes, each equipped with 16 GB of RAM and two quad-core Intel Harpertown E5450 processors running at 3.00 Ghz. The nodes are connected via InfiniBand and access a high speed Lustre file system.⁷

Throughout all our experiments, we use one process per node with 8 threads each. Furthermore, we choose m_{blk} —the width of X_{blk} — as large as possible to fit in the combined main memory.

Processing huge numbers of SNPs out-of-core Since GWAS-1D-MPI incorporates the double-buffering technique introduced in Section 2.2, it can process datasets with arbitrarily large numbers of SNPs m without introducing any overhead due to I/O operations. To support this claim, we perform a series of experiments on $np = 1, 2, 4$, and 8 nodes (8, 16, 32, and 64 cores) to solve a problem of size $n = 30,000$ and $p = 4$ with increasing dataset size m . The performance of these experiments is presented in Fig. 5, where the vertical lines mark the points at which the 16 GB of RAM per node are insufficient to store all m vectors X_{Ri} . The plot shows a very smooth behavior with m (dominated by the triangular solve in Algorithm 3, line 10) well beyond this in-core memory limit.

Increasing the population size n We now turn to the main goal of GWAS-1D-MPI: performing computations on problems whose matrix $M \in \mathbb{R}^{n \times n}$ exceeds the capacity of the main memory. For this purpose, we use $m = 10,000$, $p = 4$ and execute GWAS-1D-MPI on $np = 1, 2, 4$, and 8 nodes (8, 16, 32, and 64 cores) with increasing matrix size n . The performance of these executions (Fig. 6) is dominated by the cubic complexity of the Cholesky factorization of M (Algorithm 3, line 2). The vertical lines indicate where the nodes' aggregate main memory would be exceeded by the combined size of the distributed M and the buffers for processing the X_{Ri} one by one. The plot shows that our implementation succeeds in overcoming these memory limitations through increasing the number of nodes.

3. Multi-trait GWAS

In an important class of GWAS (analysis of “omics” phenotypes), the studies involve many traits y_j [12–15]. In this case, the set of generalized least squares problems in Eq. (1) extends into the second dimension j :

$$b_{ij} := \left(X_i^T M_j^{-1} X_i \right)^{-1} X_i^T M_j^{-1} y_j, \quad \text{with } i = 1, \dots, m \quad \text{and } j = 1, \dots, t. \quad (4)$$

⁷ Due to the favorable ratio between computation and data movement, we recommend using algorithms GWAS-1D-MPI (Algorithm 3) and GWAS-2D-MPI (Algorithm 7, in Section 3) not only on traditional compute clusters, but on Cloud systems too.




Fig. 5. Performance of the single-trait solver GWAS-1D-MPI as a function of m . $n = 30,000$, $p = 4$, and m ranges from 10^3 to 10^7 . The vertical lines indicate the limits for in-core versions of the parallel algorithm imposed by the combined RAM sizes.




Fig. 6. Performance of the single-trait solver GWAS-1D-MPI as a function of n . $p = 4$, $m = 10,000$, and n ranges from 5,000 to 120,000. The vertical lines indicate the limits imposed by the combined RAM sizes.

This extra dimension is not only reflected in the traits y_j , but it also introduces varying matrices M_j . Such symmetric positive definite M_j 's share the common structure

$$M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I \right), \quad (5)$$

where the so called kinship matrix $\Phi \in \mathbb{R}^{n \times n}$ is independent of j . This structure plays a critical role in the design of fast algorithms for multi-trait GWAS.

