Can Numerical Linear Algebra make it in Nature?

Advice on collaborations with computational biologists

Paolo Bientinesi
in collaboration with Diego Fabregat, Elmar Peise and Yurii Aulchenko

AICES, RWTH Aachen
pauldj@aices.rwth-aachen.de

June 9, 2014
Householder Symposium XIX
Spa, Belgium

The problem: **Genome-Wide Association Studies**

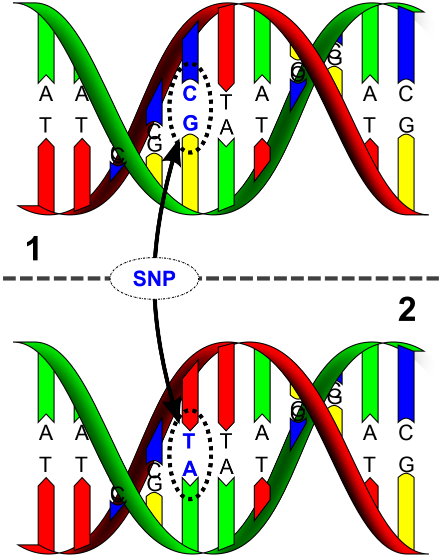

Source: David Hall

GWAS:
Correlation between a difference in the genome sequence (SNP) and a difference in the phenotype (observations)

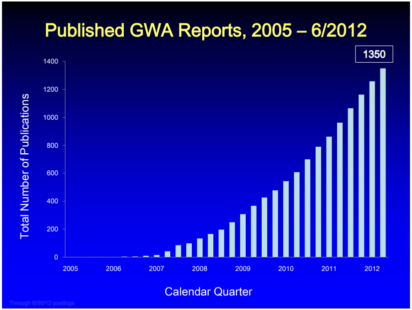

Source: Teri Manolio

The method

- “Mixed models”

The method

- "Mixed models"
- Linear regression with non-independent outcomes

The method

- "Mixed models"
- Linear regression with non-independent outcomes
- Generalized least-square problems

The method

- “Mixed models”
- Linear regression with non-independent outcomes
- Generalized least-square problems

$$b := \left(X^T M^{-1} X\right)^{-1} X^T M^{-1} y$$

- y: **phenotype** (outcome; vector of observations)

  E.g.: height, blood pressure for a set of people

- X: **genome measurements** and covariates (design matrix; predictors)

  E.g.: sex and age over height

- M: **dependencies** between observations

  E.g.: tall parents have tall children

- b: **relation** between a variation in the outcome (y) and a variation in the genome sequence (X)

$$b := \left(X^T M^{-1} X \right)^{-1} X^T M^{-1} y$$

- $X \in \mathcal{R}^{n \times p}$ “SNP”
- $y \in \mathcal{R}^{n}$ “trait”
- $b \in \mathcal{R}^{p}$ “genetic effect”
- $M \in \mathcal{R}^{n \times n}$ “covariance matrix”

- $n \approx 1{,}000 - 50{,}000$
- $p \in [1, ..., 20]$
- M: SPD

(Wrong) Problem definition

Isolated problem instances

$$b := \left(X^T M^{-1} X\right)^{-1} X^T M^{-1} y$$

"to be repeated millions of times"

(Wrong) Problem definition

Isolated problem instances

$$b := \left(X^T M^{-1} X\right)^{-1} X^T M^{-1} y$$

"to be repeated millions of times"

$$\Downarrow$$

for $i = 1, \ldots, m$ $\qquad m \approx 10^6 - 10^7$

$$b_i := \left(X_i^T M_i^{-1} X_i\right)^{-1} X_i^T M_i^{-1} y_i$$

(Wrong) Problem definition

Isolated problem instances

$$b := \left(X^T M^{-1} X\right)^{-1} X^T M^{-1} y$$

"to be repeated millions of times"

$$\Downarrow$$

for $i = 1, \ldots, m$ $\qquad m \approx 10^6 - 10^7$

$$b_i := \left(X_i^T M_i^{-1} X_i\right)^{-1} X_i^T M_i^{-1} y_i$$

$$\Downarrow$$

for $i = 1, \ldots, m$

$L_i L_i^T = M_i$ $\qquad$ `CHOL`

$X_i := L_i^{-1} X_i$ $\qquad$ `TRSM`

$y_i := L_i^{-1} y_i$ $\qquad$ `TRSV`

$b_i := \texttt{OLS}(X_i, y_i)$ $\qquad O(n^3 m)$

(Wrong) Problem definition

Isolated problem instances

$$b := \left(X^T M^{-1} X\right)^{-1} X^T M^{-1} y$$

"to be repeated millions of times"

