

Символьное обращение плохо обусловленных матриц в Грид-среде сервисов доступа к системе компьютерной алгебры МАХИМа.

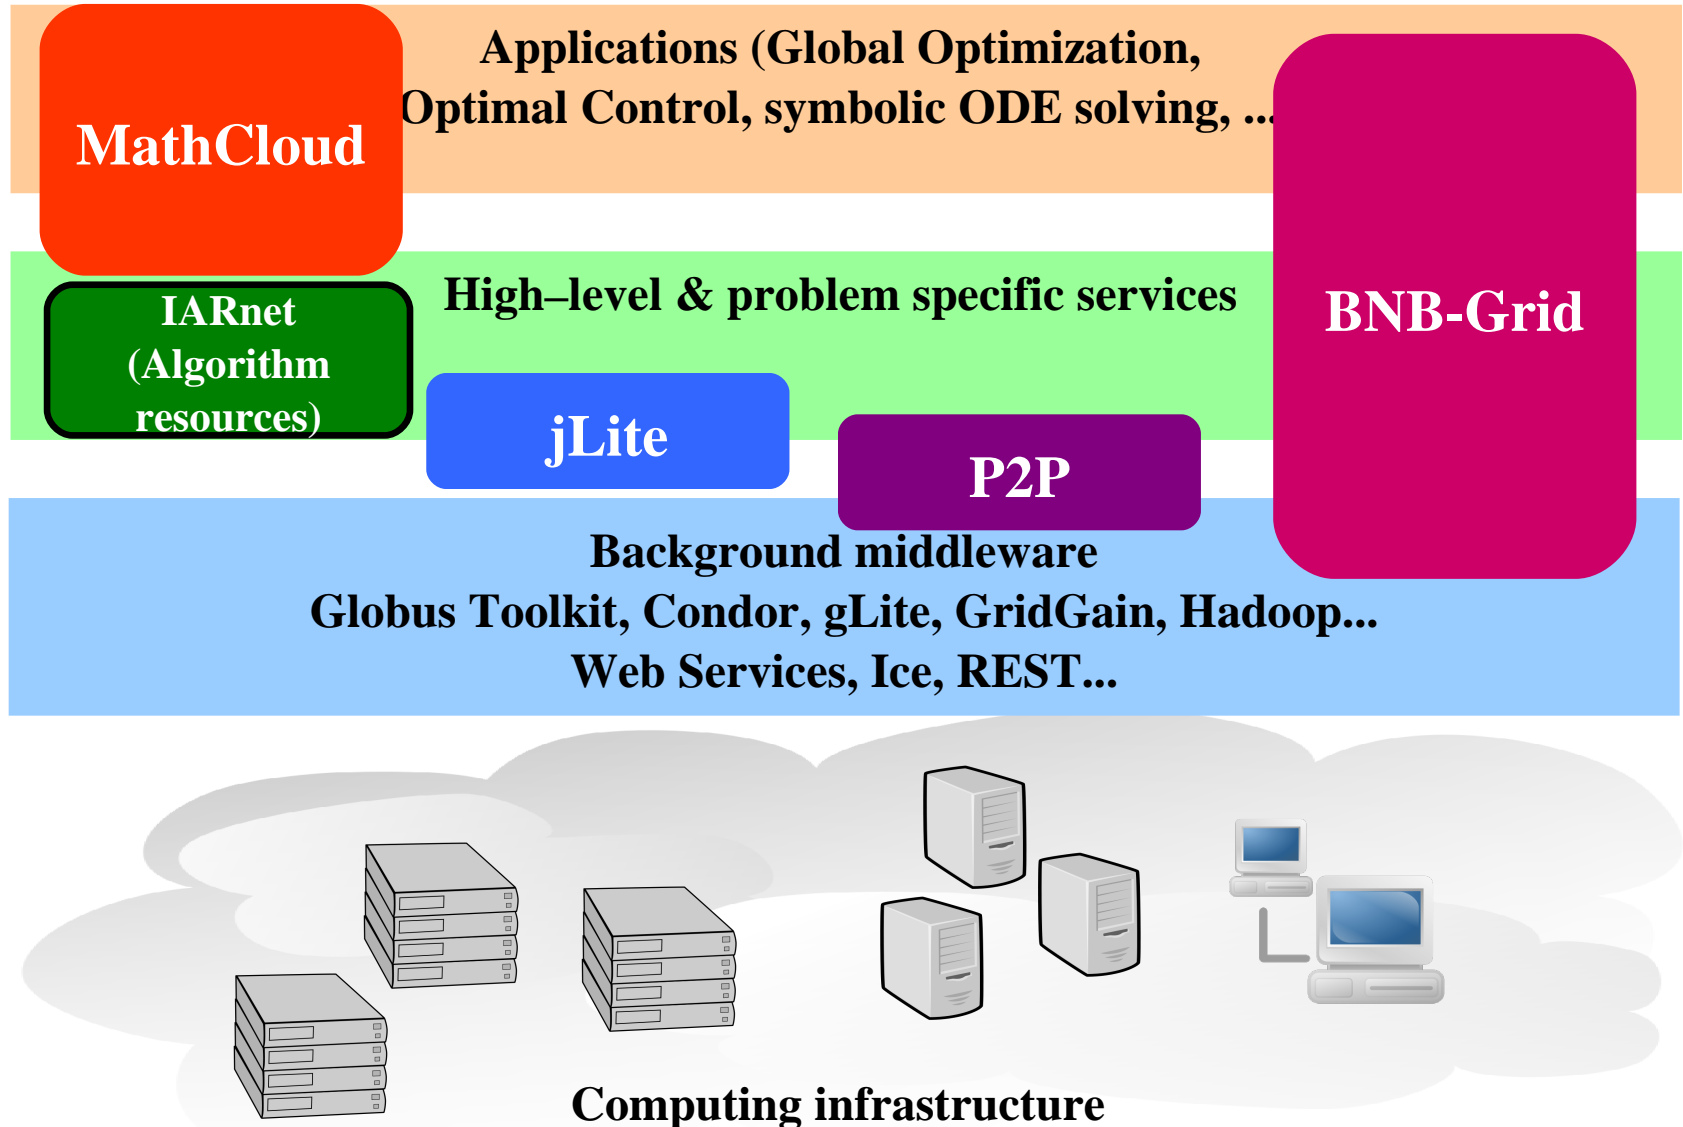
Ill-conditioned matrices symbolic inversion in Desktop Grid of CAS Maxima services.

