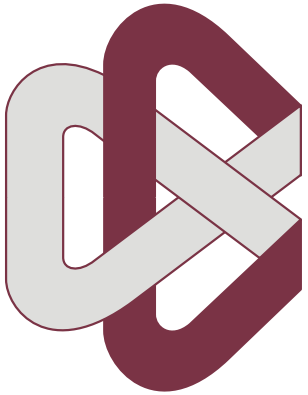
FEMT, an open source library and tools for solving large systems of equations in parallel



CIMAT

Miguel Vargas-Félix, Salvador Botello-Rionda

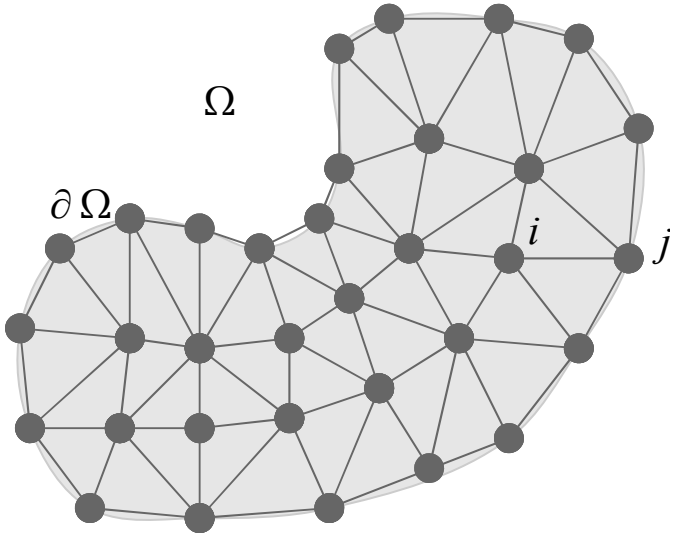
May 30, 2012

FEMT

- FEMT is an open source multi-platform library and tools (Windows, GNU/Linux, Mac OS, BSD), released under the GNU Library General Public License.
- It contains routines to handle and solve large linear systems of equations resulting from finite element and finite volume discretizations.
- It includes several solvers that run in parallel in multi-core computers using OpenMP, or in clusters of computers with MPI.
- These solvers can be used in stationary or dynamic problems.
- A sparse matrix of $1'000{,}000 \times 1'000{,}000$ can be factorized with Cholesky in 60 seconds in a 8-core computer.
- Using a cluster with 116 cores it can solve a system with 150'000,000 equations in 4.8 hours.
- It has been programmed in modern standard C++. It supports Microsoft Visual C++ 2008, GNU Compiler Collection 4.3 or Intel C++ Compiler 10.1.
- FEMT also includes several programs that allow access to library routines through pipes. This makes really easy to use FEMT from Fortran, C, Python, C#, etc.

Sparse matrices

Relation between adjacent nodes is captured as entries in a matrix. Because a node has adjacency with only a few others, the resulting matrix has a very sparse structure.



$$A = \begin{pmatrix} \circ & \circ & \circ & \circ & \circ & \circ & \circ & \circ & \cdots \\ \circ & a_{ii} & \circ & a_{ji} & \circ & \circ & \circ & \circ & \cdots \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & \circ & \cdots \\ \circ & a_{ij} & \circ & a_{jj} & \circ & \circ & \circ & \circ & \cdots \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & \circ & \cdots \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & \circ & \cdots \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & \circ & \cdots \\ \circ & \circ & \circ & \circ & \circ & \circ & \circ & \circ & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

When assembled, we have to solve a linear system of equations $\mathbf{A}\,\mathbf{x} = \mathbf{b}$.

$$A = \begin{pmatrix} 8 & 4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 3 & 0 & 0 \\ 2 & 0 & 1 & 0 & 7 & 0 \\ 0 & 9 & 3 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 5 \end{pmatrix},$$

8	4	
1	2	
1	3	
3	4	
2	1	7
1	3	5
9	3	1
2	3	6
5		
6		

$$\boldsymbol{v}_4^A = (9,3,1)$$

$$\boldsymbol{j}_4^A = (2,3,6)$$

Solvers for sparse systems on multi-core computers

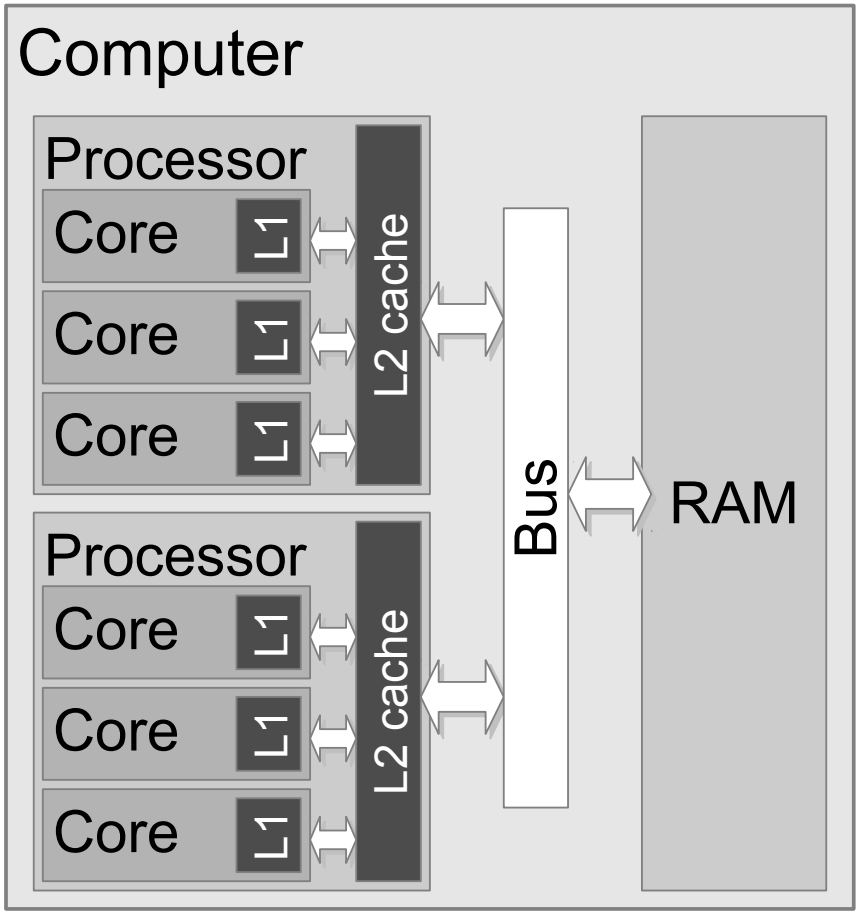
- Cholesky
- Cholesky (LDL')
- LU (Doolittle version)
- Conjugate gradient
- Conjugate gradient + Jacobi preconditioner
- Conjugate gradient + incomplete Cholesky factorization preconditioner
- Conjugate gradient + incomplete Cholesky LDL' factorization preconditioner
- Conjugate gradient + factorized sparse approximate inverse preconditioner
- Biconjugate gradient
- Biconjugate gradient + Jacobi preconditioner
- Biconjugate gradient + incomplete LU factorization preconditioner
- Biconjugate gradient + minimal residual sparse approximate inverse preconditioner

Solver for sparse matrices on clusters of computers

- Schur substructuring method

Parallelization using multi-core computers

Several processing units (cores) that access the same memory. All the cores fight for the RAM. We must design our programs to reduce this conflict.



In modern computers the processor is a lot faster than the memory, between them a high speed memory is used to improve data access. The ***cache***.

The most important issue to achieve high performance is to use the cache efficiently.

Access to	Cycles
Register	≤ 1
L1	~ 3
L2	~ 14
Main memory	~ 240

- Work using continuous memory blocks.
- Access memory in sequence.
- Each core should work in an independent memory area.

Algorithms to solve system of equations should take care of this.

Cholesky factorization for sparse matrices

For full matrices the computational complexity of Cholesky factorization $\mathbf{A}=\mathbf{L}\,\mathbf{L}^{\mathrm{T}}$ is $O\left(n^3\right)$.

To calculate entries of $\mathbf{L}$

$$L_{ij}=\frac{1}{L_{jj}}\left(A_{ij}-\sum_{k=1}^{j-1}L_{ik}L_{jk}\right),\ \text{for } i>j$$

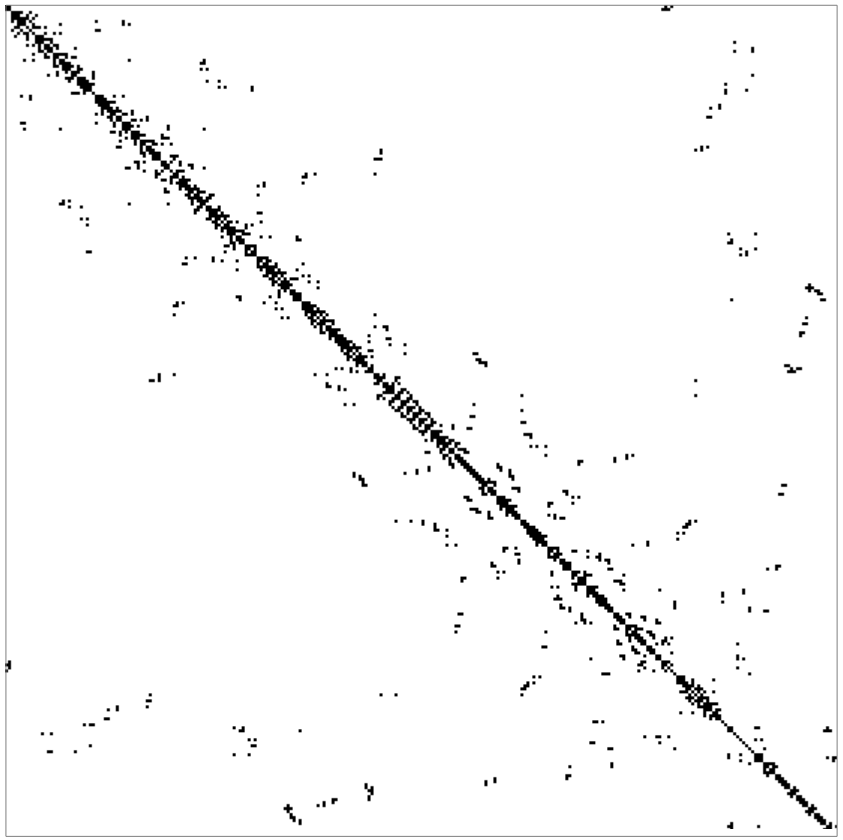
$$L_{jj}=\sqrt{A_{jj}-\sum_{k=1}^{j-1}L_{jk}^2}.$$

