

# High-performance and automatic computing

Paolo Bientinesi

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

October 22nd, 2013  
Goethe Universität Frankfurt am Main



Deutsche  
Forschungsgemeinschaft

**DFG**

- 1 HPAC: Introduction
- 2 Application: Genome-Wide Association Studies
- 3 Performance experiments
- 4 Conclusions

- **High-performance computing**

numerical computations, parallel architectures

→ time-to-solution, efficiency, scalability, ...

- **High-performance computing**

numerical computations, parallel architectures

→ time-to-solution, efficiency, scalability, . . .

- **Automatic computing**

optimization, search, but also derivation & deduction

→ range of algorithms, productivity

- **High-performance computing**

numerical computations, parallel architectures

→ time-to-solution, efficiency, scalability, . . .

- **Automatic computing**

optimization, search, but also derivation & deduction

→ range of algorithms, productivity

- **Applications**

→ application-specific properties & needs, large scale, full code

## methods & applications

- P. Bientinesi
- E. Di Napoli                      Electronic structure calculations
- M. Petschow                      Parallel eigensolvers
- D. Fabregat                      Automation, Computational biology
- D. Tameling\*                      Molecular dynamics
- E. Peise                              Performance modeling & prediction
- L. Beyer                              Density Functional Theory
- F. Kürten, Y. Madzhunkov, A. Frank, P. Springer, ...

## methods & applications

- P. Bientinesi
- E. Di Napoli                      Electronic structure calculations
- M. Petschow                      Parallel eigensolvers
- D. Fabregat                      Automation, Computational biology
- D. Tameling\*                      Molecular dynamics
- E. Peise                              Performance modeling & prediction
- L. Beyer                              Density Functional Theory
- F. Kürten, Y. Madzhunkov, A. Frank, P. Springer, ...



D. Fabregat



E. Peise




Y. Aulchenko

- 1 HPAC: Introduction
- 2 Application: Genome-Wide Association Studies**
- 3 Performance experiments
- 4 Conclusions




# Genome-Wide Association Studies



Source: David Hall

# Genome-Wide Association Studies



Source: David Hall

Yurii

Paolo

Yurii

“Mixed models”

Paolo

???

# Genome-Wide Association Studies

Yurii

Paolo

“Mixed models”

???

Linear regression with non-independent outcomes

???

# Genome-Wide Association Studies

Yurii

Paolo

“Mixed models”

???

Linear regression with non-independent outcomes

???

Generalized least-square problems

...

Yurii

Paolo

“Mixed models”

???

Linear regression with non-independent outcomes

???

Generalized least-square problems

...

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

- Inputs:  $M \in \mathbb{R}^{n \times n}$ ,  $X \in \mathbb{R}^{n \times p}$ ,  $y \in \mathbb{R}^n$
- Output:  $b \in \mathbb{R}^p$

★**To be repeated millions of times**★

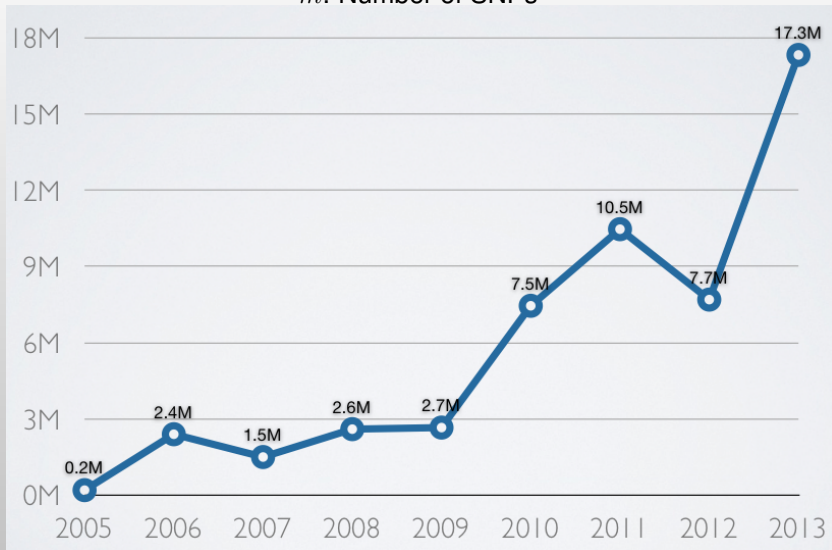




$n$ : Population size



*m*: Number of SNPs



## Problem definition (1)

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

“to be repeated millions of times”

# Problem definition (1)

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

“to be repeated millions of times”

↓

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

for  $i = 1, \dots, m$

# Problem definition (1)

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

“to be repeated millions of times”