Vladimir V. Voloshinov

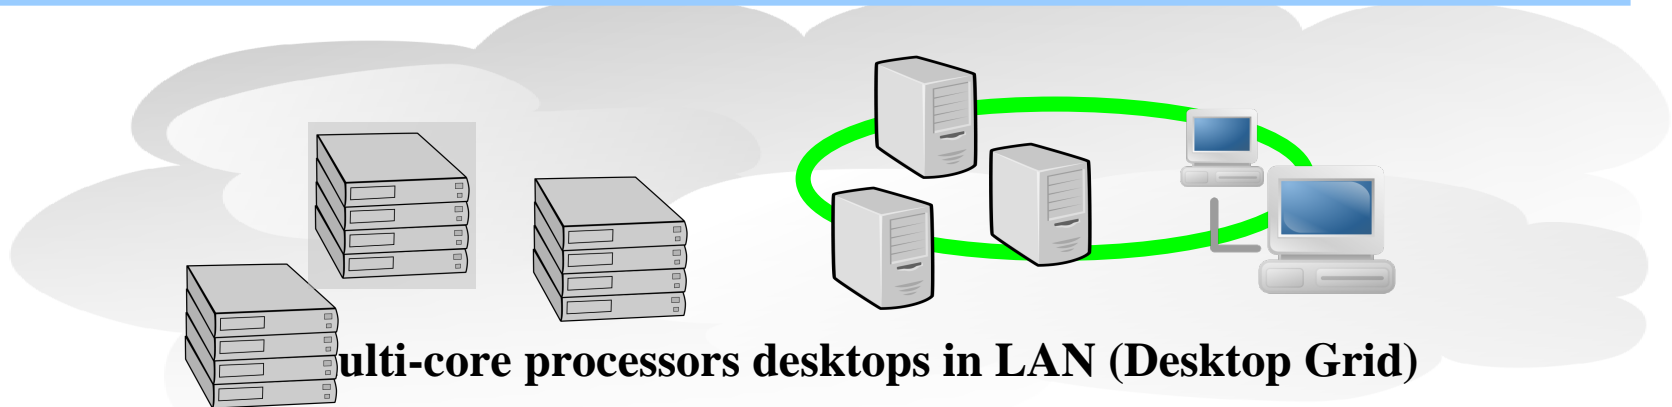
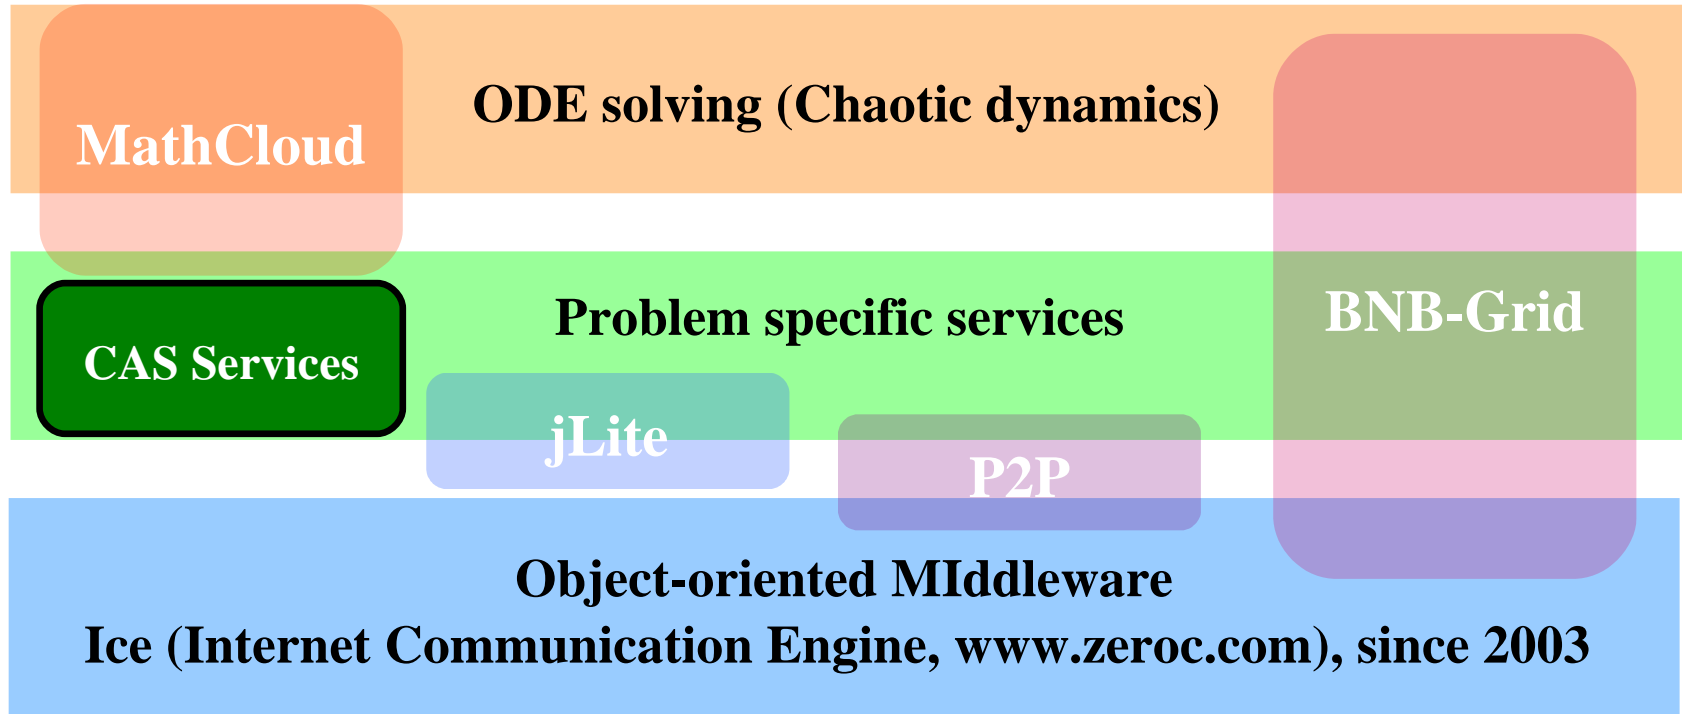
Center of Grid-technologies and Distributed Computing,
Institute of System Analysis RAS, <http://dcs.isa.ru>
Moscow, 2009

Supported by RFBR, grant #08-07-00430-a

Area of our studies



Relates to the report



CAS (Computer Algebra System) Maxima (1).

Has been started in Massachusetts Institute of Technology, by prof. William Schelter. Since 1998 - GNU Public License.

**The core - GCL (GNU Common Lisp), Windows, Linux
Single-threaded Lisp interpreter**

Has almost the same basic “symbolic” capabilities as Maple and Mathematica: differentiation&integration, series, ODE solving, matrices and linear algebra, polynomials, sets, lists, tensors...

<http://maxima.sourceforge.net/> (GPL)

<http://maxima.sourceforge.net/ru/>

CAS (Computer Algebra System) Maxima (2).

Inherent feature of float arithmetic computing — accumulation of rounding errors in intermediate operations.

Maxima handles rational number x as pairs of {numerator, denominator}, $\{p,q\}: x = \frac{p}{q}$

All arithmetic operations are performed without loss of accuracy.

Theoretically, regardless memory (and elapsed time!) unlimited accuracy may be expected.

If necessary, only final result may be rounded to float representation.

CAS (Computer Algebra System) Maxima (3).

Procedural programming language (MaximaScript).

MaximaScript «wraps» Lisp.

Examples of script:

Some
initialisation.

```
/*N:100;*/  
/*TID:"H";*/  
path2common:"../common/";  
...  
load(eigen);  
colInvByLUkd(LU,N,k):=lu_backsub(LU,unitVect(k,N));  
luFileName(TID,N):=simplode([path2common,"LU",TID,N,".lsp"]);  
invColFileName(TID,N,k):=simplode(["invTrCol",TID,N,"_",k,".lsp"]);  
M:hilbert_matrix(N);
```

Composing inverse
matrix from its
columns written to
files before.

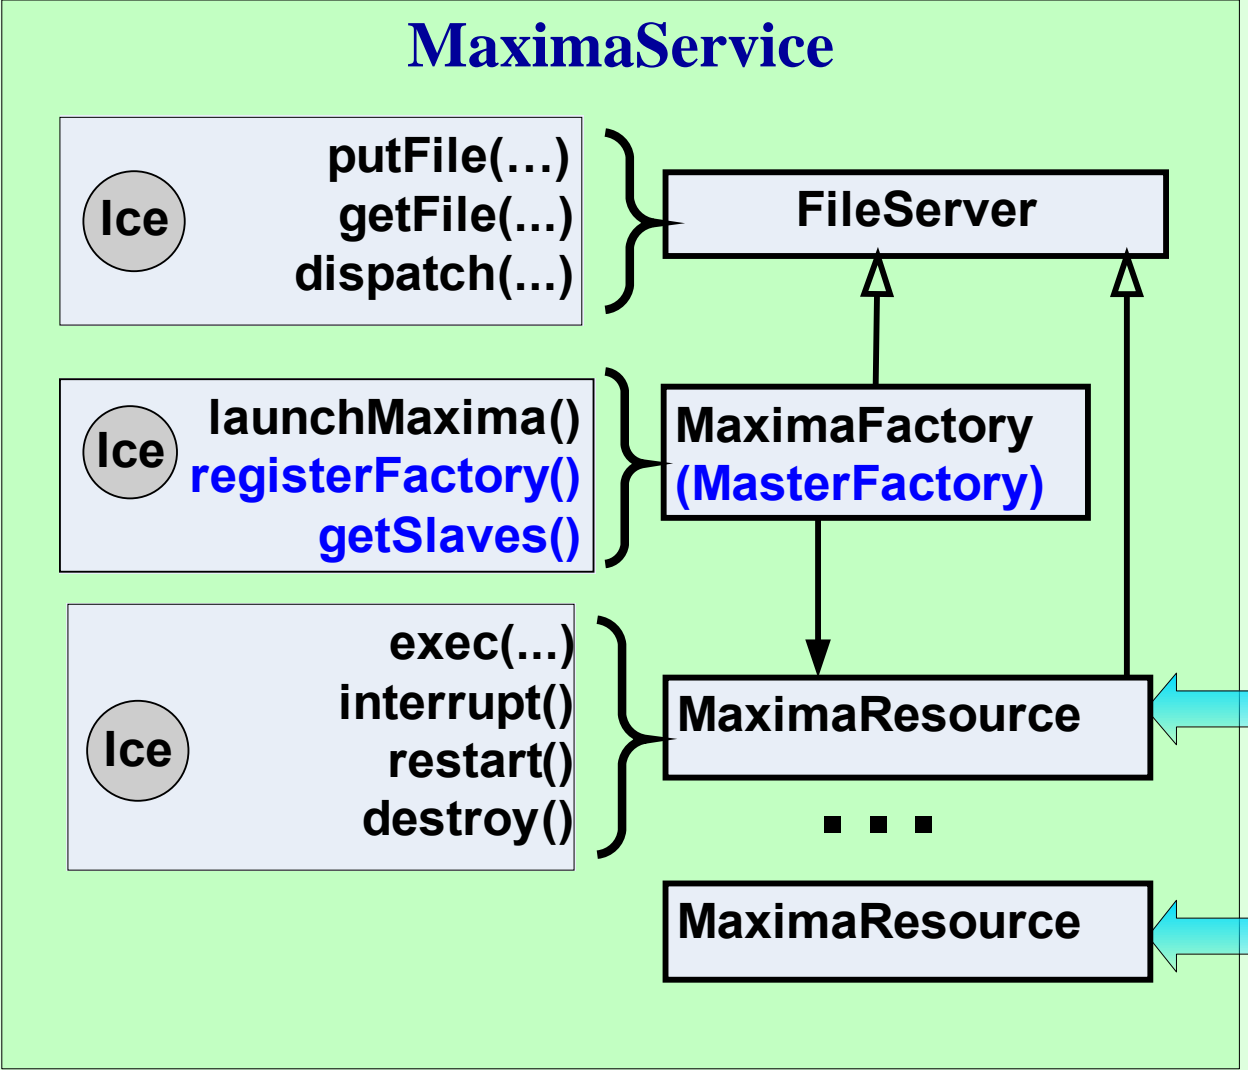
```
invColFileName(TID,N,k):=simplode(["invTrCol",TID,N,"_",k,".lsp"])$  
colName(TID,N,k):=concat('invTrCol,TID,N,"_",k)$  
readInv(TID,N,rpath):=block([k,fname,cmd],  
  tmpM:zeromatrix(0,N),  
  for k:1 thru N step 1 do (  
    fname:simplode([rpath,invColFileName(TID,N,k)]),  
    load(fname),  
    cmd:simplode(["tmpM:addrow(tmpM,"",colName(TID,N,k),"")]),  
    eval_string(cmd)  
  ),  
  return(tmpM)  
);
```

Maxima remote access service.