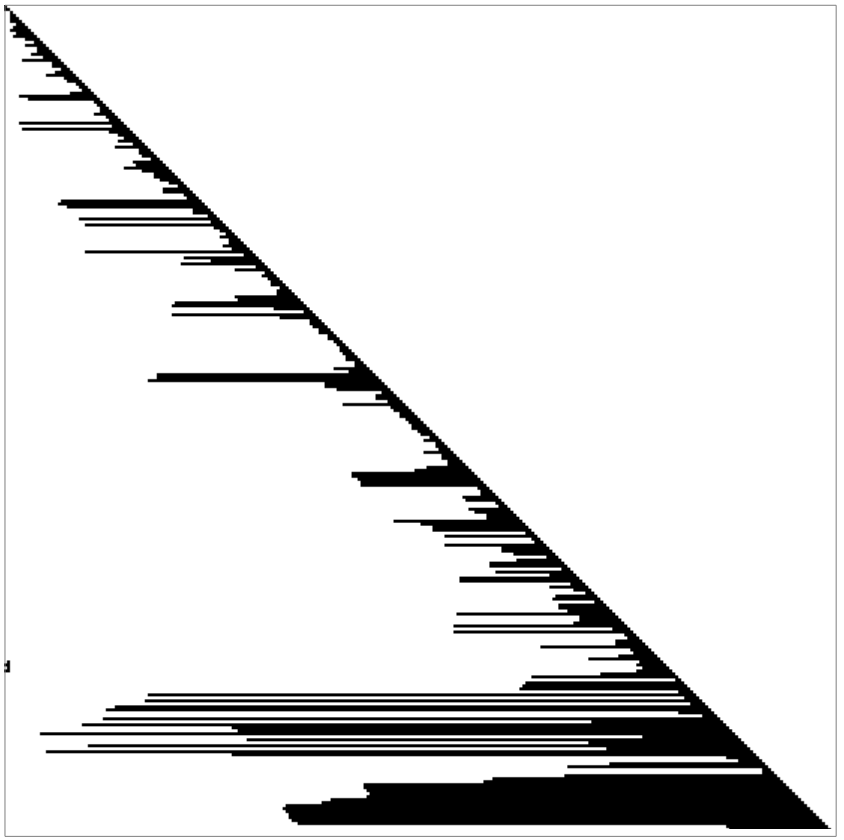
We use four strategies to reduce time and memory usage when performing this factorization on sparse matrices:

1. Reordering of rows and columns of the matrix to reduce fill-in in $\mathbf{L}$. This is equivalent to use a permutation matrix to reorder the system $\left(\boldsymbol{P}\,\boldsymbol{A}\,\boldsymbol{P}^{\mathrm{T}}\right)\left(\boldsymbol{P}\,\boldsymbol{x}\right)=\left(\boldsymbol{P}\,\boldsymbol{b}\right)$.
2. Use symbolic Cholesky factorization to obtain an exact $\mathbf{L}$ factor (non zero entries in $\mathbf{L}$).
3. Organize operations to improve cache usage.
4. Parallelize the factorization.

Matrix reordering

We want to reorder rows and columns of $\boldsymbol{A}$, in a way that the number of non-zero entries of $\boldsymbol{L}$ are reduced. $\eta(\boldsymbol{L})$ indicates the number of non-zero entries of $\boldsymbol{L}$.

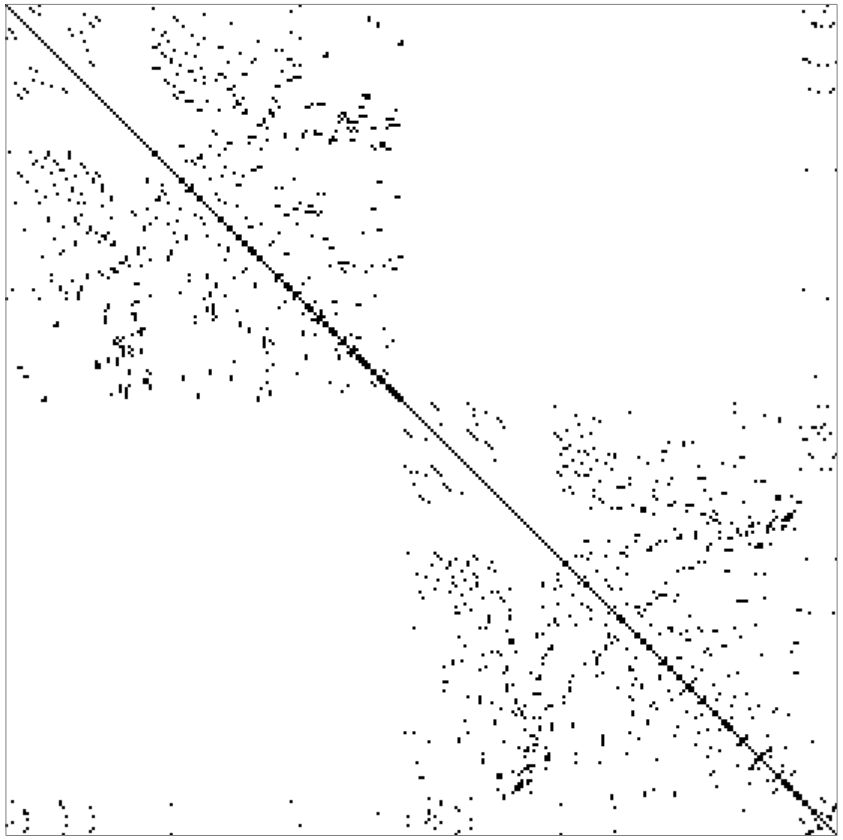
$\boldsymbol{A}=$ 

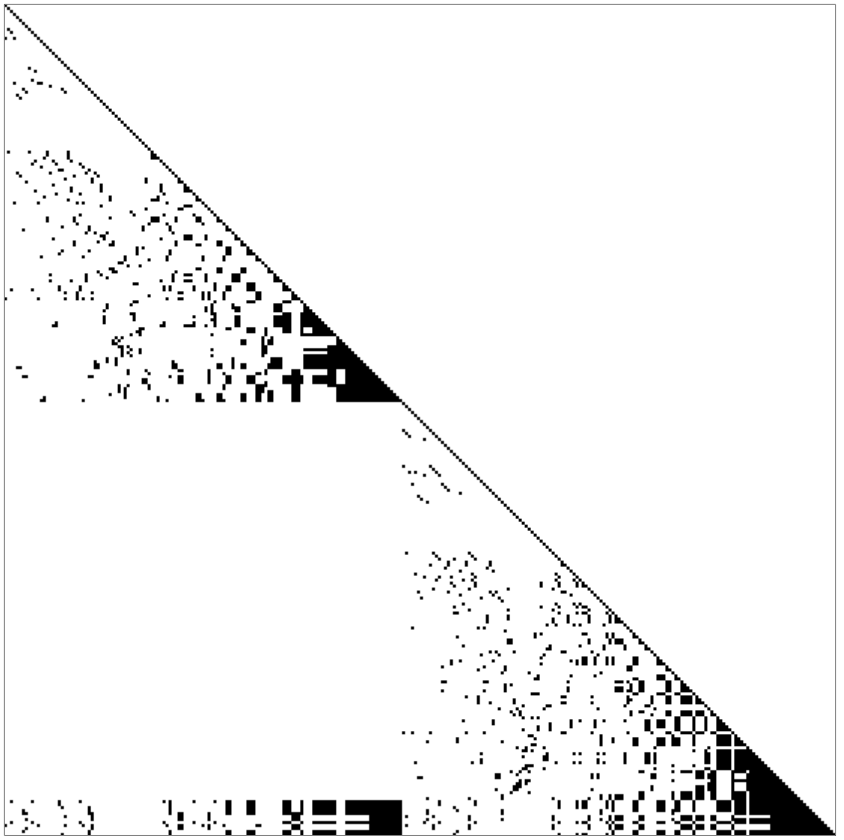
$\boldsymbol{L}=$ 

The stiffness matrix to the left $\boldsymbol{A}\in\mathbb{R}^{556\times 556}$, with $\eta(\boldsymbol{A})=1810$. To the right the lower triangular matrix $\boldsymbol{L}$, with $\eta(\boldsymbol{L})=8729$.

There are several heuristics like the minimum degree algorithm [Geor81] or a nested dissection method [Kary99].

By reordering we have a matrix A' with $\eta(A')=1810$ and its factorization L' with $\eta(L')=3215$. Both factorizations solve the same system of equations.

$A' =$ 

$L' =$ 

We reduce the factorization fill-in by

$$\frac{\eta(L')=3215}{\eta(L)=8729}=0.368.$$

To determine a “good” reordering for a matrix A that minimize the fill-in of L is an NP complete problem [Yann81].

Symbolic Cholesky factorization

The algorithm to determine the L_{ij} entries that area non-zero is called symbolic Cholesky factorization [Gall90].

Let be, for all columns $j=1\ldots n$,

$$\boldsymbol{a}_j \overset{\text{def}}{=} \left\{k>j \;\middle|\; A_{kj}\neq 0\right\},$$
$$\boldsymbol{l}_j \overset{\text{def}}{=} \left\{k>j \;\middle|\; L_{kj}\neq 0\right\}.$$

The sets $\boldsymbol{r}_j$ will register the columns of $\boldsymbol{L}$ which structure will affect the column j of $\boldsymbol{L}$.

$\boldsymbol{r}_j \leftarrow \emptyset,\ j \leftarrow 1\ldots n$
for $j \leftarrow 1\ldots n$
  $\boldsymbol{l}_j \leftarrow \boldsymbol{a}_j$
  for $i\in\boldsymbol{r}_j$
    $\boldsymbol{l}_j \leftarrow \boldsymbol{l}_j\cup\boldsymbol{l}_i\setminus\{j\}$
  end_for
  $p \leftarrow \begin{cases}\min\{i\in\boldsymbol{l}_j\} & \text{si } \boldsymbol{l}_j\neq\emptyset\\ j & \text{otro caso}\end{cases}$
  $\boldsymbol{r}_p \leftarrow \boldsymbol{r}_p\cup\{j\}$
end_for

$$\boldsymbol{A}=\begin{pmatrix} a_{11} & a_{12} & & & & a_{16}\\ a_{21} & a_{22} & a_{23} & a_{24} & & \\ & a_{32} & a_{33} & & a_{35} & \\ & a_{42} & & a_{44} & & \\ & & a_{53} & & a_{55} & a_{56}\\ a_{61} & & & & a_{65} & a_{66}\end{pmatrix} \qquad \boldsymbol{L}=\begin{pmatrix} l_{11} & & & & & \\ l_{21} & l_{22} & & & & \\ & l_{32} & l_{33} & & & \\ & l_{42} & l_{43} & l_{44} & & \\ & & l_{53} & l_{54} & l_{55} & \\ l_{61} & l_{62} & l_{63} & l_{64} & l_{65} & l_{66}\end{pmatrix}$$

$\boldsymbol{a}_2=\{3,4\}$ $\qquad$ $\boldsymbol{l}_2=\{3,4,6\}$

This algorithm is very efficient, its complexity in time and space has an order of $O\left(\eta\left(\boldsymbol{L}\right)\right)$.

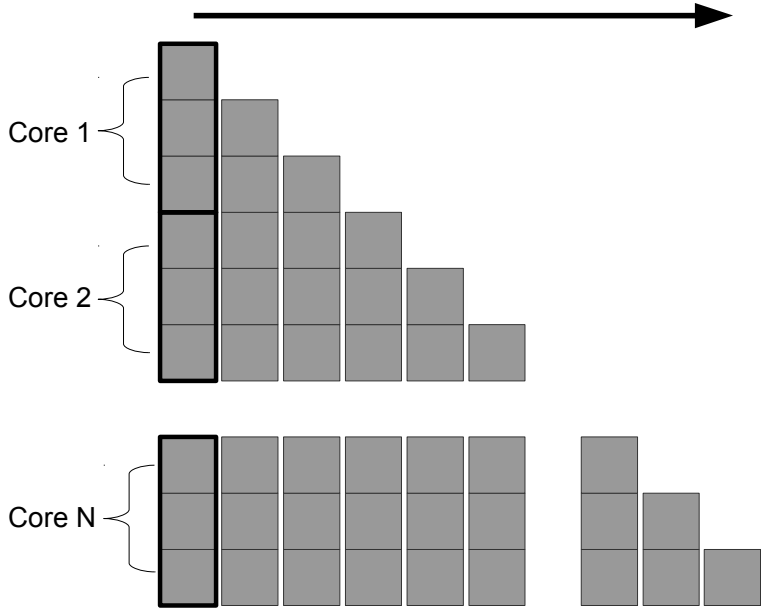
Parallelization of the factorization

The calculation of the non-zero L_{ij} entries can be done in parallel if we fill $\boldsymbol{L}$ column by column [Heat91].