3.1. The algorithm

In contrast to our single-trait algorithms, which are based on the Cholesky decomposition of M , the key to fast algorithms for the multi-trait case is the eigendecomposition of Φ :

$$1 \quad Z \Lambda Z^T := \Phi.$$

Here, Z and Λ are, respectively, the orthonormal collection of eigenvectors and the diagonal matrix of eigenvalues of Φ . Substituting this decomposition, Eq. (5) becomes

$$M_j = \sigma_j^2 \left(h_j^2 \underbrace{Z^T \Lambda Z}_{\Phi} + (1 - h_j^2) \underbrace{Z^T Z}_I \right) = Z^T \sigma_j^2 \left(h_j^2 \Lambda + (1 - h_j^2) I \right) Z.$$

This means that (1) the eigenvectors of M_j and Φ are the same; (2) the eigenvalues of M_j are those of Φ shifted and scaled:

$$2 \quad \Lambda_j := \sigma_j^2 (h_j^2 \Lambda + (1 - h_j^2) I)$$

As a result,⁸ M_j^{-1} can be expressed as $M_j^{-1} = (Z \Lambda_j Z^T)^{-1} = Z^{-T} \Lambda_j^{-1} Z^{-1} = Z \Lambda^{-1} Z^T$. Plugging this into Eq. (4) yields

$$\begin{aligned} b_{ij} &= (X_i^T Z \Lambda_j^{-1} Z^T X_i)^{-1} X_i^T Z \Lambda_j^{-1} Z^T y_j \\ &= ((Z^T X_i)^T \Lambda_j^{-1} (Z^T X_i))^{-1} (Z^T X_i)^T \Lambda_j^{-1} (Z^T y_j). \end{aligned}$$

This expression shows that the assignments

$$\begin{aligned} 3 \quad \tilde{X}_i &:= Z^T X_i && \text{(matrix times matrix)} \\ 4 \quad \tilde{y}_j &:= Z^T y_j && \text{(matrix times vector)} \end{aligned}$$

can, respectively, be computed independently of j and i . As a result, we have

$$b_{ij} = (\tilde{X}_i^T \Lambda_j^{-1} \tilde{X}_i)^{-1} \tilde{X}_i^T \Lambda_j^{-1} \tilde{y}_j.$$

By decomposing

$$5 \quad K_j K_j^T := \Lambda_j^{-1} \quad \text{(reciprocal square root of diagonal)}$$

and assigning

$$\begin{aligned} 6 \quad \bar{X}_{ij} &:= K_j^T \tilde{X}_i && \text{(many-vector scaling)} \\ 7 \quad \bar{y}_j &:= K_j^T \tilde{y}_j, && \text{(vector scaling)} \end{aligned}$$

the problem then reduces to

$$b_{ij} = (\bar{X}_{ij}^T \bar{X}_{ij})^{-1} \bar{X}_{ij}^T \bar{y}_j.$$

This ordinary least squares problem is of the same form as encountered in Section 2.1; hence it is solved in the same way.

Algorithm 4. Optimized algorithm for multi-trait studies

```

1  ZΛZT := Φ
2  X̃L := ZTXL
3  for i in {1, ..., m}
4    X̃Ri := ZTXRi
5  end
6  for j in {1, ..., t}
7    ỹj := ZTyj
8    Λj := σj2(hj2Λ + (1 - hj2)I)
9    KjKjT := Λj-1
10   X̄Lj := KjTX̃L, ȳj := KjTỹj
11   STLj := X̄LjTX̄Lj, b̄Tj := X̄LjTȳj
12   for i in {1, ..., m}
13     X̄Rij := KjTX̃Ri
14     SBLij := X̄RijTX̄Lj, SBRij := X̄RijTX̄Rij
15     b̄Bij := X̄RijTȳj
16     set Sij := ( STLj | *
                  SBLij | SBRij ), b̄ij := ( b̄Tj
                                          b̄Bij )
17     bij := Sij-1b̄ij
18   end
19 end
```

⁸ Using the orthonormality of Z : $Z^{-1} = Z^T$.