Node (unified services container IARnet2, Ice)

Basic service management

Resource access control and security



Available implementation.

**<http://code.google.com/p/remote-maxima>, GPL
Language C++, platforms Windows XP, Linux, MacOS**

**Ice, <http://www.zeroc.ice>, version 3.2.0 and later
Maxima version 5.12.0 and later.**

Required additional libraries:

Boost 1.36.0, <http://www.boost.org/>

Boost.Process library,

<http://www.highscore.de/boost/process.zip>

CxxTest, <http://cxxtest.tigris.org>

Hilbert matrices

Matrices H_N of the type: $H_N = \{h_{m,n}\}_{m=1,n=1}^{N,N}$, где $h_{m,n} = \frac{1}{m+n-1}$

Gram matrix of the basic polynomials in $L_2[0,1]$ $h_{mn} = \int_0^1 t^{m-1} \cdot t^{n-1} dt$

Well-known example of ill-conditioned matrices.

Condition number of H_N is growing exponentially w.r.t N

$$\text{cond}(H_N) = \|H_N\| \cdot \|(H_N)^{-1}\| \sim e^{3.5 \cdot N}$$

$$\|A_{N \times N}\| = \left(\sum_{m=1, n=1}^{N, N} A_{i,j}^2 \right)^{1/2}$$

| N | $\text{cond}(H_N)$ |
|-----|----------------------|
| 10 | $1.6 \cdot 10^{13}$ |
| 50 | $1.5 \cdot 10^{74}$ |
| 70 | $5.5 \cdot 10^{104}$ |
| 100 | $4.1 \cdot 10^{150}$ |
| 150 | $1.2 \cdot 10^{227}$ |

Values of $\text{cond}(H_N)$ for some N (calculated exactly in Maxima)

«Traditionally», «well-conditioned» matrices should have cond less than 1000.

Possible speedup reasoning for «*LU-inversion*» of H_N (1)

Let M be $[N \times N]$ matrix. LU-factor $\{L, U, P\}$: L - lower-triangular; U - upper-triangular; P – permutation matrix. $L \cdot U = P \cdot M$.

Let E_N be unit matrix. To obtain M^{-1} - solve (by forward and backward substitution): $L \cdot U \cdot X = P \cdot E_N \Rightarrow X = M^{-1}$

Maxima has function *lu_factor*(M) (Gaussian elimination) and *lu_backsub*(LU, B) (LU — returned *lu_factor*(M)) to solve $L \cdot U \cdot X = P \cdot B$ for any rectangular matrix $B[N \times M]$.

invert_by_lu(M):=*lu_backsub*(*lu_factor*(M), E_N)

If E_N is sliced vertically into K submatrices

$$\mathbf{E}_N = \left[\mathbf{E}_N^{1:n_1} \mid \mathbf{E}_N^{n_1+1:n_2} \mid \mathbf{E}_N^{n_2+1:n_3} \mid \dots \mid \mathbf{E}_N^{n_{K-1}+1:N} \right], \quad 1 = n_0 < n_1 < n_2 < n_3 < \dots < n_{K-1} < n_K = N$$

then parallel call of *lu_backsub*($LU, \mathbf{E}_N^{n_{k-1}+1:n_k}$) gives K set of inverse columns $M^{-1}[n_{k-1}:n_k]$.

$K=N$ means parallel calculations of M^{-1} columns.

Possible speedup reasoning for «*LU-inversion*» of H_N (2)

Durations of two phases of inversion :

LU-factorization (Gaussian elimination) and calculation of inverse matrix's columns.

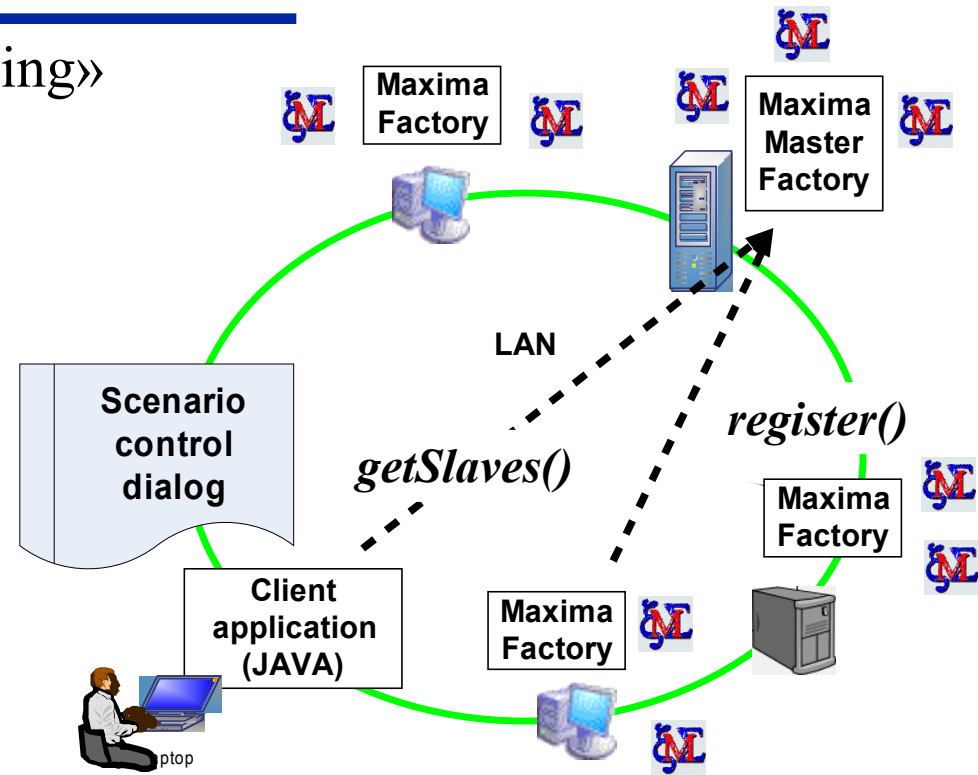
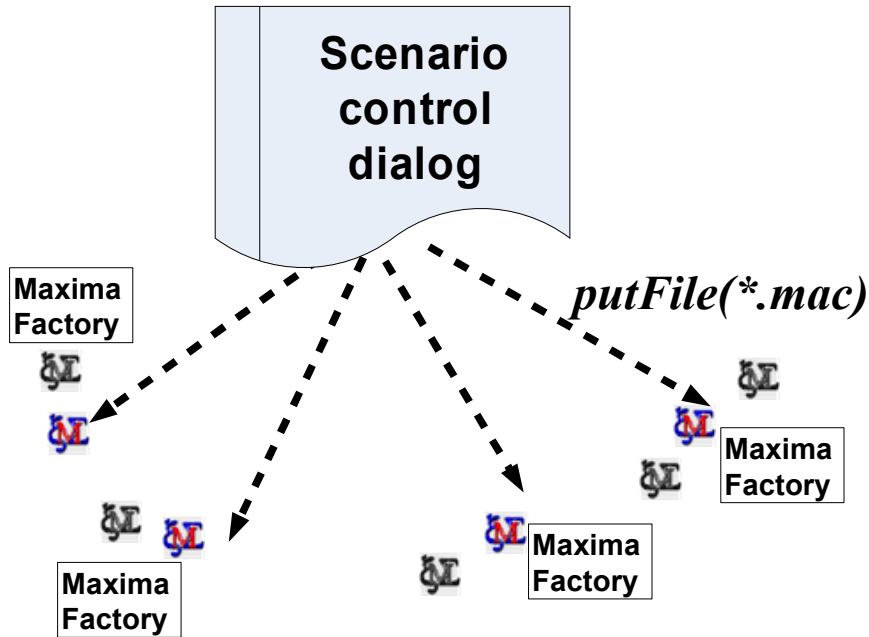
| N | $invert_by_lu(H_N), T_{inv},$ sec | $LUP=lu_factor(H_N), T_{LU}$ sec | $lu_backsub(LUP, E_N), T_{bs}$ sec |
|-----|--|--------------------------------------|--|
| 100 | 27 | 9 | 17 |
| 150 | 122 | 43 | 76 |
| 200 | 368 | 128 | 229 |
| 250 | 826 | 306 | 513 |
| 300 | 1684 | 631 | 1043 |

«Columns phase» (T_{bs}) takes almost twice of LU-factorization.