Let $J(i)$ be the indexes of the non-zero values of the row i of $\boldsymbol{L}$. Formulae to calculate L_{ij} are:

$$L_{ij} = \frac{1}{L_{jj}} \left(A_{ij} - \sum_{\substack{k \in (J(i) \cap J(j)) \\ k<j}} L_{ik} L_{jk} \right), \text{ for } i > j$$

$$L_{jj} = \sqrt{A_{jj} - \sum_{\substack{k \in J(j) \\ k<j}} L_{jk}^2}.$$



The paralellization was made using the OpenMP schema.

LU factorization for sparse matrices

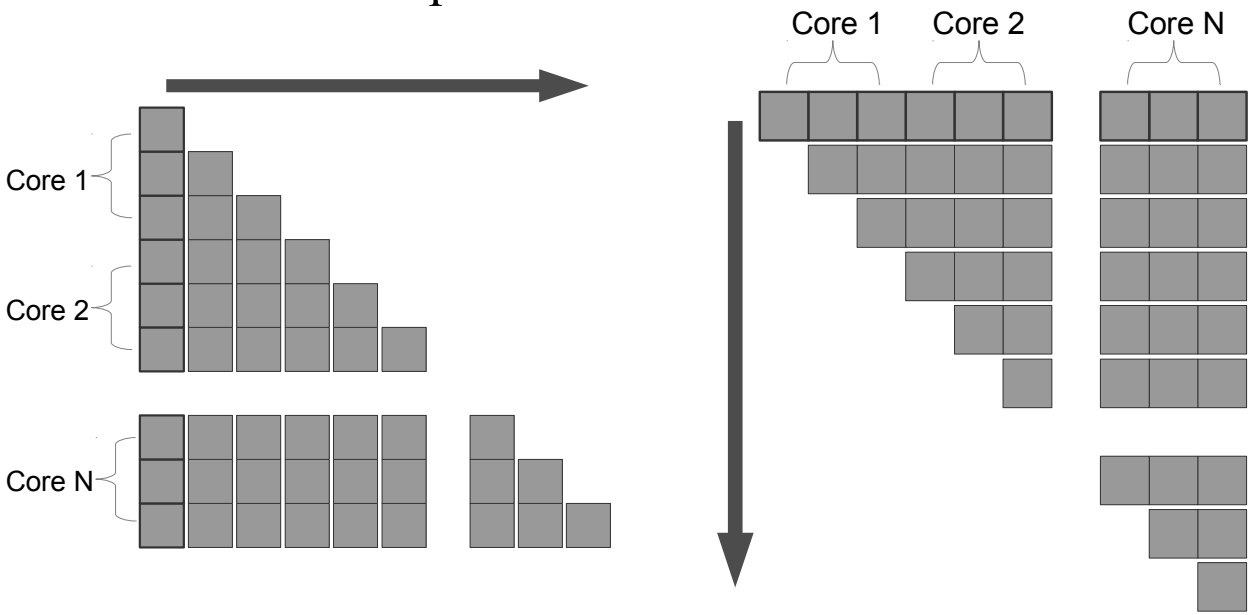
Symbolic Cholesky factorization could be use to determine the structure of the LU factorization if the matrix has symmetric structure,

$$U_{ij} = A_{ij} - \sum_{\substack{k \in (J(i) \cap J(j)) \\ k<j}} L_{ik} U_{jk} \quad \text{for } i > j,$$

$$L_{ji} = \frac{1}{U_{ii}} \left(A_{ji} - \sum_{\substack{k \in (J(j) \cap J(i)) \\ k<i}} L_{jk} U_{ik} \right) \quad \text{for } i > j,$$

$$U_{ii} = A_{ii} - \sum_{\substack{k \in J(i) \\ k<i}} L_{ik} U_{ik}, \; L_{ii} = 1.$$

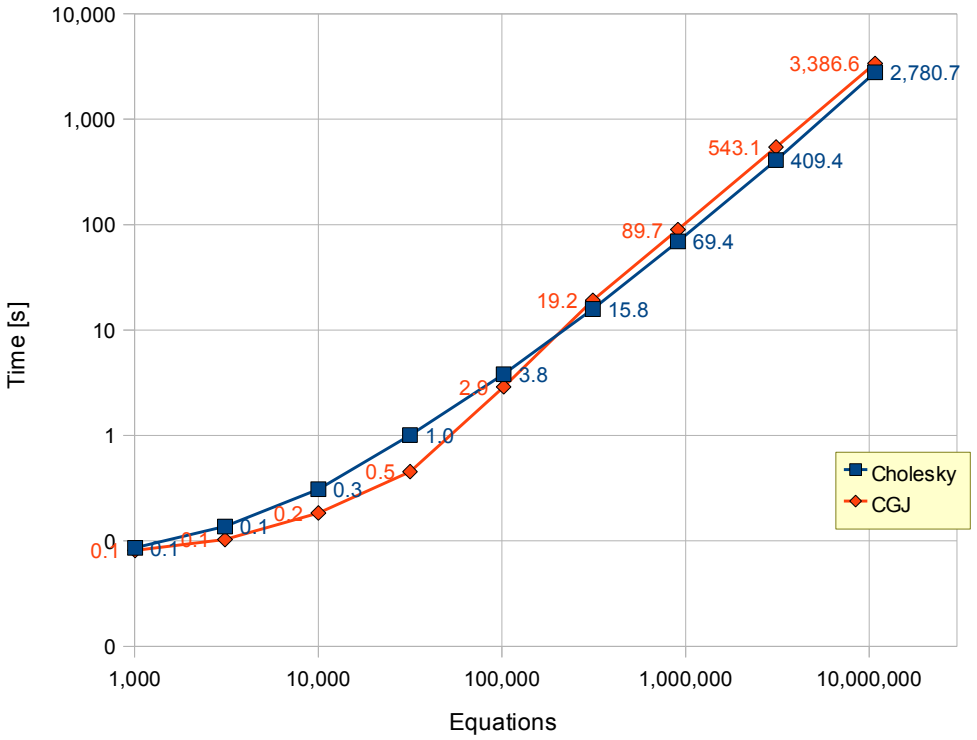
Filling of $\boldsymbol{L}$ and $\boldsymbol{U}$ can also be done in parallel.

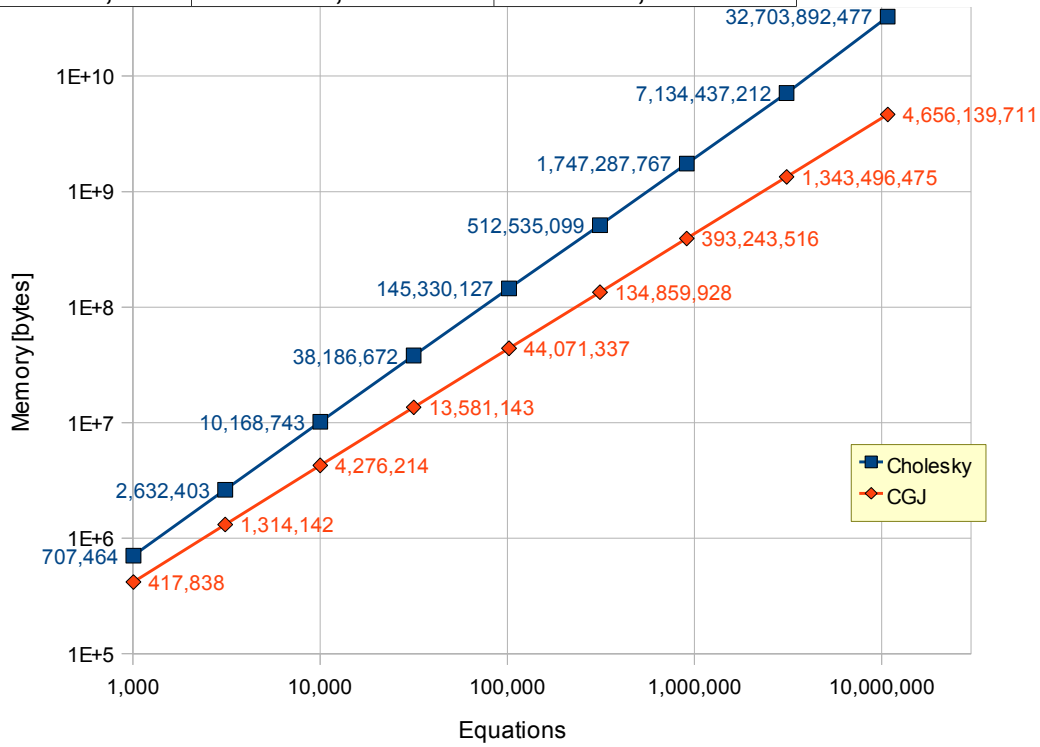


How efficient is it?

The next table shows results solving a 2D Poisson equation problem, comparing Cholesky and conjugate gradient with Jacobi preconditioning. Several discretizations where used.

Equations	nnz(A)	nnz(L)	Cholesky [s]	CGJ [s]
1,006	6,140	14,722	0.086	0.081
3,110	20,112	62,363	0.137	0.103
10,014	67,052	265,566	0.309	0.184
31,615	215,807	1’059,714	1.008	0.454
102,233	705,689	4’162,084	3.810	2.891
312,248	2’168,286	14’697,188	15.819	19.165
909,540	6’336,942	48’748,327	69.353	89.660
3’105,275	21’681,667	188’982,798	409.365	543.110
10’757,887	75’202,303	743’643,820	2,780.734	3,386.609





Parallel preconditioned conjugate gradient

For large and sparse systems of equations it is necessary to choose preconditioners that are also sparse.

$\boldsymbol{x}_0$, initial approximation
$\boldsymbol{r}_0 \leftarrow \boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_0$, initial gradient
$\boldsymbol{q}_0 \leftarrow \boldsymbol{M}\boldsymbol{r}_0$
$\boldsymbol{p}_0 \leftarrow \boldsymbol{q}_0$, initial descent direction
$k \leftarrow 0$
while $\|\boldsymbol{r}_k\| > \varepsilon$

$$\alpha_k \leftarrow -\frac{\boldsymbol{r}_k^{\mathrm{T}}\boldsymbol{q}_k}{\boldsymbol{p}_k^{\mathrm{T}}\boldsymbol{A}\boldsymbol{p}_k}$$

$$\boldsymbol{x}_{k+1} \leftarrow \boldsymbol{x}_k + \alpha_k \boldsymbol{p}_k$$

$$\boldsymbol{r}_{k+1} \leftarrow \boldsymbol{r}_k - \alpha_k \boldsymbol{A}\boldsymbol{p}_k$$

$$\boldsymbol{q}_{k+1} \leftarrow \boldsymbol{M}\boldsymbol{r}_{k+1}$$

$$\beta_k \leftarrow \frac{\boldsymbol{r}_{k+1}^{\mathrm{T}}\boldsymbol{q}_{k+1}}{\boldsymbol{r}_k^{\mathrm{T}}\boldsymbol{q}_k}$$

$$\boldsymbol{p}_{k+1} \leftarrow \boldsymbol{q}_{k+1} + \beta_k \boldsymbol{p}_k$$

$$k \leftarrow k+1$$

Jacobi preconditioner

The diagonal part of $\boldsymbol{M}^{-1}$ is stored as a vector,

$$\boldsymbol{M}^{-1} = \left(\operatorname{diag}\left(\boldsymbol{A}\right)\right)^{-1}.$$

Parallelization of this preconditioner is straightforward, because the calculus of each entry of $\boldsymbol{q}_k$ is independent.