$$\Downarrow$$

for $i = 1, \ldots, m$ $\qquad m \approx 10^6 - 10^7$

$$b_i := \left(X_i^T M_i^{-1} X_i\right)^{-1} X_i^T M_i^{-1} y_i$$

$$\Downarrow$$

for $i = 1, \ldots, m$

$L_i L_i^T = M_i$	CHOL
$X_i := L_i^{-1} X_i$	TRSM
$y_i := L_i^{-1} y_i$	TRSV
$b_i :=$ OLS(X_i, y_i)	$O(n^3 m)$

Problem definition #1

One-dimensional sequence of GLS problems

for $i = 1, \ldots, m$ $\qquad m \approx 10^6 - 10^7$

$$b_i := \left(X_i^T M^{-1} X_i\right)^{-1} X_i^T M^{-1} y$$

and

$X_i = [X_L | X_{Ri}],$ where $X_{Ri} \in \mathcal{R}^{n \times 1}$

Problem definition #1

One-dimensional sequence of GLS problems

for $i = 1, \dots, m$ $\qquad m \approx 10^6 - 10^7$

$$b_i := \left(X_i^T M^{-1} X_i\right)^{-1} X_i^T M^{-1} y$$

and

$X_i = [X_L | X_{Ri}]$, where $X_{Ri} \in \mathcal{R}^{n \times 1}$

$$\Downarrow$$

$$b_i := \left(\left[\frac{X_L^T}{X_{Ri}^T} \right] M^{-1} [X_L | X_{Ri}] \right)^{-1} \left[\frac{X_L^T}{X_{Ri}^T} \right] M^{-1} y$$

$$b_i := \left[\begin{array}{c|c} \overline{X}_L^T \overline{X}_L & \star \\ \hline \overline{X}_{Ri}^T \overline{X}_L & \overline{X}_{Ri}^T \overline{X}_{Ri} \end{array} \right]^{-1} \left[\frac{\overline{X}_L^T}{\overline{X}_{Ri}^T} \right] L^{-1} y$$

Problem definition #1

One-dimensional sequence of GLS problems

for $i = 1, \ldots, m \qquad m \approx 10^6 - 10^7$

$$b_i := \left(X_i^T M^{-1} X_i\right)^{-1} X_i^T M^{-1} y$$

and

$X_i = [X_L | X_{Ri}]$, where $X_{Ri} \in \mathcal{R}^{n \times 1}$

$$\Downarrow$$

$LL^T = M;$
$\overline{X}_L := L^{-1} X_L;$
$\overline{y} := L^{-1} y;$
$S_{TL} := \overline{X}_L^T \overline{X}_L;$
$\overline{b}_T := \overline{X}_L^T y;$

for $i = 1, \ldots, m$
$\quad \overline{X}_{Ri} := L^{-1} X_{Ri}$ TRSV
$\quad S_{BLi} := X_{Ri}^T \overline{X}_L$
$\quad S_{BRi} := \overline{X}_{Ri}^T \overline{X}_{Ri}$
$\quad \vdots$
$\quad b_i := S_i^{-1} \overline{b}_i \qquad O(n^2 m)$

Problem definition #2

Two-dimensional sequence of GLS problems

$$b_{ij} := \left({X_i}^T {M_j}^{-1} X_i\right)^{-1} {X_i}^T {M_j}^{-1} y_j$$

$$\Downarrow$$

for $i = 1, \ldots, m$

for $j = 1, \ldots, t$

$$b_{ij} := \left(X_i^T M_j^{-1} X_i\right)^{-1} X_i^T M_j^{-1} y_j$$

and

$X_i = [X_L | X_{Ri}]$, with $X_{Ri} \in \mathcal{R}^{n \times 1}$

and SPD(M_j)