↓

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

for  $i = 1, \dots, m$

↓

## Problem size

$M_i \in \mathbb{R}^{n \times n}$	$1000 \leq n \leq 20k+$	7.5MBs – 3GBs
$X_i \in \mathbb{R}^{n \times p}$	$3 \leq p \leq 20$	30 – 625KBs
$y_j \in \mathbb{R}^n$		8 – 780KBs
$b_i \in \mathbb{R}^p$		24 – 160 Bytes
<b>Total</b>	$10^6 \leq m \leq 10^8$	7.5 – 3000 TBs

## Problem definition (2)

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

## Problem definition (2)

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

↓

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

and  $X_i = [X_L | X_{Ri}]$ ,

for  $i = 1, \dots, m$

↓

### Problem size

$M \in \mathbb{R}^{n \times n}$      $1000 \leq n \leq 100k$     7.5MBs – 74.5GBs

$X_{Ri} \in \mathbb{R}^n$     8 – 780KBs

$b_i \in \mathbb{R}^p$     24 – 160 Bytes

---

Total     $10^6 \leq m \leq 10^8$     74GBs – 7 TBs

## Problem definition (3)

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



## Problem definition (3)

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

↓

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


and

$$M_j = \sigma_j(\Phi + h_j I),$$


for  $i = 1, \dots, m$  and  $j = 1, \dots, t$

Moreover, either  $t = 1$  or  $t \leq 10^5$ .

# GWAS: complete problem definition



# GWAS: complete problem definition



## 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
<b>Total</b>	$m \leq 10^8, t \leq 10^5$	1.5 – 100s TBs

# Vision: domain-specific compiler

$$\alpha, \beta, \gamma \in \mathbb{R}$$

$$\mu := (\gamma * \alpha^{-1} * \gamma)^{-1} * \beta$$

↓

**compiler**

↓

$$\text{code for } \mu := \frac{\alpha\beta}{\gamma^2}$$

# Vision: domain-specific compiler

$$\alpha, \beta, \gamma \in \mathbb{R}$$

$$\mu := (\gamma * \alpha^{-1} * \gamma)^{-1} * \beta$$

↓

**compiler**

↓

$$\text{code for } \mu := \frac{\alpha\beta}{\gamma^2}$$

Nov. 1954: *The IBM Mathematical FORMula TRANslating System, FORTTRAN*

“a set of programs to accept a concise formulation of a problem and to produce automatically a solution of the problem”

$$Y \in \mathbb{R}^{n \times p}, \quad A \in \mathbb{R}^{n \times n}, \quad b \in \mathbb{R}^n$$

$$v := (Y^T * A^{-1} * Y)^{-1} * b$$

⇓

< **lin. alg. compiler** >

⇓

algorithm, code

# Automatic generation

$$Y \in \mathbb{R}^{n \times p}, \quad A \in \mathbb{R}^{n \times n}, \quad b \in \mathbb{R}^n$$

$$v := (Y^T * A^{-1} * Y)^{-1} * b$$



< **lin. alg. compiler** >



algorithm, code

Matrix algebra

Sequences of operations

Inference of properties

Cost analysis

Building blocks

Code generation

# Algorithms generated

Algorithm 1

$$LL^T = M$$

$$X := L^{-1}X$$

$$S := X^T X$$

$$GG^T = S$$

$$y := L^{-1}y$$

$$b := X^T y$$

$$b := G^{-1}b$$

$$b := G^{-T}b$$

Algorithm 2

$$LL^T = M$$

$$X := L^{-1}X$$

$$QR := X$$

$$y := L^{-1}y$$

$$b := Q^T y$$

$$b := R^{-1}b$$

...

Algorithm 20

$$ZWZ^T = \Phi$$

$$D := (hW + (1-h)I)^{-1}$$

$$KK^T = D$$

$$X := Z^T X$$

$$X := K^T X$$

$$QR := X$$

$$y := L^{-1}y$$

$$b := Q^T y$$

$$b := R^{-1}b$$

...



# Many algorithms! Predictions?

# Many algorithms! Predictions?

## Flop count – rough estimate

	Alg. 1	Alg. 2	Alg. 20
Single instance ( $t = 1$ )	$O(n^3)$	$O(n^3)$	$O(n^3)$
2D sequence ( $t \gg 1$ )	$O(tn^3 + mtn^2)$	$O(tn^3 + mtn^2)$	$O(n^3 + mtn)$

# Many algorithms! Predictions?

## Flop count – rough estimate

	Alg. 1	Alg. 2	Alg. 20
Single instance ( $t = 1$ )	$O(n^3)$	$O(n^3)$	$O(n^3)$
2D sequence ( $t \gg 1$ )	$O(tn^3 + mtn^2)$	$O(tn^3 + mtn^2)$	$O(n^3 + mtn)$

### Analytic models

Roman Iakymchuk

### Model-based prediction

Elmar Peise

# Algorithm → implementations

## operands

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

# Algorithm → implementations

## operands

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

Does  $M$  fit in memory?