In the case of unlimited number of resources ~300% speedup may be expected.

«LU» scenario.

0 phase: start nodes, scripts «dispatching»



Windows Task Manager

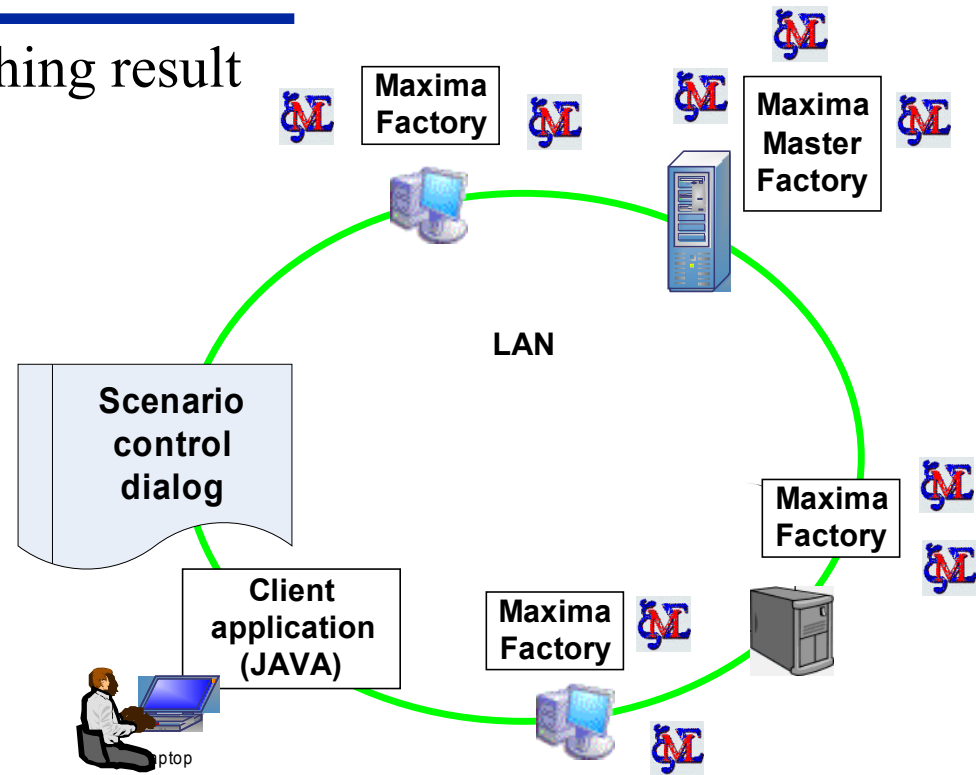
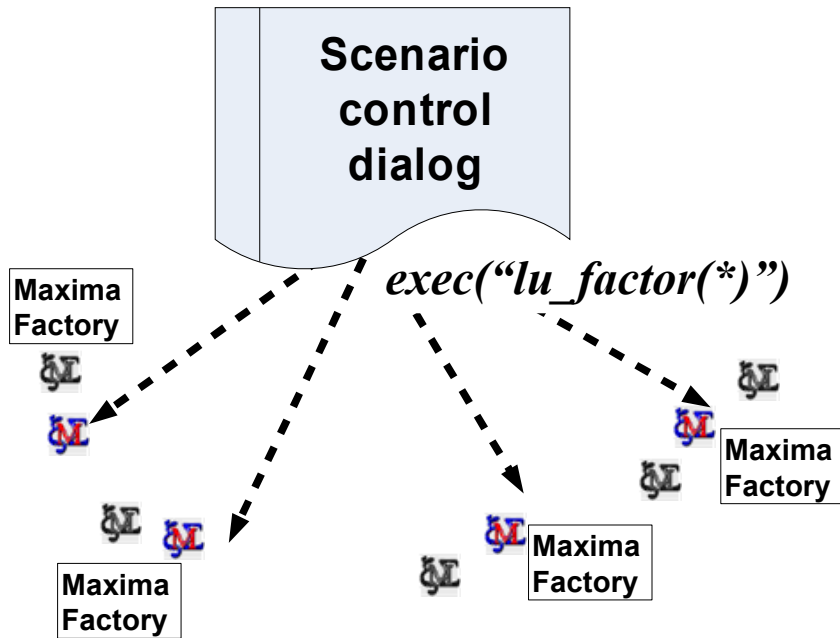
File Options View Help

Applications Processes Performance Networking Users

| Image Name | PID | User N... | Session ID | CPU | CPU Time | Mem Us... | VM Size | Handles | Threads | I/O Read Bytes |
|---------------------|------|-----------|------------|-----|----------|-----------|-----------|---------|---------|----------------|
| System Idle Process | 0 | SYSTEM | 0 | 75 | 21:33:35 | 28 K | 0 K | 0 | 4 | 0 |
| maxima.exe | 2920 | | 0 | 25 | 0:00:39 | 230 400 K | 289 612 K | 42 | 1 | 1 248 977 |
| Firefox Reader.exe | 5148 | | 0 | 00 | 0:00:01 | 9 676 K | 4 516 K | 119 | 3 | 172 |
| chrome.exe | 4676 | | 0 | 00 | 0:00:02 | 25 464 K | 20 056 K | 59 | 3 | 302 229 |

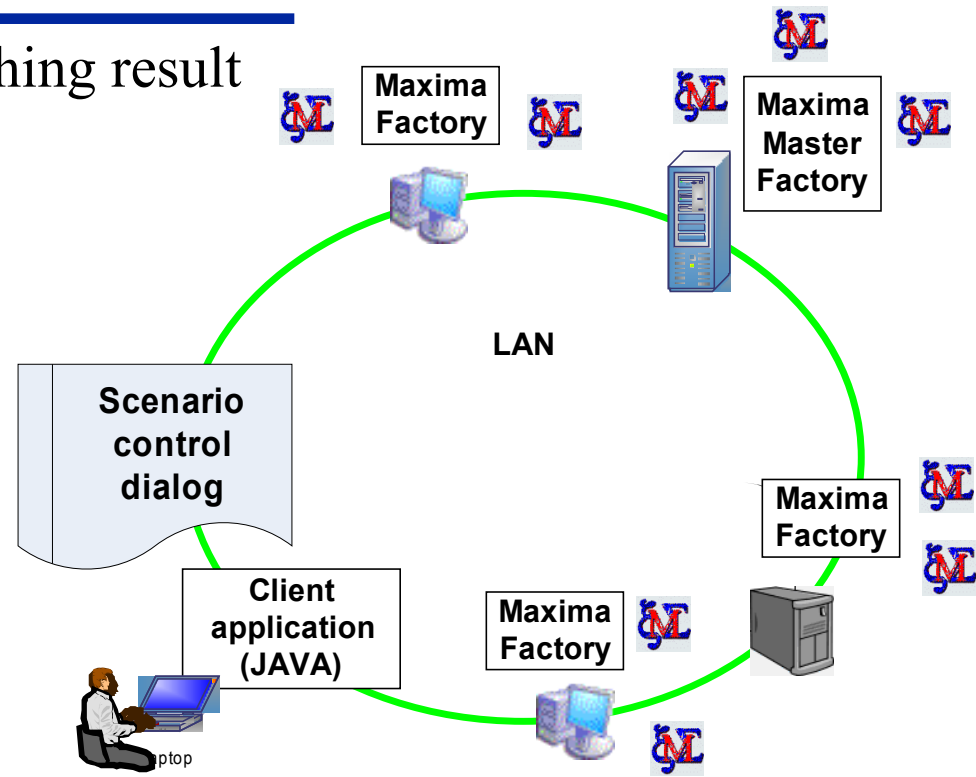
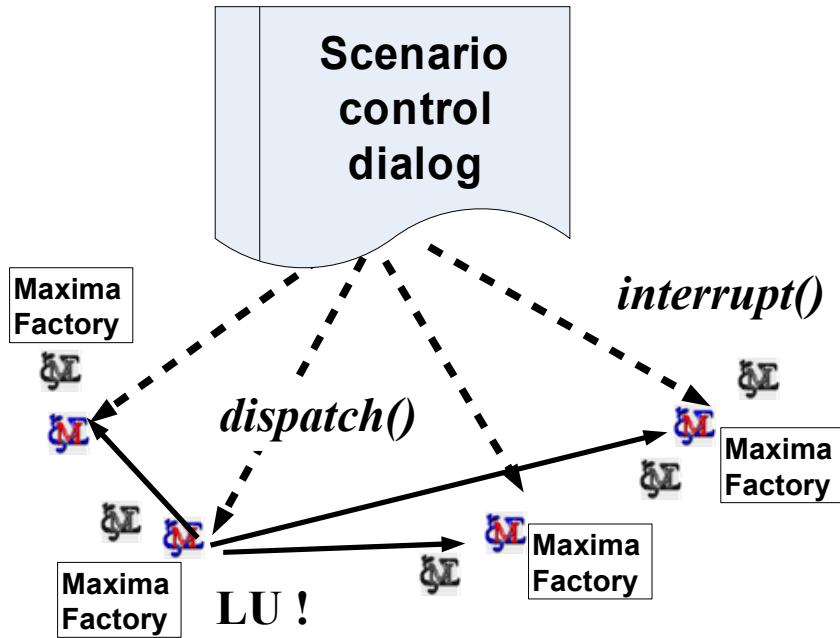
«LU» scenario.