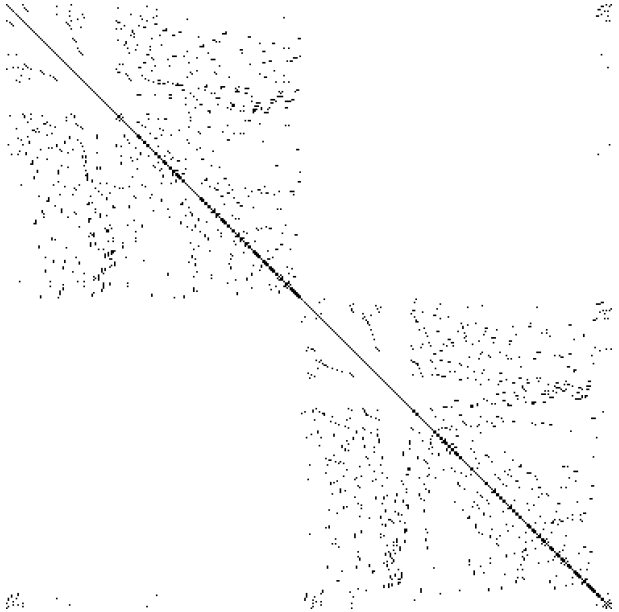
Incomplete Cholesky factorization preconditioner

This preconditioner has the form

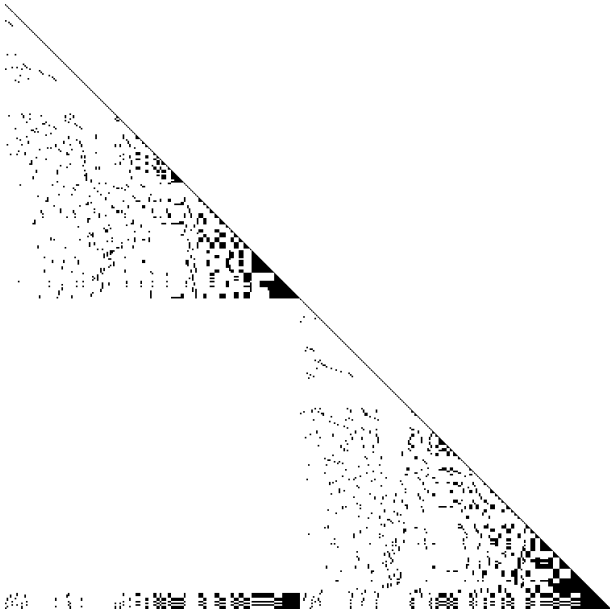
$$\boldsymbol{M}^{-1}=\boldsymbol{G}_l\boldsymbol{G}_l^{\mathrm{T}},$$

where $\boldsymbol{G}_l$ is a lower triangular sparse matrix that have a structure similar to the Cholesky factorization of $\boldsymbol{A}$.

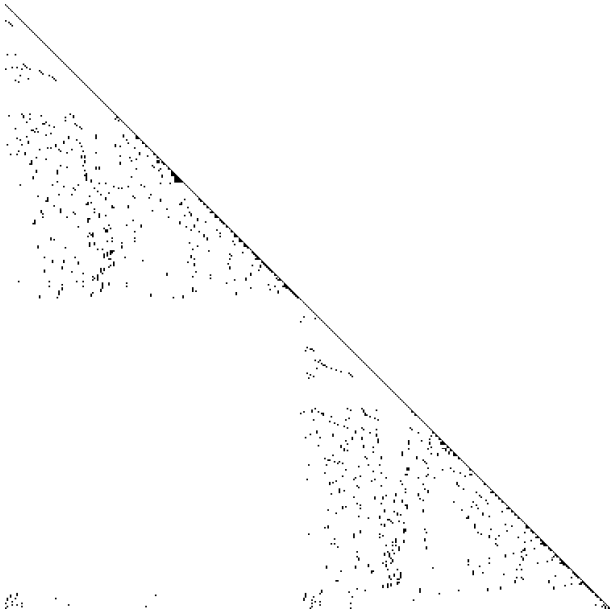
- The structure of $\boldsymbol{G}_0$ is equal to the structure of the lower triangular form of $\boldsymbol{A}$.
- The structure of $\boldsymbol{G}_m$ is equal to the structure of $\boldsymbol{L}$ (complete Cholesky factorization of $\boldsymbol{A}$).
- For $0<l<m$ the structure of $\boldsymbol{G}_l$ is creating having a number of entries between $\boldsymbol{L}$ and the lower triangular form of $\boldsymbol{A}$, making easy to control the sparsity of the preconditioner.



$\mathbf{A}$



$\mathbf{G}_m=\mathbf{L}$



$\mathbf{G}_{10}$

Entries are filled using

$$G_{ij} = \frac{1}{G_{jj}} \left(A_{ij} - \sum_{\substack{k \in (J(i) \cap J(j)) \\ k<j}} G_{ik} G_{jk} \right), \text{ for } i > j$$

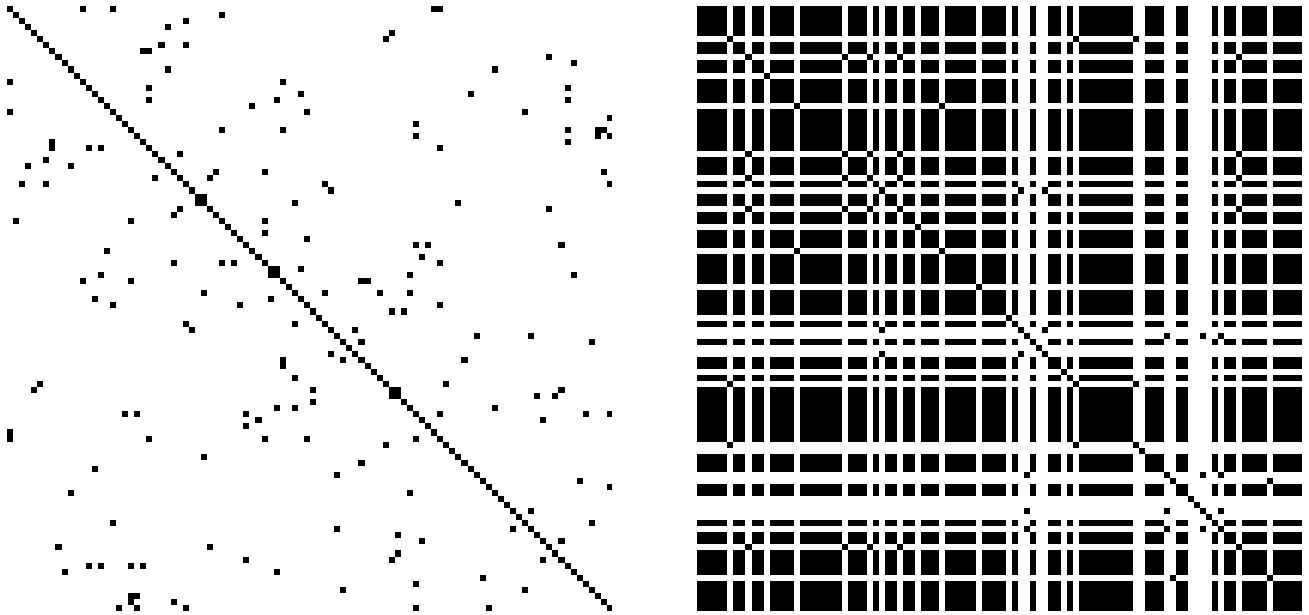
$$G_{jj} = \sqrt{A_{jj} - \sum_{\substack{k \in J(j) \\ k<j}} G_{jk}^2}.$$

This preconditioner is not always stable [Golu96 p535].

The use of this preconditioner implies to solve a system of equations in each CG step using a backward and a forward substitution algorithm, this operations are fast given the sparsity of $\boldsymbol{G}_l$. Unfortunately the dependency of values makes these substitutions very hard to parallelize.

Factorized sparse approximate inverse preconditioner

The aim of this preconditioner is to construct $\boldsymbol{M}$ to be an approximation of the inverse of $\boldsymbol{A}$ with the property of being sparse. The inverse of a sparse matrix is not necessary sparse.



A way to create an approximate inverse is to minimize the Frobenius norm of the residual $\boldsymbol{I}-\boldsymbol{A}\boldsymbol{M}$,

$$F(\boldsymbol{M})=\|\boldsymbol{I}-\boldsymbol{A}\boldsymbol{M}\|_{\mathrm{F}}^2.$$

The Frobenius norm is defined as

$$\|\boldsymbol{A}\|_{\mathrm{F}}=\sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n}|a_{ij}|^2}=\sqrt{\mathrm{tr}\left(\boldsymbol{A}^{\mathrm{T}}\boldsymbol{A}\right)}.$$

It is possible to separate $F(\boldsymbol{M}) = \|\boldsymbol{I} - \boldsymbol{A}\boldsymbol{M}\|_{\mathrm{F}}^2$ into decoupled sums of 2-norms for each column [Chow98],

$$F(\boldsymbol{M}) = \|\boldsymbol{I} - \boldsymbol{A}\boldsymbol{M}\|_{\mathrm{F}}^2 = \sum_{j=1}^{n} \|\boldsymbol{e}_j - \boldsymbol{A}\boldsymbol{m}_j\|_2^2,$$

where $\boldsymbol{e}_j$ is the j-th column of $\boldsymbol{I}$ and $\boldsymbol{m}_j$ is the j-th column of $\boldsymbol{M}$. With this separation we can parallelize the construction of the preconditioner.

The factorized sparse approximate inverse preconditioner [Chow01] creates a preconditioner

$$\boldsymbol{M} = \boldsymbol{G}_l^{\mathrm{T}} \boldsymbol{G}_l,$$

where $\boldsymbol{G}$ is a lower triangular matrix such that

$$\boldsymbol{G}_l \approx \boldsymbol{L}^{-1},$$

where $\boldsymbol{L}$ is the Cholesky factor of $\boldsymbol{A}$. l is a positive integer that indicates a level of sparsity of the matrix.

Instead of minimizing $F(\boldsymbol{M}) = \|\boldsymbol{I} - \boldsymbol{A}\boldsymbol{M}\|_{\mathrm{F}}^2$, we minimize $\|\boldsymbol{I} - \boldsymbol{G}_l\boldsymbol{L}\|_{\mathrm{F}}^2$, it is noticeable that this can be done without knowing $\boldsymbol{L}$.

This preconditioner has these features:

- $\boldsymbol{M}$ is SPD if there are no zeroes in the diagonal of $\boldsymbol{G}_l$.
- The algorithm to construct the preconditioner is parallelizable.
- This algorithm is stable if $\boldsymbol{A}$ is SPD.