As in the 1D case, we take advantage of $X_i = (X_L|X_{Ri})$ and propagate this structure to $\tilde{X}_i = (\tilde{X}_L|\tilde{X}_{Ri})$, $\bar{X}_{ij} = (\bar{X}_{Lj}|\bar{X}_{Rij})$, $\bar{S}_{ij} = \begin{pmatrix} \bar{S}_{TLj} & * \\ \bar{S}_{BLij} & \bar{S}_{BRij} \end{pmatrix}$, and $\bar{b}_{ij} = \begin{pmatrix} \bar{b}_{Tj} \\ \bar{b}_{Bij} \end{pmatrix}$. Extracting all objects independent of the indices i and j from the corresponding loops, we obtain the mathematically optimized Algorithm 4. This optimization reduces the complexity of the algorithm from $O(n^3 + n^2(m+t)p + mtnp^2)$ to $O(n^3 + n^2(m+t) + mtnp)$.

3.2. Out-of-core

Algorithm 5. Out-of-core algorithm for multi-trait studies. The y_j , X_{Ri} , and b_{ij} are streamed from and to disk in blocks. Asynchronous I/O operations are highlighted in green. (Function innerloops is given in Algorithm 6)

```

1   $Z\Lambda Z^T := \Phi$ 
2   $\tilde{X}_L := Z^T X_L$ 
3  for each  $blk_m$  (asynchronously loading  $X_{blk_m}$ )
4     $\tilde{X}_{blk_m} := Z^T X_{blk_m}$ 
5  end (asynchronously storing  $\tilde{X}_{blk_m}$ )
6  for each  $blk_t$  (asynchronously loading  $y_{blk_t}$ )
7     $\tilde{y}_{blk_t} := Z^T y_{blk_t}$ 
8  end (asynchronously storing  $\tilde{y}_{blk_t}$ )
9  for each  $tile_t$  (asynchronously loading  $\tilde{y}_{tile_t}$ )
10   for each  $tile_m$  (asynchronously loading  $\tilde{X}_{tile_m}$ )
11      $b_{tile} := \text{innerloops}(\tilde{X}_{tile_m}, \tilde{y}_{tile_t})$ 
12   end (asynchronously storing  $b_{tile}$ )
13 end

```

Algorithm 6. Function innerloops computes a tile of b_{ij} 's from corresponding tiles of \tilde{X}_{Ri} 's and \tilde{y}_j 's

```

1  function innerloops( $X_{tile}, \tilde{y}_{tile}$ )
2    for j in  $\{1, \dots, \text{width}(\tilde{y}_{tile})\}$ 
3      set  $\tilde{y}_j := \tilde{y}_{tile}[j]$ 
4       $\Lambda_j := \sigma_j^2(h_j^2\Lambda + (1-h_j^2)I)$ 
5       $K_j K_j^T := \Lambda_j^{-1}$ 
6       $\bar{X}_{Lj} := K_j^T \tilde{X}_L, \bar{y}_j := K_j^T \tilde{y}_j$ 
7       $S_{TLj} := \bar{X}_{Lj}^T \bar{X}_{Lj}, \bar{b}_{Tj} := \bar{X}_{Lj}^T \bar{y}_j$ 
8      for i in  $\{1, \dots, \text{width}(\tilde{X}_{tile})\}$ 
9        set  $\bar{X}_{Rj} := \tilde{X}_{tile}[i]$ 
10        $\bar{X}_{Rij} := K_j^T \bar{X}_{Ri}$ 
11        $S_{BLij} := \bar{X}_{Rij}^T \bar{X}_{Lj}, S_{BRij} := \bar{X}_{Rij}^T \bar{X}_{Rij}$ 
12        $\bar{b}_{Bij} := \bar{X}_{Rij}^T \bar{y}_j$ 
13       set  $S_{ij} := \begin{pmatrix} S_{TLj} & * \\ S_{BLij} & S_{BRij} \end{pmatrix}, \bar{b}_{ij} := \begin{pmatrix} \bar{b}_{Tj} \\ \bar{b}_{Bij} \end{pmatrix}$ 
14        $b_{ij} := S_{ij}^{-1} \bar{b}_{ij}$ 
15       set  $b_{tile}[i, j] := b_{ij}$ 
16     end
17   end
18   return  $b_{tile}$ 
19 end