I phase: LU-factorization and dispatching result



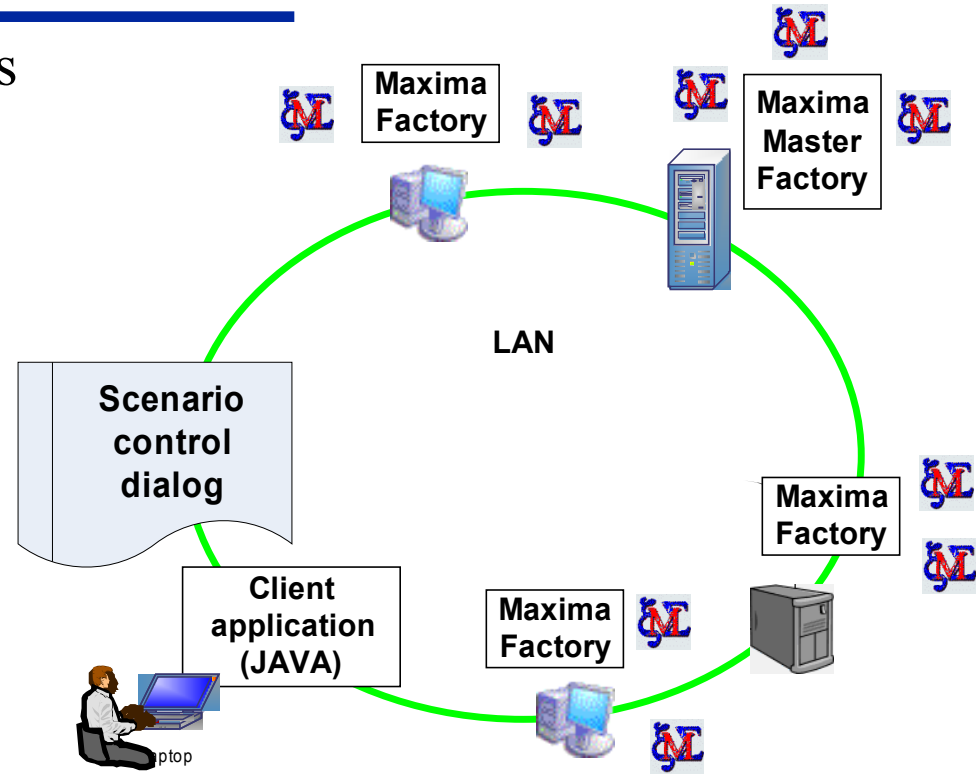
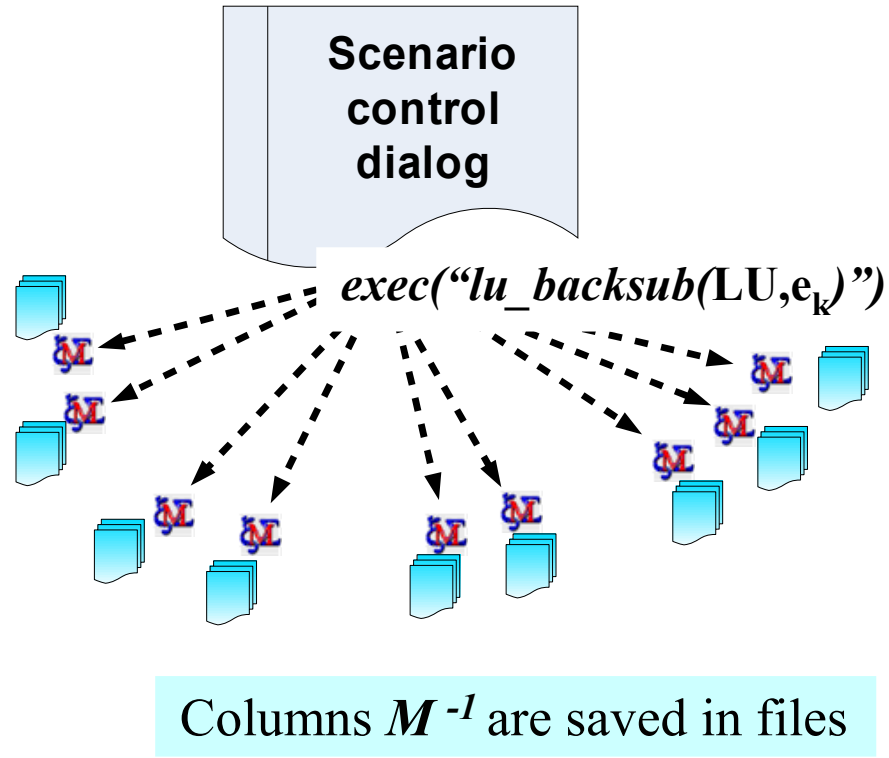
«LU» scenario.

I phase: LU-factorization and dispatching result



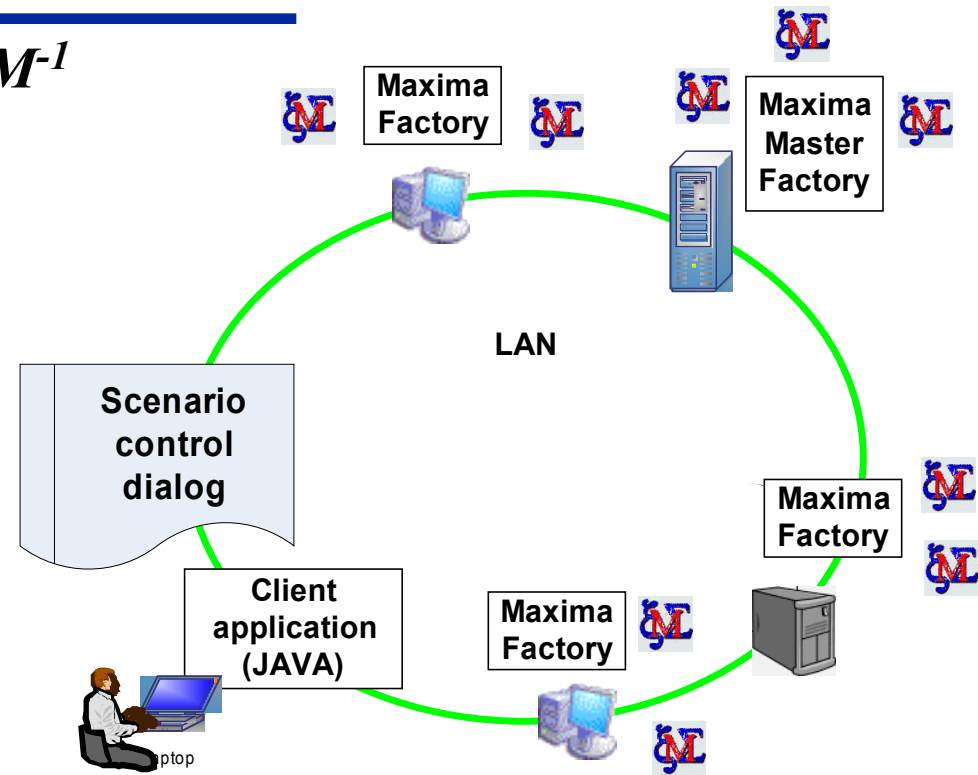
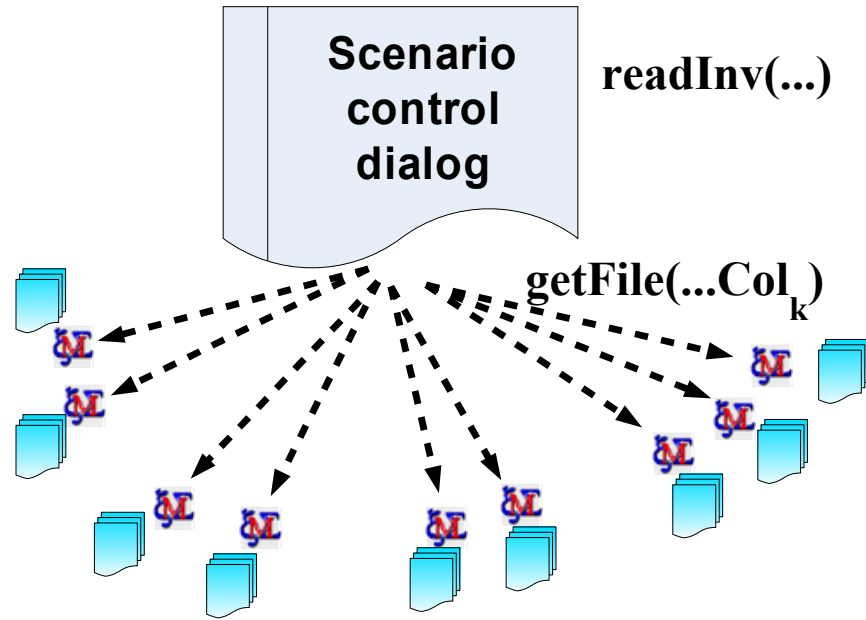
«LU» scenario.