$m \approx 10^6 - 10^7$; $\quad t = 1$ or $\approx 10^3 - 10^5$

Overview of the full problem

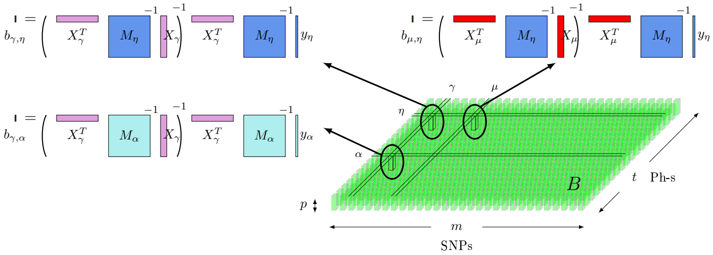

Overview of the full problem

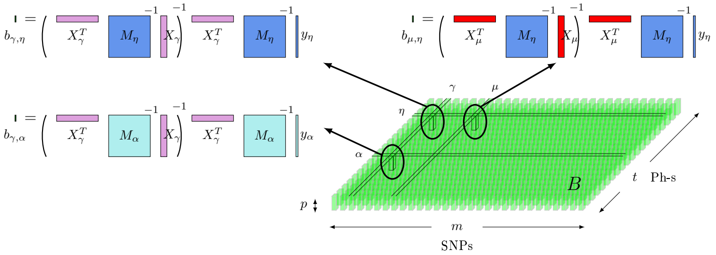

Problem size

$M \in \mathbb{R}^{n \times n}$	$1000 \leq n \leq 100k$	7.5MBs – 74.5GBs
$X_{Ri}, y_j \in \mathbb{R}^n$		8 – 780KBs
$b_{ij} \in \mathbb{R}^p$	$3 \leq p \leq 20$	24 – 160 Bytes
Total	$m \leq 10^8,\ t \leq 10^5$	**1.5 – 100s Terabytes**

$$\begin{cases} b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j \\ M_j = \ ? \end{cases}$$

$$\begin{cases} b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j \\ M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I \right) \end{cases}$$

$$
\begin{cases} b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j \\ M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I \right) \end{cases}
$$

$$
\Phi = Q \Lambda Q^T
$$

$$
\Rightarrow M_j = Q \left(\alpha_j \Lambda + \beta_j I \right) Q^T
$$

$$\begin{cases} b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j \\ M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I \right) \end{cases}$$

$$\Phi = Q \Lambda Q^T$$

$$\Rightarrow M_j = Q \, (\alpha_j \Lambda + \beta_j I) \, Q^T$$

$$\Rightarrow M_j^{-1} = Q \, (\alpha_j \Lambda + \beta_j I)^{-1} \, Q^T$$

$$
\begin{cases} b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j \\ M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I \right) \end{cases}
$$

$$
\Phi = Q \Lambda Q^T
$$

$$
\Rightarrow M_j = Q \left(\alpha_j \Lambda + \beta_j I \right) Q^T
$$

$$
\Rightarrow M_j^{-1} = Q \left(\alpha_j \Lambda + \beta_j I \right)^{-1} Q^T
$$

$$
b_{ij} :=
$$

$$
(X_i^T Q D_j^{-1} Q^T X_i)^{-1} X_i^T Q D_j^{-1} Q^T y_j
$$

$$\begin{cases} b_{ij} := (X_i^T M_j^{-1} X_i)^{-1} X_i^T M_j^{-1} y_j \\ M_j = \sigma_j^2 \left(h_j^2 \Phi + (1 - h_j^2) I\right) \end{cases}$$

$$\Phi = Q \Lambda Q^T$$

$$\Rightarrow M_j = Q \, (\alpha_j \Lambda + \beta_j I) \, Q^T$$

$$\Rightarrow M_j^{-1} = Q \, (\alpha_j \Lambda + \beta_j I)^{-1} \, Q^T$$

$b_{ij} :=$
$(X_i^T Q D_j^{-1} Q^T X_i)^{-1} X_i^T Q D_j^{-1} Q^T y_j$

$Q \Lambda Q^T = \Phi$
for 1 $\leq$ i $\leq$ m
    $X_i' := Q^T X_i$ GEMM
for 1 $\leq$ j $\leq$ t
    $y_j' := Q^T y_j$ GEMV
for 1 $\leq$ j $\leq$ t
    $D_j := \sigma_j^2 (h_j^2 \Lambda + (1 - h_j^2) I)$
    $K_j K_j^T = D_j^{-\frac{1}{2}}$
    $v_j := K_j^T y_j'$
    for 1 $\leq$ i $\leq$ m
        $W_{ij} := K_j^T X_i'$
        $S_{ij} := W_{ij}^T W_{ij}$
        $b_{ij} := W_{ij}^T v_j$
        $b_{ij} := S_{ij}^{-1} b_{ij}$