```

To allow processing of arbitrarily large numbers of SNPs m and traits t , we make use of double-buffering mechanisms equivalent to those discussed in Section 2.2, leading to Algorithm 5: In lines 3 through 5, the matrix–vector products $\tilde{X}_{Ri} := Z^T X_{Ri}$ are combined into far more efficient matrix–matrix products on blocks X_{blk_m} of vectors X_{Ri} . While one \tilde{X}_{blk_m} is computed, the previous \tilde{X}_{blk_m} and the next X_{blk_m} are, respectively, stored and loaded simultaneously. Subsequently (lines 6

– 8), the same mechanism is used to compute $\tilde{y}_j := Z^T y_j$ in blocks y_{blk_t} . This process results in two temporary files containing, respectively, all \tilde{X}_{Ri} 's and \tilde{y}_j 's. These files are of the same size as the inputs X_{Ri} and y_j , i.e., $n \cdot m$ and $n \cdot t$ doubles.

The main loops of the algorithm (lines 9 – 13) employ blocking and double-buffering along both m and t . Thereby, the result is computed in tiles b_{tile} of vectors b_{ij} . While one of these tiles is computed, both the next set of vectors \tilde{X}_{Ri} (and \tilde{y}_j) is loaded in blocks \tilde{X}_{tile_m} (and \tilde{y}_{tile_t}) and the previous b_{tile} is stored asynchronously.

In total, the algorithm involves four blocking factors corresponding to blk_t , blk_m , $tile_t$, and $tile_m$. In order to make the matrix products involving y_{blk_t} and X_{blk_m} as efficient as possible, it comes naturally to choose their sizes large. On the other hand, in order to maximize the computation per I/O ratio, the widths of \tilde{y}_{tile_t} and \tilde{X}_{tile_m} should be chosen such that b_{tile} is roughly square.

A highly efficient shared memory implementation of Algorithm 5 is presented in [10]; it is shown to be several orders of magnitude faster than comparable software packages.

3.3. Distributed memory

Due to the size of the main memory, the kinship matrix $\Phi \in \mathbb{R}^{n \times n}$ limits the multi-trait shared memory implementation. (Very much as the single-trait shared memory implementation was limited by the covariance matrix $M \in \mathbb{R}^{n \times n}$.) To overcome this limitation, we present an ELEMENTAL-based distributed memory solution: Algorithm 7.

Algorithm 7. Distributed memory algorithm for multi-trait studies. Asynchronous I/O operations are depicted in green, distributed matrices and operations in blue, and quantities that differ across processes in red. (Function innerloops is given in Algorithm 6)

```

1   $Z \Lambda Z^T := \Phi$ 
2   $\tilde{X}_L := Z^T X_L$ 
3  copy  $\tilde{X}_L := \tilde{X}_L$ 
4  for each  $blk_t$  (asynchronously loading  $y_{blk_t}$ )
5    set  $y_{blk_t} := \text{combine}(y_{blk_t})$ 
6     $\tilde{y}_{blk_t} := Z^T y_{blk_t}$ 
7    set  $\tilde{y}_{blk_t} := \text{localpart}(\tilde{y}_{blk_t})$ 
8  end (asynchronously storing  $\tilde{y}_{blk_t}$ )
9  for each  $blk_m$  (asynchronously loading  $X_{blk_m}$ )
10   set  $X_{blk_m} := \text{combine}(X_{blk_m})$ 
11    $\tilde{X}_{blk_m} := Z^T X_{blk_m}$ 
12   set  $\tilde{X}_{blk_m} := \text{localpart}(\tilde{X}_{blk_m})$ 
13 end (asynchronously storing  $\tilde{X}_{blk_m}$ )
14 for each  $tile_t$  (asynchronously loading  $\tilde{y}_{tile_t}$ )
15   for each  $tile_m$  (asynchronously loading  $\tilde{X}_{tile_m}$ )
16      $b_{tile} := \text{innerloops}(\tilde{X}_{tile_m}, \tilde{y}_{tile_t})$ 
17   end (asynchronously storing  $b_{tile}$ )
18 end

```

To overcome the aforementioned limitation, this algorithm distributes Φ and its eigenvectors Z across multiple processes. Consequently, applying the same technique used for X_{blk} in the single-trait Algorithm 3, the blocks of vectors X_{blk_m} and y_{blk_t} –to which Z is applied– are also distributed.