II phase: inverse columns calculations



«LU» scenario.

III phase: collect columns and build M^{-1}

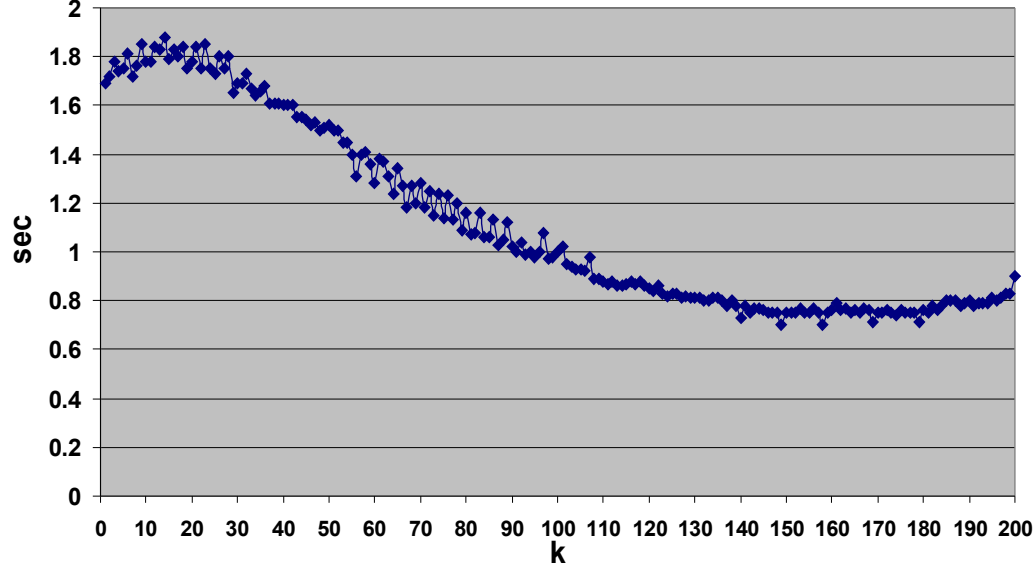


Columns M^{-1} are sent in files

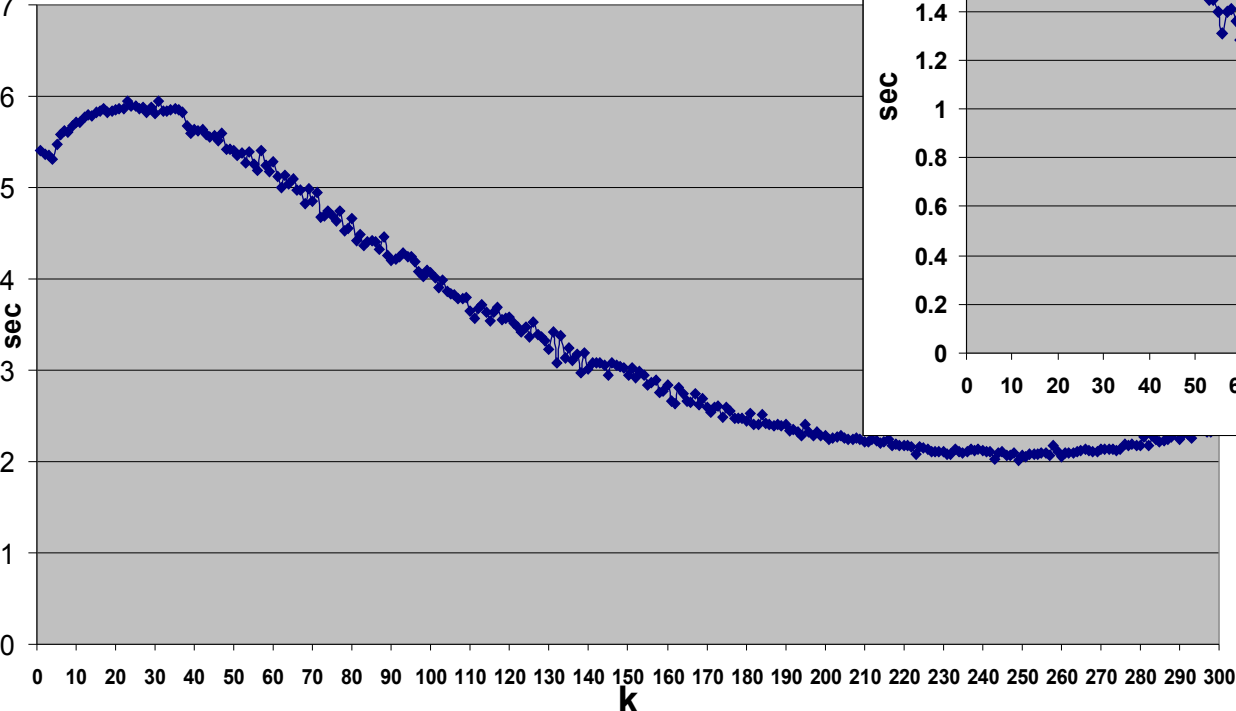
Load balancing problem in «LU» scenario

| Host id | Processor | Cores involved | Memory | OS | Perform. |
|---------|------------------------------------|--------------------|--------|-------|----------|
| 0 | Intel Core Quad Q6600 2.4 GHz | 3 | 3.8 Gb | Linux | 1.1 |
| 1 | Intel Core Quad Q6600 2.4 GHz | 3 | 2 Gb | WinXP | 1.0 |
| 2 | Intel Centrino Dual-Core 1.466 GHz | 2 | 1 Gb | WinXP | 0.6 |
| 3 | Pentium 4, 3 GHz | 2 (HyperThreading) | 1 Gb | WinXP | 0.6 |
| 4 | Pentium 4 2.8 GHz | 2 (HyperThreading) | 512 Mb | WinXP | 0.5 |

Duration of k-th column calc., inv(H200)



Duration of k-th column calculation, inv(H300)



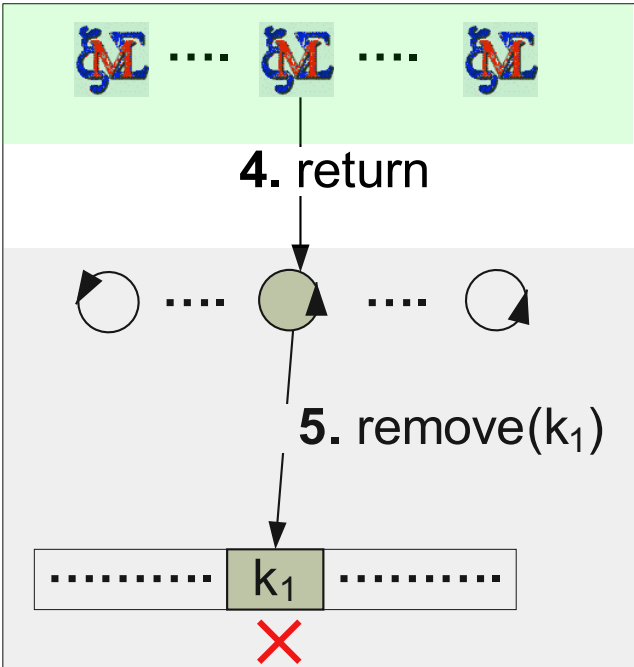
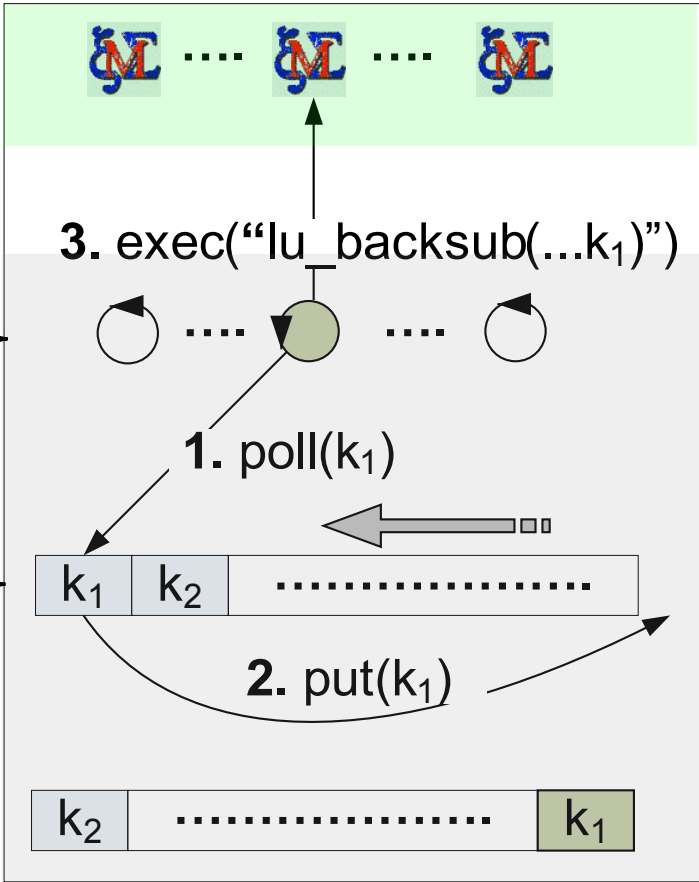
Load balancing implementation at Phase II (inv. columns calc)