Parallel preconditionated biconjugated gradient

The algorithm is [Meie92]:

ε, tolerance

$\boldsymbol{x}_0$, initial coordinate

$\boldsymbol{g}_0 \leftarrow \boldsymbol{A}\boldsymbol{x}_0 - \boldsymbol{b}$, initial gradient

$\tilde{\boldsymbol{g}}_0 \leftarrow \boldsymbol{g}_0^{\mathrm{T}}$, initial pseudo-gradient

$\boldsymbol{q}_0 \leftarrow \boldsymbol{M}^{-1}\boldsymbol{g}_0$

$\tilde{\boldsymbol{q}}_0 \leftarrow \tilde{\boldsymbol{g}}_0 \boldsymbol{M}^{-1}$

$\boldsymbol{p}_0 \leftarrow -\boldsymbol{q}_0$, initial descent direction

$\bar{\boldsymbol{p}}_0 \leftarrow -\tilde{\boldsymbol{q}}_0$, initial pseudo-direction of descent

$k \leftarrow 0$

while $\left\|\boldsymbol{g}_k\right\| > \varepsilon$

$\boldsymbol{w} \leftarrow \boldsymbol{A}\boldsymbol{p}_k$

$\tilde{\boldsymbol{w}} \leftarrow \tilde{\boldsymbol{p}}_k \boldsymbol{A}$

$$\alpha_k \leftarrow -\frac{\tilde{\boldsymbol{q}}_k \boldsymbol{g}_k}{\tilde{\boldsymbol{p}}_k \boldsymbol{w}}$$

$$\boldsymbol{x}_{k+1} \leftarrow \boldsymbol{x}_k + \alpha_k \boldsymbol{p}_k$$

$$\boldsymbol{g}_{k+1} \leftarrow \boldsymbol{g}_k + \alpha \boldsymbol{w}$$

$$\tilde{\boldsymbol{g}}_{k+1} \leftarrow \tilde{\boldsymbol{g}}_k + \alpha \tilde{\boldsymbol{w}}$$

$$\boldsymbol{q}_{k+1} \leftarrow \boldsymbol{M}^{-1}\boldsymbol{g}_{k+1}$$

$$\tilde{\boldsymbol{q}}_{k+1} \leftarrow \tilde{\boldsymbol{g}}_{k+1}\boldsymbol{M}^{-1}$$

$$\beta_k \leftarrow \frac{\tilde{\boldsymbol{g}}_{k+1}\boldsymbol{q}_{k+1}}{\tilde{\boldsymbol{g}}_k \boldsymbol{q}_k}$$

$$\boldsymbol{p}_{k+1} \leftarrow -\boldsymbol{q}_{k+1} + \beta_{k+1}\boldsymbol{p}_k$$

$$\tilde{\boldsymbol{p}}_{k+1} \leftarrow -\tilde{\boldsymbol{q}}_{k+1} + \beta_{k+1}\tilde{\boldsymbol{p}}_k$$

$$k \leftarrow k+1$$

Preconditioning with incomplete LU factorization

This preconditioner is analog to the incomplete Cholesky preconditioner, this one is formed by two matrices

$$\boldsymbol{M} = \boldsymbol{G}_k \boldsymbol{H}_k,$$

where $\boldsymbol{G}_k$ is a sparse lower triangular matrix, and $\boldsymbol{H}_k$ is an sparse upper triangular matrix. They have a similar structure to the factorization LU of $\boldsymbol{A} \in \mathbb{R}^{n \times n}$.

- The structure of $\boldsymbol{G}_0$ is equal to the structure of the lower triangular part of $\boldsymbol{A}$.
- The structure of $\boldsymbol{H}_0$ is equal to the structure of the upper triangular part of $\boldsymbol{A}$.
- The structure of $\boldsymbol{G}_m$ is equal to the strucure of $\boldsymbol{L}$ (from the LU factorization of $\boldsymbol{A}$.)
- The structure of $\boldsymbol{H}_m$ is equal to the strucure of $\boldsymbol{U}$ (from the LU factorization of $\boldsymbol{A}$.)
- For $0 < k < m$ the structure of $\boldsymbol{G}_k$ is created to have a number of entries between $\boldsymbol{G}_0$ and $\boldsymbol{G}_m$, making easy to control the sparsity of the preconditioner. Similarly for $\boldsymbol{H}_k$.

Values of $\boldsymbol{G}_k$ and $\boldsymbol{H}_k$ are filled using the formulae to calculate the LU factorization

$$H_{ij} = A_{ij} - \sum_{k=1}^{j-1} G_{ik} H_{kj}, \; G_{ji} = \frac{1}{H_{ii}} \left(A_{ji} - \sum_{k=1}^{i-1} G_{jk} H_{ki} \right) \text{ for } i > j,$$

$$H_{ii} = A_{ii} - \sum_{k=1}^{i-1} G_{ik} H_{ki}, \; G_{ii} = 1.$$

Approximate inverse preconditioner

An appoximate inverse for non-simmetric matrices can be constructed using the minimal resiual [Chow98]. It is constructed by minimizing the Frobenius norm of the residual $\boldsymbol{I}-\boldsymbol{A}\boldsymbol{M}$,

$$F(\boldsymbol{M})=\|\boldsymbol{I}-\boldsymbol{A}\boldsymbol{M}\|_{\mathrm{F}}^2.$$

The algorithm to build $\boldsymbol{M}$ is:

Let $\boldsymbol{M}=\boldsymbol{M}_0$
for $j \leftarrow 1 \ldots n$
  $\boldsymbol{m}_j \leftarrow \boldsymbol{M}\boldsymbol{e}_j$
  for $i \leftarrow 1 \ldots s$
    $\boldsymbol{r}_j \leftarrow \boldsymbol{e}_j-\boldsymbol{A}\boldsymbol{m}_j$
    $\alpha_j \leftarrow \dfrac{\boldsymbol{r}_j^{\mathrm{T}}\boldsymbol{A}\boldsymbol{r}_j}{\left(\boldsymbol{A}\boldsymbol{r}_j\right)^{\mathrm{T}}\boldsymbol{A}\boldsymbol{r}_j}$
    $\boldsymbol{m}_j \leftarrow \boldsymbol{m}_j+\alpha_j\boldsymbol{r}_j$
    Truncate less significative entries of $\boldsymbol{m}_j$

This is an iterative method, here s is small. If no sparse vectors are used, then this algorithm will have an order of $O\left(n^2\right)$.

Approximate inverse via minimal residual iteration with dropping in the search direction algorithm is shown next.

Let $\boldsymbol{M} = \boldsymbol{M}_0$
for $j \leftarrow 1 \ldots n$
  Let $\boldsymbol{m}_j \leftarrow \boldsymbol{M}\,\boldsymbol{e}_j$ be an sparse vector
  Let $\boldsymbol{d}_j$ a sparse vector with the same strucure than $\boldsymbol{m}_j$
  for $i \leftarrow 1 \ldots s$
    $\boldsymbol{r}_j \leftarrow \boldsymbol{e}_j - \boldsymbol{A}\,\boldsymbol{m}_j$
    $\boldsymbol{d}_j \leftarrow \boldsymbol{r}_j$ only take the entries of $\boldsymbol{r}_j$ with index in $\boldsymbol{d}_j$
    $$\alpha_j \leftarrow \frac{\boldsymbol{r}_j^{\mathrm{T}}\boldsymbol{A}\,\boldsymbol{d}_j}{\left(\boldsymbol{A}\,\boldsymbol{d}_j\right)^{\mathrm{T}}\boldsymbol{A}\,\boldsymbol{d}_j}$$
    $\boldsymbol{m}_j \leftarrow \boldsymbol{m}_j + \alpha_j \boldsymbol{d}_j$
    if $\eta\left(\boldsymbol{m}_j\right) < lfil$
      add $\max\left|\left(\boldsymbol{r}_j\right)_k\right|$, such that k does not exists in the structure of $\boldsymbol{m}_j$

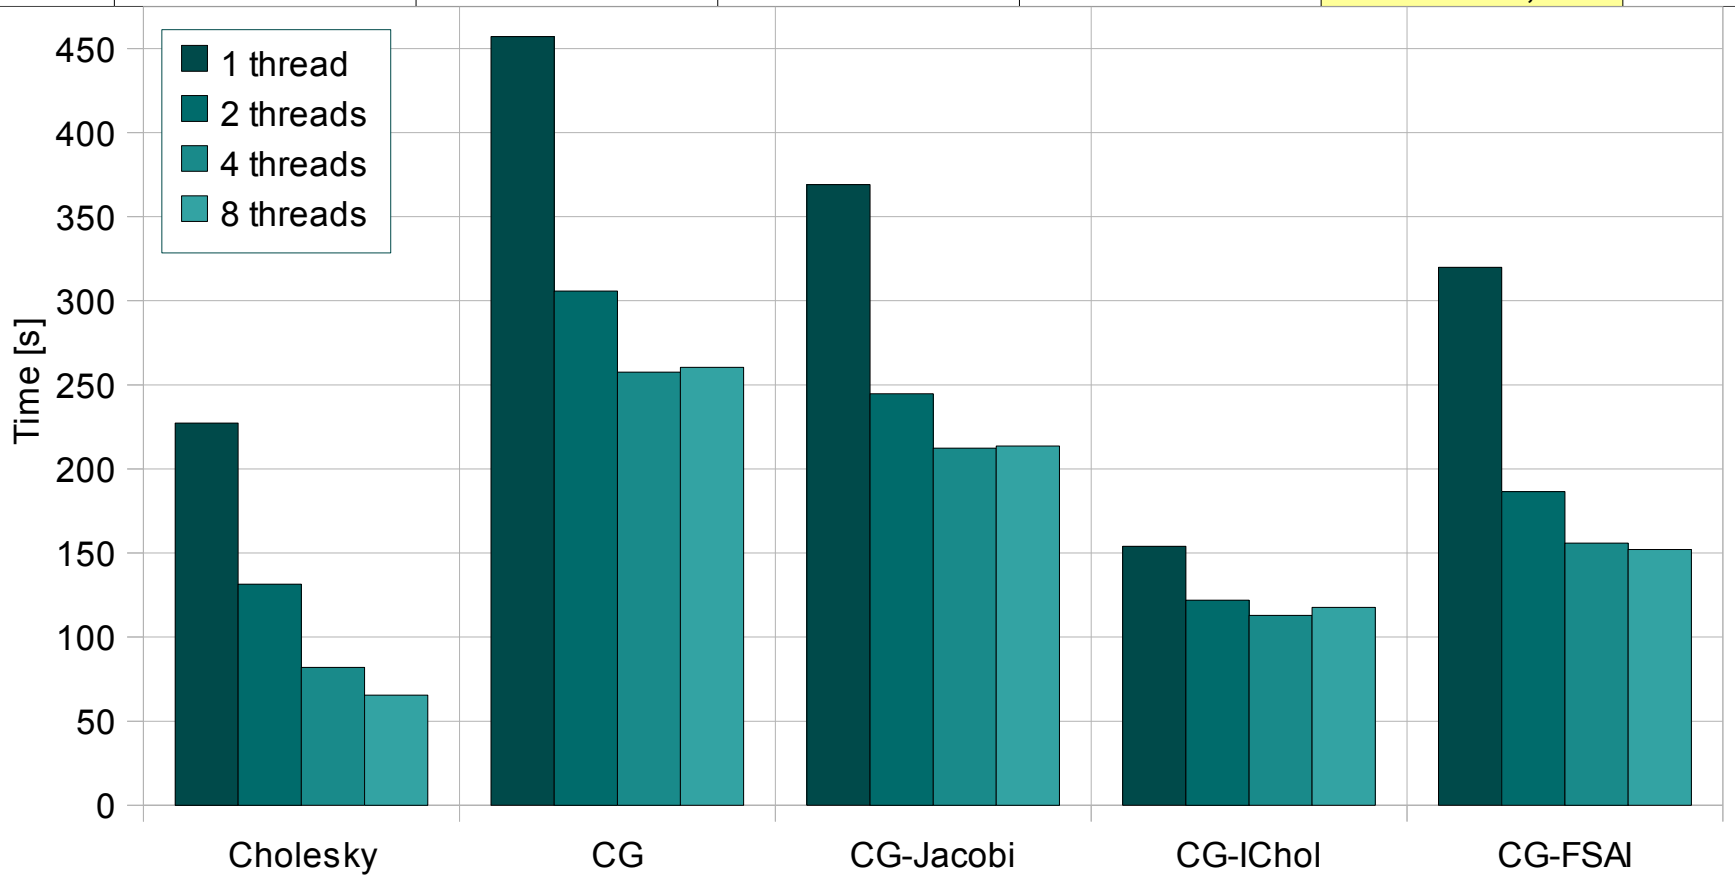
Cons:

- The final structure of $\boldsymbol{M}$ is not symetric.
- $\boldsymbol{M}$ can be singular if s is small.