The remainder of Algorithm 7 (lines 14 onward) does not involve any large matrices that necessitate distributing. Hence, since each process can work on separate tiles b_{tile} , this part of the algorithm is embarrassingly parallel.

3.4. Performance results: GWAS-2D-MPI

The performance experiments for GWAS-2D-MPI, i.e., the implementation of Algorithm 7, are carried out with the same setup used for GWAS-1D-MPI (Section 2.3.3). Again, the implementation makes use of Elemental—specifically of the EleMRRR eigensolver [16]—and of Intel's MKL library.

GWAS-2D-MPI vs. GWAS-1D-MPI. In the scenario of multi-trait studies ($t > 1$), the main advantage of GWAS-2D-MPI over multiple runs of GWAS-1D-MPI is that redundant calculations are avoided. Complexity-wise, the difference is apparent: $O(n^3 + mn^2t)$ vs. $O(n^3t + mn^2t)$, respectively, for GWAS-2D-MPI and GWAS-1D-MPI. The following experiment is designed to illustrate such a gap. Fixing $n = 30,000$, $p = 4$, $m = 10,000$, and increasing t from 1 to 100, we compare the runtime for GWAS-2D-MPI with that for t independent runs of GWAS-1D-MPI. As Fig. 7 shows, regardless of the number of cores used, GWAS-2D-MPI offers the best timings for $t \geq 35$. Most importantly, the difference in slope indicates that the difference between the approaches will grow larger as t increases. Indeed, at $t = 5000$ GWAS-2D-MPI outperforms t executions of GWAS-1D-MPI by more than two orders of magnitude.

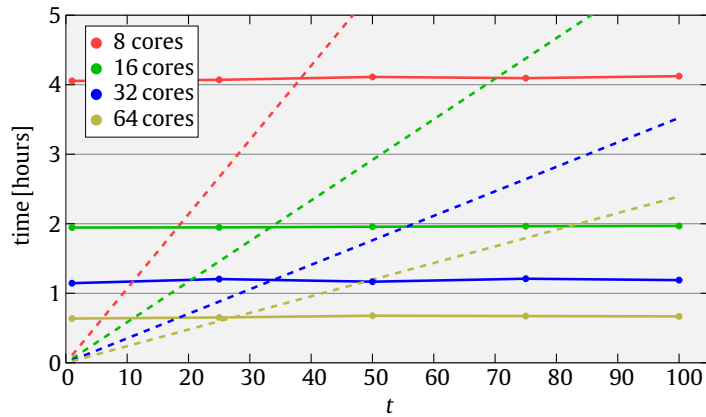


Fig. 7. Performance of the multi-trait solver GWAS-2D-MPI (—) compared to t runs of single-trait solver GWAS-1D-MPI (- - -) as a function of t . $n = 30,000, p = 4, m = 10,000$, and t ranges from 1 to 100.

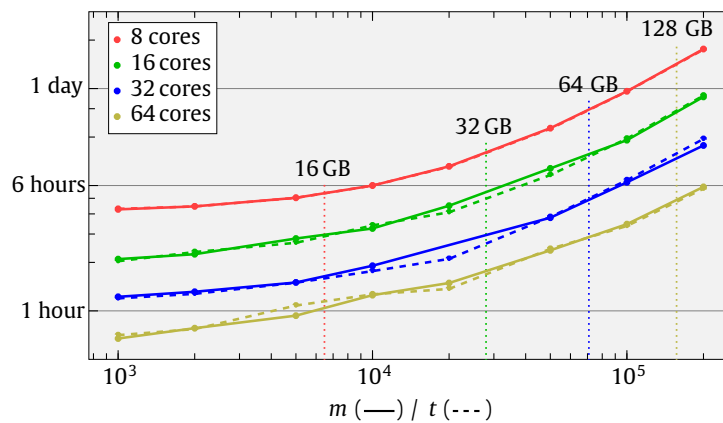


Fig. 8. Performance of the multi-trait solver GWAS-2D-MPI as a function of m and t . $n = 30,000, p = 4$, and, while either m or t is fixed at 10,000, the other ranges from 1,000 to 200,000. The vertical lines indicate the limits for in-core versions of the parallel algorithm imposed by the combined RAM sizes. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)