# Algorithm → implementations

## operands

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

## Does $M$ fit in memory?

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

## operands

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

### Does $M$ fit in memory?

- YES ⇒ single node + multithreading  
streaming HD↔CPU, double buffering, in-core implementation

### Does $M$ fit in GPU-memory?

- Yes ⇒ accelerator  
streaming HD↔CPU↔GPU, triple+double buffering, GPU implementation

## operands

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

### Does $M$ fit in memory?

- YES ⇒ single node + multithreading  
streaming HD↔CPU, double buffering, in-core implementation

### Does $M$ fit in GPU-memory?

- Yes ⇒ accelerator  
streaming HD↔CPU↔GPU, triple+double buffering, GPU implementation
- NO ⇒ distributed memory + MPI  
partitioning + streaming HD↔CPUs, double buffering, data distribution





- 1 HPAC: Introduction
- 2 Application: Genome-Wide Association Studies
- 3 Performance experiments**
- 4 Conclusions

# Results


$t = 1$

## Single-trait analysis




Single node


MPI



Single node

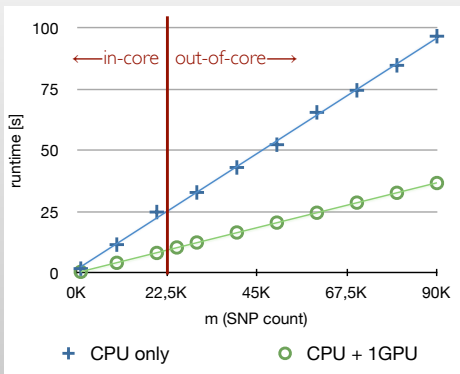


MPI




## 1 GPU

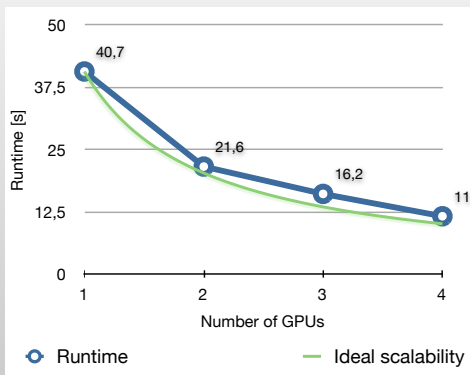
## Scalability



## 1 GPU



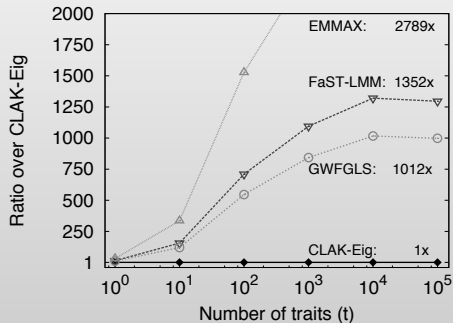
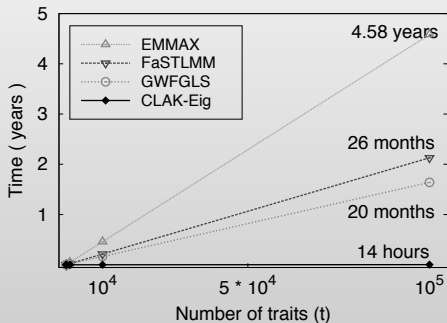
## Scalability



# Results

$t \gg 1$

Multi-trait analysis, "OMICS"-data



- 1 HPAC: Introduction
- 2 Application: Genome-Wide Association Studies
- 3 Performance experiments
- 4 Conclusions**



## HPC's perspective

- In-core efficiency
- How to sustain efficiency?

## HPC's perspective

- In-core efficiency
- How to sustain efficiency?

## App's perspective

- Many data formats
- Missing data, bogus data
- Output data. Post-processing?
- New features

## HPC's perspective

- In-core efficiency
- How to sustain efficiency?

## App's perspective

- Many data formats
- Missing data, bogus data
- Output data. Post-processing?
- New features

- HUGE gap:  
algorithm  $\leftrightarrow$  optimized implementation  
(data management, parallelism)

## HPC's perspective

- In-core efficiency
- How to sustain efficiency?

## App's perspective

- Many data formats
- Missing data, bogus data
- Output data. Post-processing?
- New features

- HUGE gap:  
algorithm ↔ optimized implementation  
(data management, parallelism)
- Development cycle: several months!

## HPC's perspective

- In-core efficiency
- How to sustain efficiency?

## App's perspective

- Many data formats
- Missing data, bogus data
- Output data. Post-processing?
- New features

- HUGE gap:  
algorithm ↔ optimized implementation  
(data management, parallelism)
- Development cycle: several months!
- How to deal with BIG problems?  
Expose knowledge, exploit knowledge