Maxima services

Client application

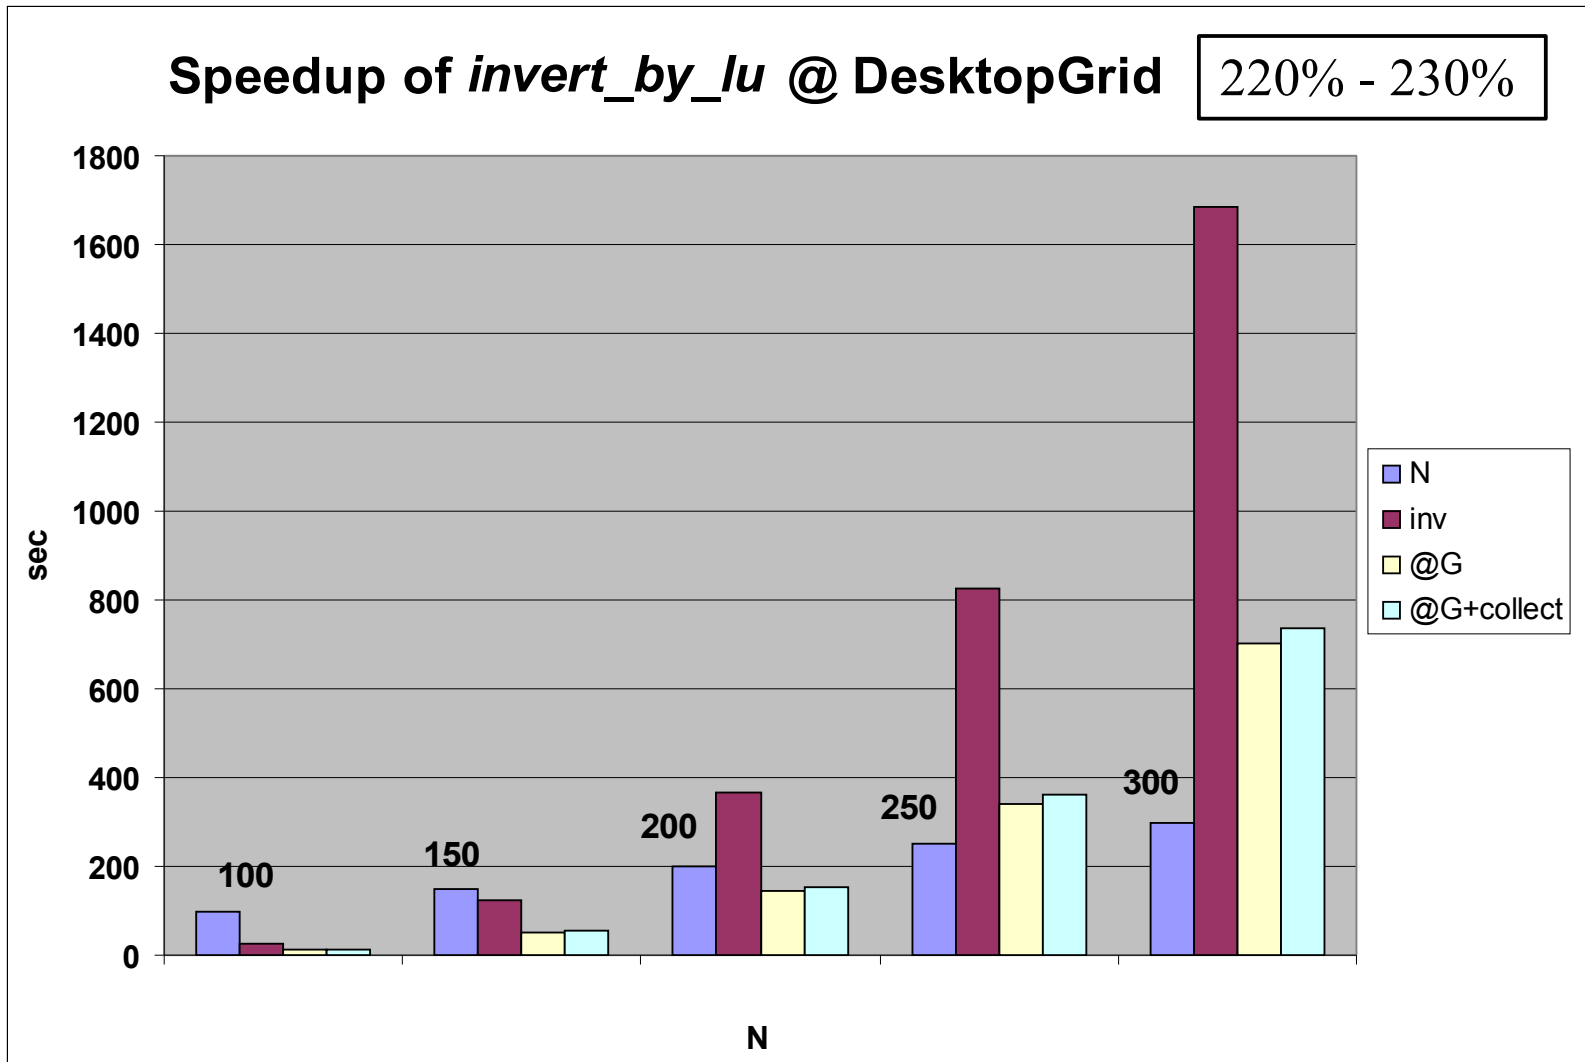
Worker threads

Thread-safe task queue



It turned out, that this heuristic gives ~1.5 times more than optimal schedule of Phase II (in deterministic post facto task assignment problem formulation)

Results of tests of distributed inversion via LU-decomposition



«Collection» of inverse matrix's columns (as files) takes ~ 2.5% of total scenario duration (including network data transmission)

Matrix inversion by Schur complement (1)

There is another inversion approach based on «block decomposition» and Schur complement

$$\mathbf{M}[N \times N], \mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{U} \\ \mathbf{V} & \mathbf{B} \end{bmatrix}, \exists \mathbf{M}^{-1}$$

$$\mathbf{A} = [N_A \times N_A], \mathbf{B} = [(N - N_A) \times (N - N_A)].$$

Schur complement: $\mathbf{S} = \mathbf{B} - \mathbf{V}\mathbf{A}^{-1}\mathbf{U}$, then $\exists \mathbf{S}^{-1}$ и

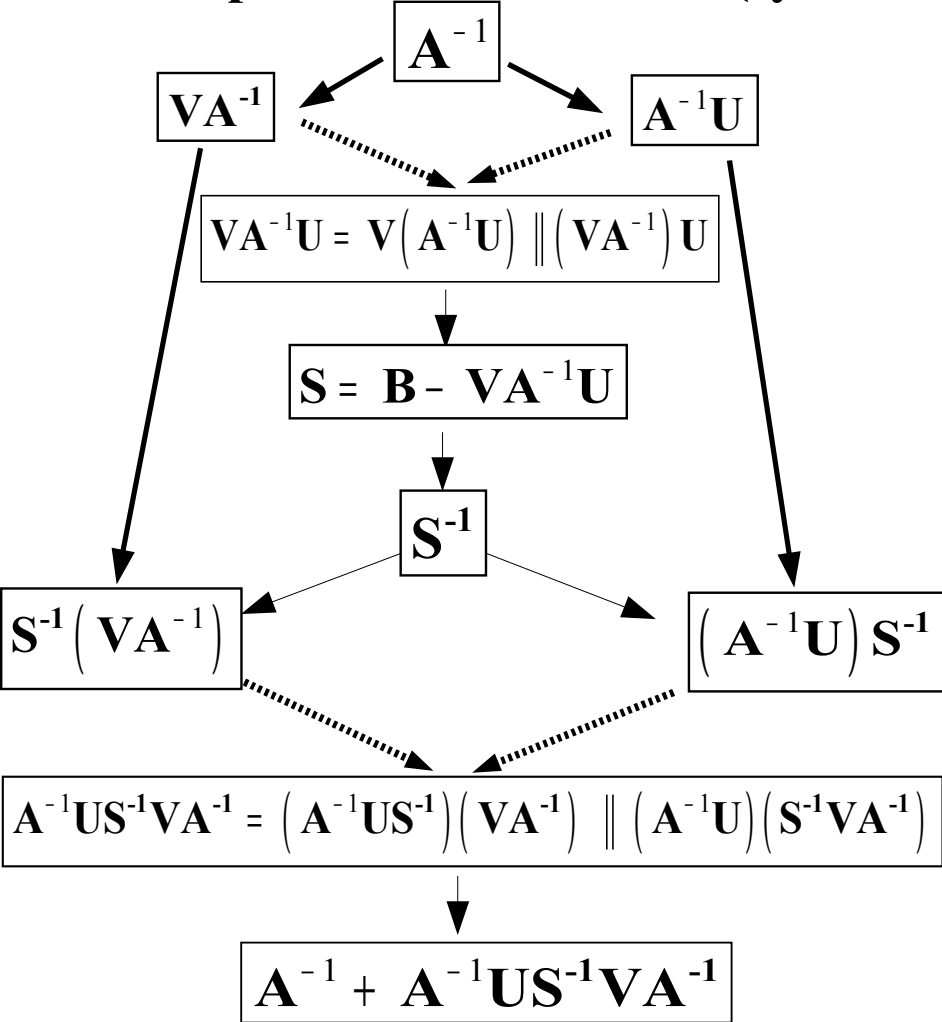
$$\mathbf{M}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{U}\mathbf{S}^{-1}\mathbf{V}\mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{U}\mathbf{S}^{-1} \\ -\mathbf{S}^{-1}\mathbf{V}\mathbf{A}^{-1} & \mathbf{S}^{-1} \end{bmatrix}.$$

Blocks processing provides more flexible for subsequent parallel and recursive algorithm because calculation of \mathbf{A}^{-1} , \mathbf{S}^{-1} and matrices multiplications may be parallelized as well.