Performance comparisons

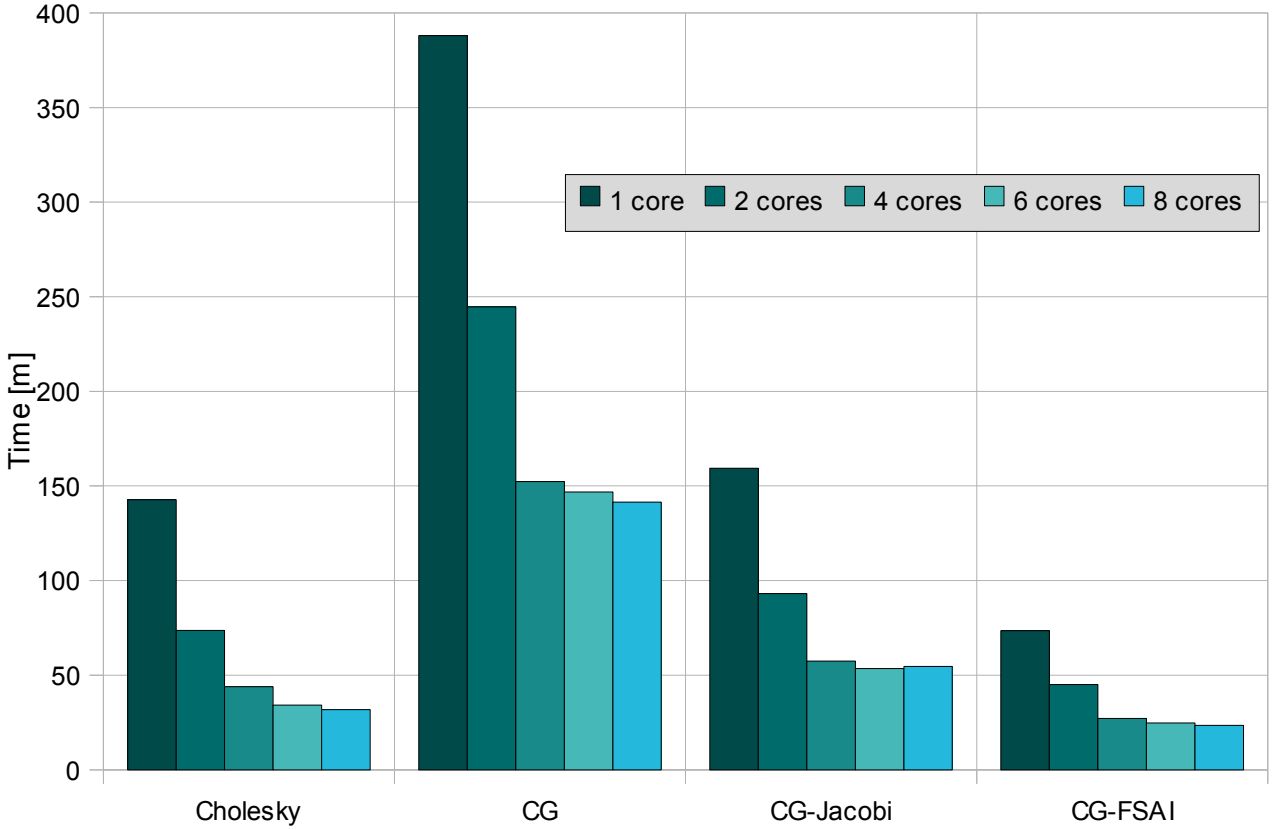
2D solid linear deformation, 501,264 elements, con 909,540 equations, $\eta(A) = 18'062{,}500$.

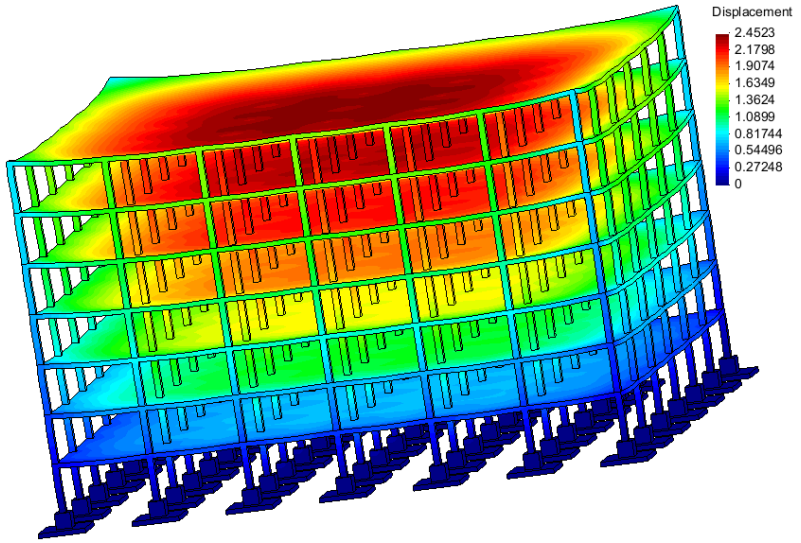
Solver	1 thread [s]	2 threads [s]	4 threads [s]	8 threads [s]	Steps	Memory
Cholesky	227.2	131.4	81.9	65.4		3,051,144,550
CG	457.0	305.8	257.5	260.4	9,251	317,929,450
CG-Jacobi	369.0	244.7	212.4	213.7	6,895	325,972,366
CG-IChol	153.9	121.9	112.8	117.6	1,384	586,380,322
CG-FSAI	319.8	186.5	155.7	152.1	3,953	430,291,930



3D solid deformed by self-weight

Elements:	264,250
Nodes:	326,228
Equations:	978,684
nnz(K):	69,255,522

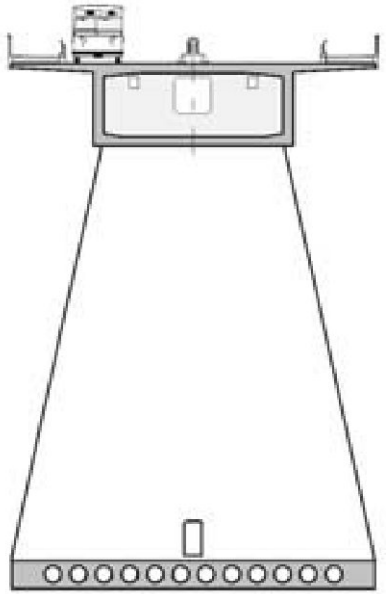




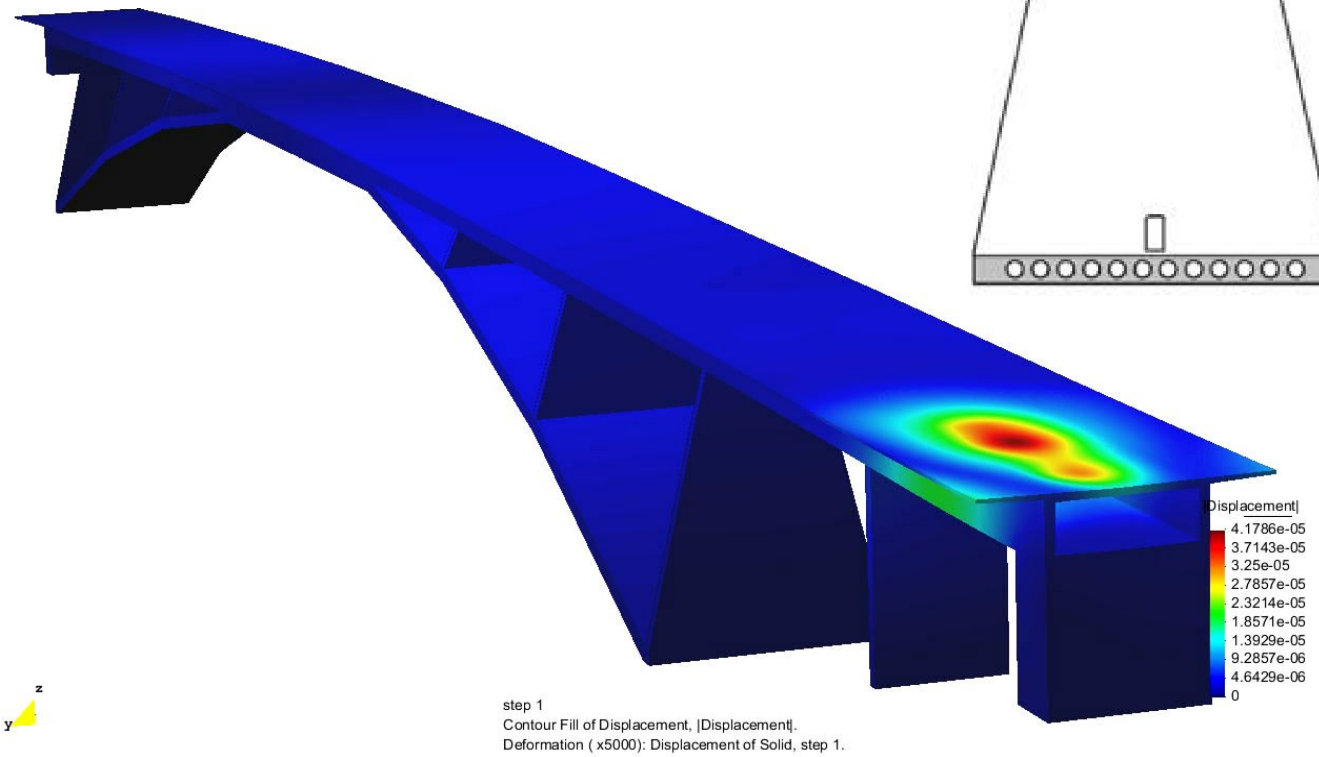
Solver	1 core [m]	2 cores [m]	4 cores [m]	6 cores [m]	8 cores [m]	Memory
Cholesky	142	73	43	34	31	19,864’132,056
CG	387	244	152	146	141	922’437,575
CG-Jacobi	159	93	57	53	54	923’360,936
CG-FSAI	73	45	27	24	23	1,440’239,572

Infante Henrique bridge over the Douro river, Portugal

Simulation of a 18 wheels 36 metric tons truck crossing the Infante D. Henrique Bridge.