Cost: $O(n^2 mt)$ vs. $O(nmt)$

(Big) Data management

operands

$X's$	input	100s GBs – 2 TBs	streaming from disk
$y's$	input	1 – 10 GBs	streaming from disk
M	input	MBs – 80 GBs	read once
$b's$	output	100s MBs or 10s TBs	streaming to disk

Tiling: tb, mb?

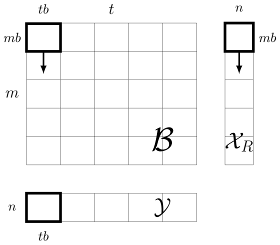

Parallelism

Does M fit in memory?

Parallelism

Does M fit in memory?

- YES $\Rightarrow$ single node + multithreading
  streaming HD$\leftrightarrow$CPU, double buffering, in-core implementation

Parallelism

Does M fit in memory?

- YES $\Rightarrow$ single node + multithreading
  streaming HD$\leftrightarrow$CPU, double buffering, in-core implementation

  ### Does M fit in GPU-memory?

  - Yes $\Rightarrow$ accelerator
    streaming HD$\leftrightarrow$CPU$\leftrightarrow$GPU, triple+double buffering, CPU+GPU implementation

Parallelism

Does M fit in memory?

- YES $\Rightarrow$ single node + multithreading
  streaming HD$\leftrightarrow$CPU, double buffering, in-core implementation

  Does M fit in GPU-memory?

  - Yes $\Rightarrow$ accelerator
    streaming HD$\leftrightarrow$CPU$\leftrightarrow$GPU, triple+double buffering, CPU+GPU implementation

- NO $\Rightarrow$ distributed memory + hybrid parallelism
  partitioning + streaming HD$\leftrightarrow$CPUs, double buffering, data distribution