Speedup evaluation at the next slide.

Matrix inversion by Schur complement (2)

Let $M[N \times N]$ be divided into four $[N/2 \times N/2]$ blocks. The cost of inverse matrix blocks' "parallel" calculation (symbol $\ll||\gg$) may be evaluated as follows.

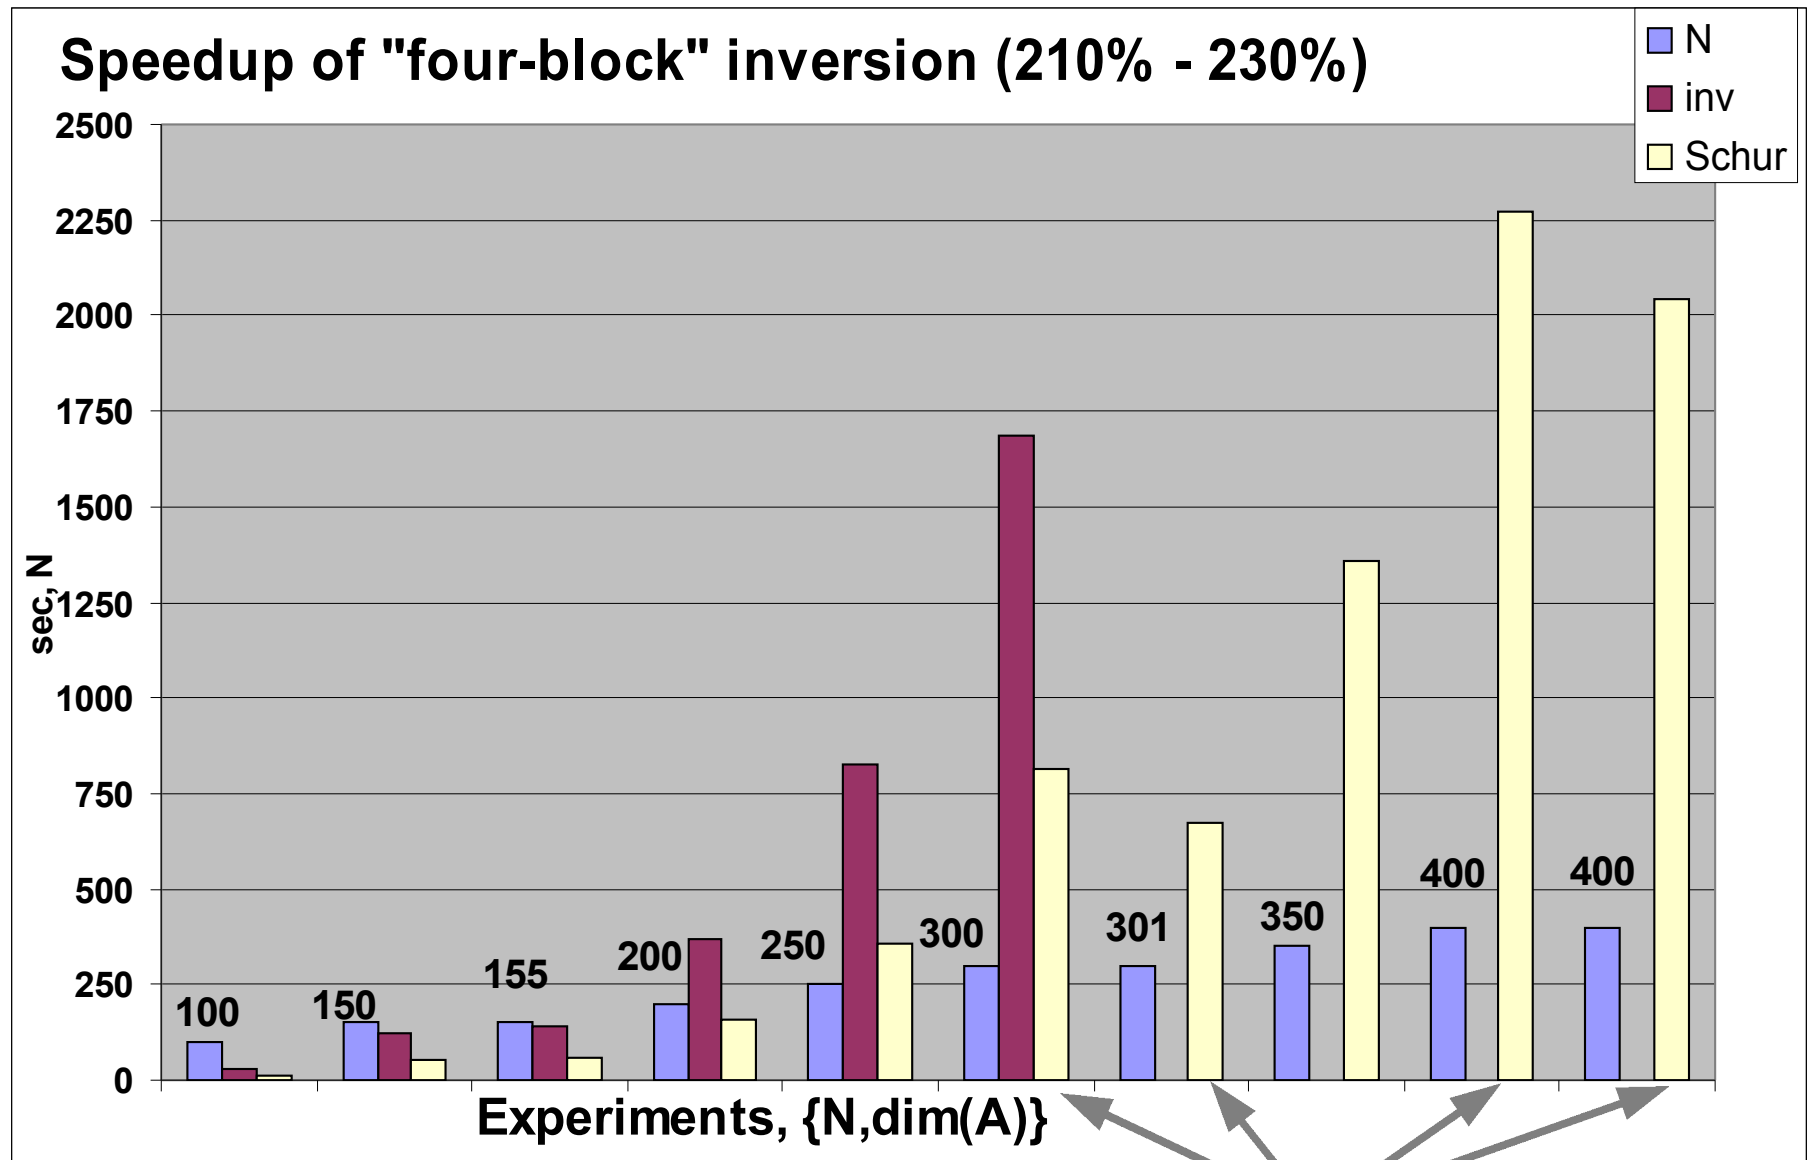


- $A^{-1} \Rightarrow \sim O((N/2)^3)$
- $VA^{-1} \parallel A^{-1}U \Rightarrow \sim O((N/2)^3)$
- $VA^{-1}U \Rightarrow \sim O((N/2)^3)$
- $B - VA^{-1}U \Rightarrow \sim O((N/2)^2)$
- $S^{-1} \Rightarrow \sim O((N/2)^3)$
- $S^{-1}(VA^{-1}) \parallel (A^{-1}U)S^{-1} \Rightarrow \sim O((N/2)^3)$
- $A^{-1}US^{-1}VA^{-1} \Rightarrow \sim O((N/2)^3)$
- $A^{-1} + A^{-1}US^{-1}VA^{-1} \Rightarrow \sim O((N/2)^2)$

So, speedup:

$$\frac{4N^3}{3N^3 + 2N^2} \approx \frac{4}{3} \quad (130\% \text{ for large } N)$$

Preliminary results of Schur complement approach (standalone «simulation»)



Speedup appreciably depends on blocks sizes.

Thank you!