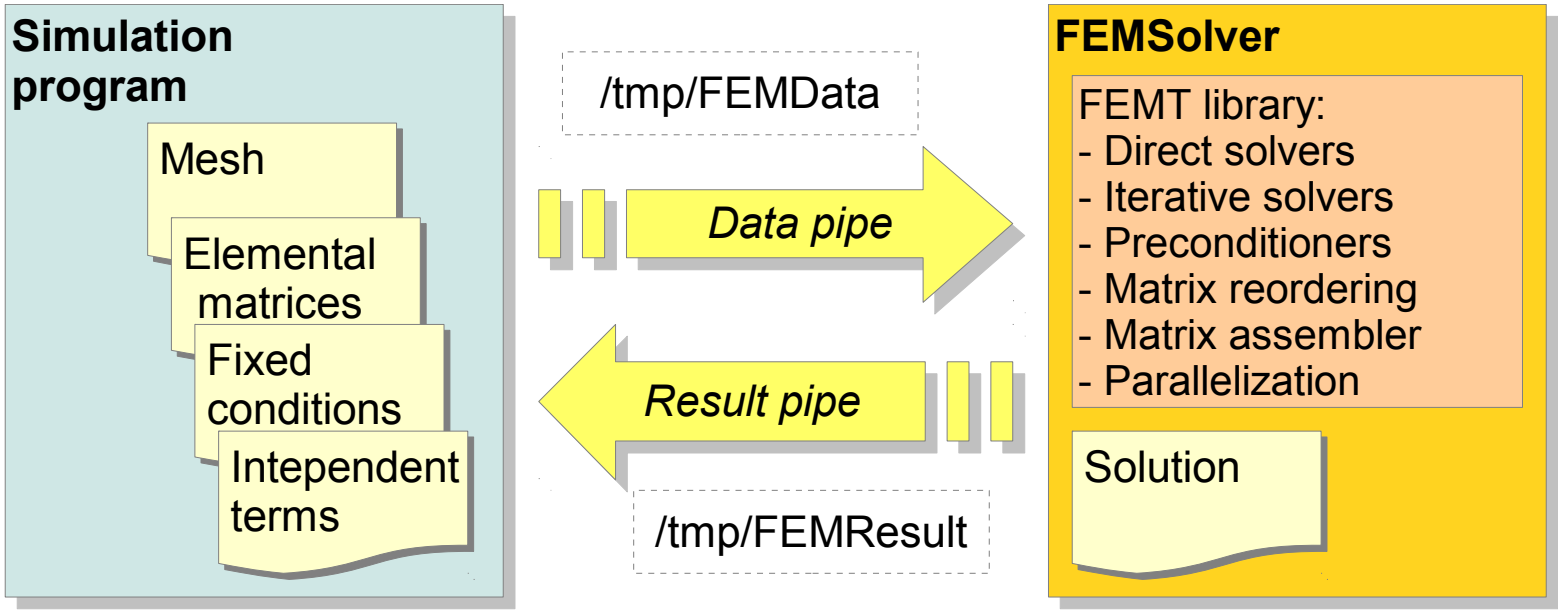
Nodes	337,195
Elements	1’413,279
Element type	Tetrahedron
HHT alpha factor	0
Rayleigh damping a	0.5
Rayleigh damping b	0.5
Degrees of freedom	1’011,585
nnz(K)	38’104,965
Time to assemble K	4.5 s
Time to reorder K	32.4 s
Factorization time	178.8 s
Time steps	372
Time per step	2.6 s
Total time	1205.1 s



Peak allocated memory: 9,537’397,868 bytes
Computer: 2 x Intel(R) Xeon(R) CPU E5620, 8 cores, 12MB cache, 32 GB of RAM

FEMSolver

FEMSolver is a program that solves finite element problems in parallel using the FEMT library on multi-core computers. It uses a very simple interface using pipes. A pipe is an object of the operationg system that can be accessed like a file but does not write data to the disk, is a fast way to communicate running programs.



This flexible schema allows an used using any programming languaje (C/C++, Fortran, Python, C#, Java, etc.) to solve large systems of equations resulting from finite element discretizations. It can use any of the solvers of the FEMT library.