Results $t = 1$

Single-trait analysis: one-dimensional sequence

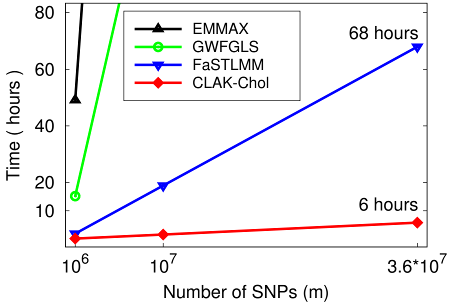

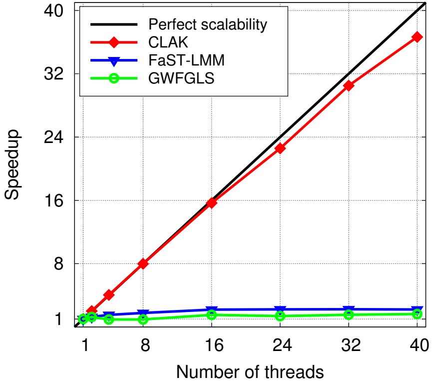

Single node

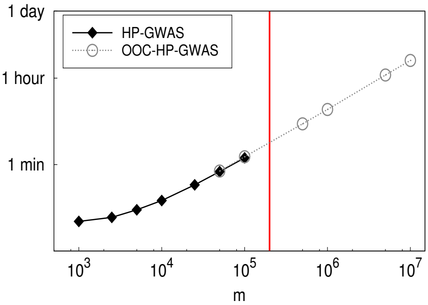

MPI

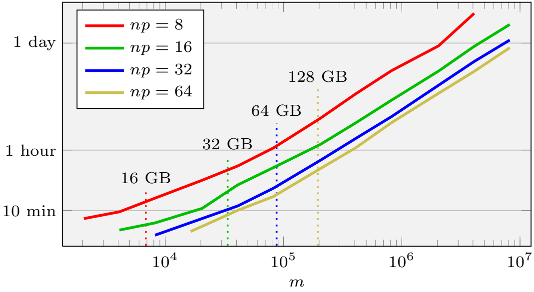

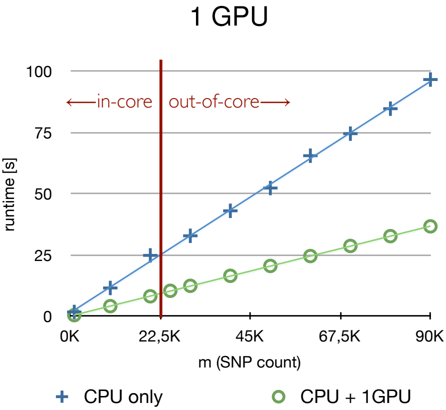
1 GPU
100
75
50
25
0
runtime [s]
←in-core
out-of-core→
0K
22,5K
45K
67,5K
90K
m (SNP count)
CPU only
CPU + 1GPU

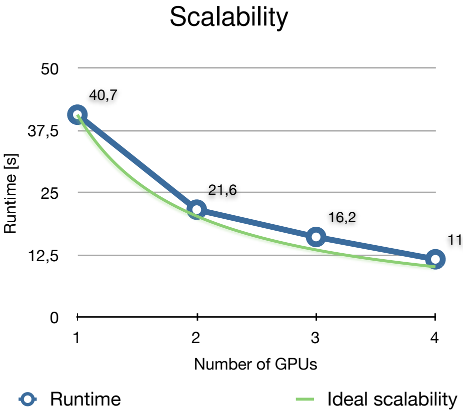
Scalability
50
37,5
25
12,5
0
Runtime [s]
40,7
21,6
16,2
11
1
2
3
4
Number of GPUs
Runtime
Ideal scalability

Multi-trait analysis: two-dimensional sequence

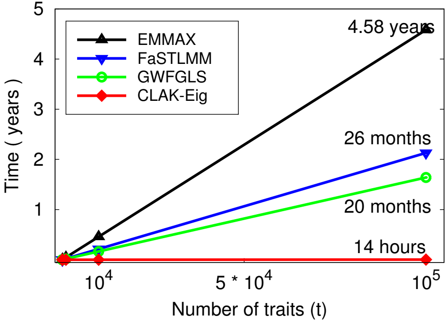

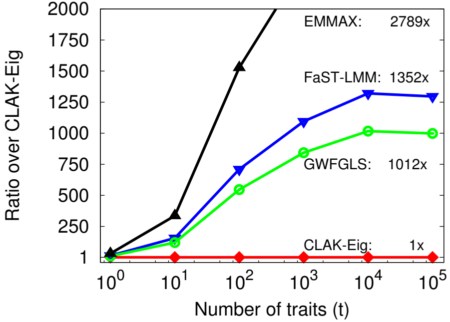

Conclusions

"Can Numerical Linear Algebra make it in Nature?"

. . .

Conclusions

“Can Numerical Linear Algebra make it in Nature?”

almost

Conclusions

"Can Numerical Linear Algebra make it in Nature?"

almost

- Multi-trait analysis:
  20 months $\rightarrow$ 14hours
- 1000x speedups
- No need for supercomputers
- Elegant solution: knowledge

Conclusions

"Can Numerical Linear Algebra make it in Nature?"

almost

- Multi-trait analysis:
  20 months $\rightarrow$ 14hours
- 1000x speedups
- No need for supercomputers
- Elegant solution: knowledge

- "Real" data vs. random data
- "Algorithms" vs. "methods"
- (Too) new methodology
- Size of readership
- New findings

Conclusions

"Can Numerical Linear Algebra make it in Nature?"

almost

- Multi-trait analysis:
  20 months $\rightarrow$ 14hours
- 1000x speedups
- No need for supercomputers
- Elegant solution: knowledge

- "Real" data vs. random data
- "Algorithms" vs. "methods"
- (Too) new methodology
- Size of readership
- New findings

- D.F. and P.B., "Computing Petaflops over Terabytes of Data: The Case of GWAS", TOMS 2014.
- D. F., Y.A. and P.B., "Sequences of Generalized. Least-Squares Problems on SMP Arch.", AMC 2014.
- L.B. and P.B., "Streaming Data from HDD to GPUs for Sustained Peak Performance", Euro-Par 2013.
- E.P., D.F., Y.A. and P.B., "Large-scale Whole Genome Association Analysis", PBio 2013 (EuroMPI'13).