Fig. 9. Performance of the multi-trait solver GWAS-2D-MPI as a function of n . $p = 4, m = t = 10,000$, and n ranges from 5000 to 85,000. The vertical lines indicate the limits imposed by the combined RAM sizes.

Large m and t Since the aforementioned double-buffering technique is applied to both the inputs (m SNPs X_{Ri} and t traits y_j) and the output ($m \cdot t$ vectors b_{ij}), the application of GWAS-2D-MPI is not constrained by either the number of SNPs or the number of traits. This is shown in Fig. 8: In this experiment, with $n = 30,000$ and $p = 4$ constant, one of m and t is fixed to 10,000, and the other varies between 1000 and 200,000. The plot provides evidence that GWAS-2D-MPI solves, without

penalty, GWAS that do not fit either in ($np = 1$) the local memory of one node (dashed vertical red line), or ($np > 1$) the aggregate memory from multiple nodes (green, blue, and yellow vertical lines).

In Fig. 8, one should also observe that the timings for the experiments with varying m , and for those with varying t are very similar. The reason is that the execution time of Algorithm 7 is dominated by the matrix–matrix multiplications $\tilde{X}_{blk_m} := Z^T X_{blk_m}$ and $\tilde{y}_{blk_t} := Z^T y_{blk_t}$ (Algorithm 7, lines 6 and 11), and the complexity of these operations $-O(n^2m + n^2t)$ – is perfectly symmetric with respect to m and t .

Large n . As the population size n increases, the quadratically growing memory requirement for the kinship matrix $\Phi \in \mathbb{R}^{n \times n}$ quickly surpass the memory available in a single node. To show that GWAS-2D-MPI overcomes this limitation, we fix $m = t = 10,000$ and $p = 4$, and let n grow from 5,000 to 85,000. The resulting execution times in Fig. 9 clearly follow a smooth cubic behavior in n .⁹ As for the single-trait case (Section 2.3, Fig. 6), the plot shows that the limit imposed by the combined main memory size (dashed vertical lines) can be overcome by increasing the number of nodes.

4. Conclusion

We presented parallel algorithms for the computation of linear mixed-models based genome-wide association studies (GWAS). They address the issue of growing dataset sizes due to the number of studied polymorphisms m , the population size n , and/or the number of traits t .

The first algorithm uses a double-buffering technique in order to process datasets with arbitrarily large numbers of genetic polymorphisms. Compared to other wide-spread GWAS-codes, our shared memory implementation, GWAS-1D-SMP, was shown to be at least one order of magnitude faster.

The second algorithm enables the processing of datasets involving large populations by storing the relationship matrix in the combined main memory of distributed memory architectures. GWAS-1D-MPI, the implementation of this algorithm, was shown to scale in both the population size and the number of processes used.

The third algorithm extends the second by analyzing arbitrary numbers of traits at once avoiding redundant computation. This reflects in the performance of our implementation GWAS-2D-MPI, which scales in all problem sizes, and is significantly faster than multiple runs of GWAS-1D-MPI.

Together, these algorithms form a viable basis for the challenges posed by the scale of current and future genome-wide association studies.

Acknowledgments

Financial support from the Deutsche Forschungsgemeinschaft (German Research Association) through Grant GSC 111 and Deutsche Telekom Stiftung is gratefully acknowledged. The authors thank Yurii Aulchenko for fruitful discussions on the biological background of GWAS.

References

- [1] L.A. Hindorf, P. Sethupathy, H.A. Junkins, E.M. Ramos, J.P. Mehta, F.S. Collins, T.A. Manolio, Potential etiologic and functional implications of genome-wide association loci for human diseases and traits, *Proc. Natl. Acad. Sci. U.S.A.* 106 (23) (2009) 9362–9367.
- [2] E. Boerwinkle, R. Chakraborty, C.F. Sing, The use of measured genotype information in the analysis of quantitative phenotypes in man. I. Models and analytical methods, *Ann. Hum. Genet.* 50 (Pt 2) (1986) 181–194.
- [3] J. Yu, G. Pressoir, W.H. Briggs, I. Vroh Bi, M. Yamasaki, J.F. Doebley, M.D. McMullen, B.S. Gaut, D.M. Nielsen, J.B. Holland, S. Kresovich, E.S. Buckler, A unified mixed-model method for association mapping that accounts for multiple levels of relatedness, *Nat. Genet.* 38 (2) (2006) 203–208.
- [4] R.J. Klein, C. Zeiss, E.Y. Chew, J.-Y. Tsai, R.S. Sackler, C. Haynes, A.K. Henning, J.P. SanGiovanni, S.M. Mane, S.T. Mayne, M.B. Bracken, F.L. Ferris, J. Ott, C. Barnstable, J. Hoh, Complement factor h polymorphism in age-related macular degeneration, *Science* 308 (5720) (2005) 385–389.
- [5] L.A. Hindorf, J. MacArthur, J. Morales, H.A. Junkins, P.N. Hall, A.K. Klemm, T.A. Manolio, A catalog of published genome-wide association studies. <<http://www.genome.gov/gwastudies/>> (accessed 03.2013).
- [6] T.A. Manolio, Published gwas reports, 2005–6/2012. <<http://www.genome.gov/multimedia/illustrations/PublishedGWAReports6-2012.pdf>>.
- [7] Y.S. Aulchenko, S. Ripke, A. Isaacs, C.M. van Duijn, GenABEL: an R library for genome-wide association analysis, *Bioinformatics* 23 (10) (2007) 1294–1296.
- [8] C. Lippert, J. Listgarten, Y. Liu, C.M. Kadie, R.I. Davidson, D. Heckerman, Fast linear mixed models for genome-wide association studies, *Nat. Methods* 8 (10) (2011) 833–835.
- [9] D. Fabregat-Traver, Y.S. Aulchenko, P. Bientinesi, Solving sequences of generalized least-squares problems on multi-threaded architectures, *Appl. Math. Comput. (AMC)* 234 (2014) 606–617.
- [10] D. Fabregat-Traver, P. Bientinesi, Computing petaflops over terabytes of data: the case of genome-wide association studies, *ACM Trans. Math. Softw.* 40 (4) (2014) 1–27. Article No: 27.
- [11] J. Poulson, B. Marker, R.A. van de Geijn, J.R. Hammond, N.A. Romero, Elemental: a new framework for distributed memory dense matrix computations, *ACM Trans. Math. Software* 39 (2) (2013) 13:1–13:24.
- [12] C. Gieger et al, Genetics meets metabolomics: a genome-wide association study of metabolite profiles in human serum, *PLoS Genet.* 4 (11) (2008) e1000282+.
- [13] A.A. Hicks et al, Genetic determinants of circulating sphingolipid concentrations in european populations, *PLoS Genet.* 5 (10) (2009) e1000672.
- [14] G. Lauc et al, Genomics meets glycomics: the first GWAS study of human N-glycome identifies HNF1 α as a master regulator of plasma protein fucosylation, *PLoS Genet.* 6 (12) (2010) e1001256.
- [15] G. Lauc et al, Loci associated with N-glycosylation of human immunoglobulin g show pleiotropy with autoimmune diseases and haematological cancers, *PLoS Genet.* 9 (1) (2013) e1003225.
- [16] M. Petschow, E. Peise, P. Bientinesi, High-performance solvers for dense hermitian eigenproblems, *SIAM J. Sci. Comput.* 35 (1) (2013) C1–C22.

⁹ Due to the complexity of the eigenvalue decomposition.