If the matrix remains constant, FEMSolver can be used to efficiently solve multi-step problems, like dynamic deformations, transient heat diffusion, etc.

```
      PROGRAM FEMSolverExample
      IMPLICIT NONE
C Commands
      INTEGER command_end, command_set_connectivity, command_fill_A, command_se
      PARAMETER (
     . command_end           = 0,
     . command_set_connectivity = 1,
     . command_fill_A        = 2,
     . command_set_Ae        = 3,
     . command_set_all_Ae = 4,
     . command_set_x         = 5,
     . command_set_b         = 6,
     . command_set_fixed     = 7,
     . command_solver_init = 8,
     . command_solver_run = 9)
C Number of elements
      INTEGER*4 E /21/
C Number of nodes
      INTEGER*4 M /17/
C Element type (2=Triangle, 3=Quadrilateral, 4=Tetrahedra, 5=Hexahedra)
      INTEGER*4 T /2/
C Nodes per element
      INTEGER*4 N /3/
C Degrees of freedom
      INTEGER*4 D /1/
C Connectivity
      INTEGER*4 connectivity(3*21) /
     . 3, 8, 5,
     . 3, 5, 1,
     . 5, 8, 12,
     . 7, 2, 4,
     . 7, 4, 11,
     . 4, 2, 1,
     . 17, 16, 13,
     . 17, 13, 15,
     . 13, 16, 14,
     . 12, 15, 10,
     . 10, 15, 13,
     . 12, 10, 5,
```

```
#include <stdio.h>
int main()
{
// Commands
    enum Command {
        command_end           = 0,
        command_set_connectivity = 1,
        command_fill_A        = 2,
        command_set_Ae        = 3,
        command_set_all_Ae  = 4,
        command_set_x         = 5,
        command_set_b         = 6,
        command_set_fixed     = 7,
        command_solver_init  = 8,
        command_solver_run  = 9};
// Number of elements
    int E = 21;
// Number of nodes
    int M = 17;
// Element type (2=Triangle, 3=Quadrilateral, 4=Tetrahedra, 5=Hexahedra)
    int T = 2;
// Nodes per element
    int N = 3;
// Degrees of freedom
    int D = 1;
// Connectivity
    int connectivity[21*3] = {
        3, 8, 5,
        3, 5, 1,
        5, 8, 12,
        7, 2, 4,
        7, 4, 11,
        4, 2, 1,
        17, 16, 13,
        17, 13, 15,
        13, 16, 14,
        12, 15, 10,
        10, 15, 13,
        12, 10, 5,
```

```
   . 14, 11, 9,
   . 9, 11, 4,
   . 14, 9, 13,
   . 10, 13, 6,
   . 6, 13, 9,
   . 10, 6, 5,
   . 6, 9, 4,
   . 5, 6, 4,
   . 1, 5, 4/
C Elemental matrices
      REAL*8 Ke(3*3,21) /
   . 6.47e-5, -2.62e-5, -3.85e-5, -2.62e-5, 6.47e-5, -3.85e-5, -3.85e-5, -3.85e-5, 7.70e-5
   . 1.19e-4, -6.49e-5, -5.46e-5, -6.49e-5, 6.45e-5, 3.33e-7, -5.46e-5, 3.33e-7, 5.42e-5
   . 6.45e-5, -6.49e-5, 3.33e-7, -6.49e-5, 1.19e-4, -5.46e-5, 3.33e-7, -5.46e-5, 5.42e-5
   . 6.47e-5, -2.62e-5, -3.85e-5, -2.62e-5, 6.47e-5, -3.85e-5, -3.85e-5, -3.85e-5, 7.70e-5
   . 1.19e-4, -6.49e-5, -5.46e-5, -6.49e-5, 6.45e-5, 3.33e-7, -5.46e-5, 3.33e-7, 5.42e-5
   . 6.45e-5, -6.49e-5, 3.33e-7, -6.49e-5, 1.19e-4, -5.46e-5, 3.33e-7, -5.46e-5, 5.42e-5
   . 6.47e-5, -2.62e-5, -3.85e-5, -2.62e-5, 6.47e-5, -3.85e-5, -3.85e-5, -3.85e-5, 7.70e-5
   . 1.19e-4, -6.49e-5, -5.46e-5, -6.49e-5, 6.45e-5, 3.33e-7, -5.46e-5, 3.33e-7, 5.42e-5
   . 6.45e-5, -6.49e-5, 3.33e-7, -6.49e-5, 1.19e-4, -5.46e-5, 3.33e-7, -5.46e-5, 5.42e-5
   . 6.47e-5, -2.62e-5, -3.85e-5, -2.62e-5, 6.47e-5, -3.85e-5, -3.85e-5, -3.85e-5, 7.70e-5
   . 1.18e-4, -6.49e-5, -5.33e-5, -6.49e-5, 6.52e-5, -3.33e-7, -5.33e-5, -3.33e-7, 5.37e-5
   . 6.52e-5, -6.49e-5, -3.33e-7, -6.49e-5, 1.18e-4, -5.33e-5, -3.33e-7, -5.33e-5, 5.37e-5
   . 6.47e-5, -2.62e-5, -3.85e-5, -2.62e-5, 6.47e-5, -3.85e-5, -3.85e-5, -3.85e-5, 7.70e-5
   . 1.18e-4, -6.49e-5, -5.33e-5, -6.49e-5, 6.52e-5, -3.33e-7, -5.33e-5, -3.33e-7, 5.37e-5
   . 6.52e-5, -6.49e-5, -3.33e-7, -6.49e-5, 1.18e-4, -5.33e-5, -3.33e-7, -5.33e-5, 5.37e-5
   . 6.44e-5, -2.56e-5, -3.89e-5, -2.56e-5, 6.44e-5, -3.89e-5, -3.89e-5, -3.89e-5, 7.78e-5
   . 7.67e-5, -3.96e-5, -3.72e-5, -3.96e-5, 6.60e-5, -2.64e-5, -3.72e-5, -2.64e-5, 6.36e-5
   . 6.60e-5, -3.96e-5, -2.64e-5, -3.96e-5, 7.67e-5, -3.72e-5, -2.64e-5, -3.72e-5, 6.36e-5
   . 7.49e-5, -3.85e-5, -3.65e-5, -3.85e-5, 6.64e-5, -2.80e-5, -3.65e-5, -2.80e-5, 6.45e-5
   . 6.05e-5, -6.76e-5, 7.08e-6, -6.76e-5, 1.33e-4, -6.58e-5, 7.08e-6, -6.58e-5, 5.87e-5
   . 7.18e-5, -3.59e-5, -3.59e-5, -3.59e-5, 6.67e-5, -3.07e-5, -3.59e-5, -3.07e-5, 6.67e-5/
C Vector with constrain values
      REAL*8 x(17) /0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 273, 0, 273, 0/
C Vector of independent terms
      REAL*8 b(17) /2.82e-4, 2.82e-4, 0, 0, 0, 0, 2.82e-4, 2.82e-4, 0, 0, 2.82e-4, 2.82e-4, 0,
C Vector that indicates where the constrains are
      INTEGER*1 fixed(17) /0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0/
C Result vector
```

```
      14, 11, 9,
      9, 11, 4,
      14, 9, 13,
      10, 13, 6,
      6, 13, 9,
      10, 6, 5,
      6, 9, 4,
      5, 6, 4,
      1, 5, 4};
// Elemental matrices
      double Ke[21][3*3] ={
      {6.47e-5,-2.62e-5,-3.85e-5,-2.62e-5,6.47e-5,-3.85e-5,-3.85e-5,-3.85e-5,7.70e-5},
      {1.19e-4,-6.49e-5,-5.46e-5,-6.49e-5,6.45e-5,3.33e-7,-5.46e-5,3.33e-7,5.42e-5},
      {6.45e-5,-6.49e-5,3.33e-7,-6.49e-5,1.19e-4,-5.46e-5,3.33e-7,-5.46e-5,5.42e-5},
      {6.47e-5,-2.62e-5,-3.85e-5,-2.62e-5,6.47e-5,-3.85e-5,-3.85e-5,-3.85e-5,7.70e-5},
      {1.19e-4,-6.49e-5,-5.46e-5,-6.49e-5,6.45e-5,3.33e-7,-5.46e-5,3.33e-7,5.42e-5},
      {6.45e-5,-6.49e-5,3.33e-7,-6.49e-5,1.19e-4,-5.46e-5,3.33e-7,-5.46e-5,5.42e-5},
      {6.47e-5,-2.62e-5,-3.85e-5,-2.62e-5,6.47e-5,-3.85e-5,-3.85e-5,-3.85e-5,7.70e-5},
      {1.19e-4,-6.49e-5,-5.46e-5,-6.49e-5,6.45e-5,3.33e-7,-5.46e-5,3.33e-7,5.42e-5},
      {6.45e-5,-6.49e-5,3.33e-7,-6.49e-5,1.19e-4,-5.46e-5,3.33e-7,-5.46e-5,5.42e-5},
      {6.47e-5,-2.62e-5,-3.85e-5,-2.62e-5,6.47e-5,-3.85e-5,-3.85e-5,-3.85e-5,7.70e-5},
      {1.18e-4,-6.49e-5,-5.33e-5,-6.49e-5,6.52e-5,-3.33e-7,-5.33e-5,-3.33e-7,5.37e-5},
      {6.52e-5,-6.49e-5,-3.33e-7,-6.49e-5,1.18e-4,-5.33e-5,-3.33e-7,-5.33e-5,5.37e-5},
      {6.47e-5,-2.62e-5,-3.85e-5,-2.62e-5,6.47e-5,-3.85e-5,-3.85e-5,-3.85e-5,7.70e-5},
      {1.18e-4,-6.49e-5,-5.33e-5,-6.49e-5,6.52e-5,-3.33e-7,-5.33e-5,-3.33e-7,5.37e-5},
      {6.52e-5,-6.49e-5,-3.33e-7,-6.49e-5,1.18e-4,-5.33e-5,-3.33e-7,-5.33e-5,5.37e-5},
      {6.44e-5,-2.56e-5,-3.89e-5,-2.56e-5,6.44e-5,-3.89e-5,-3.89e-5,-3.89e-5,7.78e-5},
      {7.67e-5,-3.96e-5,-3.72e-5,-3.96e-5,6.60e-5,-2.64e-5,-3.72e-5,-2.64e-5,6.36e-5},
      {6.60e-5,-3.96e-5,-2.64e-5,-3.96e-5,7.67e-5,-3.72e-5,-2.64e-5,-3.72e-5,6.36e-5},
      {7.49e-5,-3.85e-5,-3.65e-5,-3.85e-5,6.64e-5,-2.80e-5,-3.65e-5,-2.80e-5,6.45e-5},
      {6.05e-5,-6.76e-5,7.08e-6,-6.76e-5,1.33e-4,-6.58e-5,7.08e-6,-6.58e-5,5.87e-5},
      {7.18e-5,-3.59e-5,-3.59e-5,-3.59e-5,6.67e-5,-3.07e-5,-3.59e-5,-3.07e-5,6.67e-5}};
// Vector with constrain values
      double x[17] = {0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 273, 0, 273, 0};
// Vector of independent terms
      double b[17] = {2.82e-4, 2.82e-4, 0, 0, 0, 0, 2.82e-4, 2.82e-4, 0, 0, 2.82e-4, 2.82e-4,
// Vector that indicates where the constrains are
      bool fixed[17] = {0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0};
// Result vector
```

```
      REAL*8 r(17)
C Variable to indicate command
      INTEGER*1 command
C Data pipe
      INTEGER*4 FEMData /100/
C Result pipe
      INTEGER*4 FEMResult /101/
C Indexes
      INTEGER*4 i, j
C Paths for pipes
#ifdef WIN32
      CHARACTER*128 dataname /'\\.\pipe\FEMData'/
      CHARACTER*128 resultname /'\\.\pipe\FEMResult'/
#else
      CHARACTER*128 dataname /'/tmp/FEMData'/
      CHARACTER*128 resultname /'/tmp/FEMResult'/
#endif

C Open data and result pipes
      OPEN(FEMData, FILE=dataname, ACCESS='STREAM')
      OPEN(FEMResult, FILE=resultname, ACCESS='STREAM')

C Send mesh data
      command = command_set_connectivity
      WRITE(FEMData) command, M, E, T, N, D, (connectivity(i), i = 1, E*N)



C Send elemental matrices
```

```
      double r[17];
// Variable to indicate command
      char command;
// Data pipe
      FILE* FEMData;
// Result pipe
      FILE* FEMResult;
// Indexes
      int i;
// Names for the pipes
#ifdef WIN32
      const char* dataname = "\\\\.\\pipe\\FEMData";
      const char* resultname = "\\\\.\\pipe\\FEMResult";
#else
      const char* dataname = "/tmp/FEMData";
      const char* resultname = "/tmp/FEMResult";
#endif

// Open data and result pipes
      FEMData = fopen(dataname, "wb");
      FEMResult = fopen(resultname, "rb");

// Send mesh data
      command = command_set_connectivity;
      fwrite(&command, 1, 1, FEMData);
      // Send number of nodes
      fwrite(&M, sizeof(int), 1, FEMData);
      // Send number of elements
      fwrite(&E, sizeof(int), 1, FEMData);
      // Send element type
      fwrite(&T, sizeof(int), 1, FEMData);
      // Send nodes per element
      fwrite(&N, sizeof(int), 1, FEMData);
      // Send degrees of freedom
      fwrite(&D, sizeof(int), 1, FEMData);
      /// Send connectivity
      fwrite(connectivity, sizeof(int), E*N, FEMData);

// Send elemental matrices
```

```
    command = command_set_Ae;
    DO i = 1, E
      WRITE(FEMData) command, i, (Ke(j, i), j = 1, N*N)
    ENDDO



C Send vector with constrain values
    command = command_set_x
    WRITE(FEMData) command, (x(j), j = 1, M)


C Send vector of independent terms
    command = command_set_b
    WRITE(FEMData) command, (b(j), j = 1, M)


C Send vector that indicates where the constrains are
    command = command_set_fixed
    WRITE(FEMData) command, (fixed(j), j = 1, M)


C Initialize solver
    command = command_solver_init
    WRITE(FEMData) command

C Run solver and read result
    command = command_solver_run
    WRITE(FEMData) command
    FLUSH(FEMData)
    READ(FEMResult) (r(j), j = 1, M)

C Display result
    DO i = 1, M
      WRITE(6, *) r(i)
    ENDDO

C Send end session command
```

```
    command = command_set_Ae;
    for (i = 0; i < E; ++i) {
        int e = i + 1;
        fwrite(&command, 1, 1, FEMData);
        fwrite(&e, sizeof(int), 1, FEMData);
        fwrite(Ke[i], sizeof(double), N*N, FEMData);
    }

    // Send vector with constrain values
    command = command_set_x;
    fwrite(&command, 1, 1, FEMData);
    fwrite(x, sizeof(double), M, FEMData);

    // Send vector of independent terms
    command = command_set_b;
    fwrite(&command, 1, 1, FEMData);
    fwrite(b, sizeof(double), M, FEMData);

    // Send vector that indicates where the constrains are
    command = command_set_fixed;
    fwrite(&command, 1, 1, FEMData);
    fwrite(fixed, sizeof(bool), M, FEMData);

    // Initialize solver
    command = command_solver_init;
    fwrite(&command, 1, 1, FEMData);

// Run solver and read result
    command = command_solver_run;
    fwrite(&command, 1, 1, FEMData);
    fflush(FEMData);
    fread(r, sizeof(double), M, FEMResult);

// Display result
    for (i = 0; i < M; ++i) {
        printf("%f\n", r[i]);
    }

// Send end session command
```

```
    command = command_end
    WRITE(FEMData) command
    FLUSH(FEMData)

C Close pipes
    CLOSE(FEMResult)
    CLOSE(FEMData)

    END
```

```
        command = command_end;
        fwrite(&command, 1, 1, FEMData);
        fflush(FEMData);

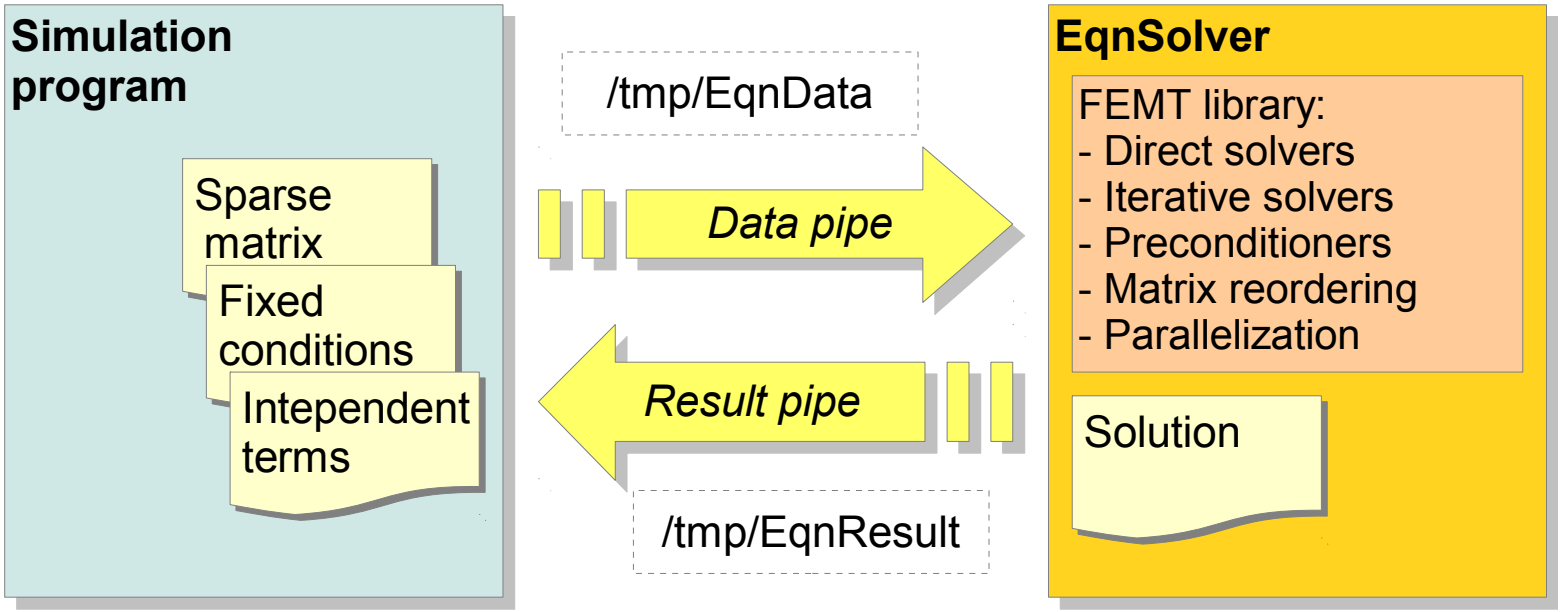
    // Close pipes
        fclose(FEMResult);
        fclose(FEMData);

        return 0;
    }
```

EqnSolver

This program was designed to solve systems of equations from finite volume.

EqnSolver works in a similar way than FEMSolver, but instead of mesh and elemental matrices, it takes as input a sparse matrix.



```
      PROGRAM EqnSolverExample

      IMPLICIT NONE
C Commands
      INTEGER command_end, command_set_size, command_set_row, command_set_x, c
      PARAMETER (
     . command_end        = 0,
     . command_set_size   = 1,
     . command_set_row    = 2,
     . command_set_x      = 3,
     . command_set_b      = 4,
     . command_set_fixed  = 5,
     . command_solver_init = 6,
     . command_solver_run  = 7)
C Number of equations
      INTEGER*4 M /22/
C Matrix values
      REAL*8 values(120) /
     . 1.2326e-4, -3.6051e-5, -3.3997e-5, -5.3215e-5,
     . -3.6051e-5, 1.7051e-4, -4.3495e-5, -6.0878e-5, -3.0087e-5,
     . -3.3997e-5, 1.7032e-4, -5.5748e-5, -4.5665e-5, -3.4909e-5,
     . -5.3215e-5, -4.3495e-5, -5.5748e-5, 3.4166e-4, -7.1143e-5, -6.8663e-5, -4.9398e-5,
     . -6.0878e-5, -7.1143e-5, 3.4267e-4, -5.8604e-5, -4.4280e-5, -4.3102e-5, -6.4661e-5,
     . -4.5665e-5, -6.8663e-5, 3.4155e-4, -5.6230e-5, -4.4270e-5, -5.8971e-5, -6.7748e-5,
     . -3.0087e-5, -5.8604e-5, 1.2279e-4, -3.4102e-5,
     . -3.4909e-5, -5.6230e-5, 1.2292e-4, -3.1777e-5,
     . -4.9398e-5, -4.4280e-5, -4.4270e-5, 3.2961e-4, -4.7225e-5, -5.2776e-5, -5.2312e-5, -
     . -5.8971e-5, -3.1777e-5, 1.7029e-4, -4.4635e-5, -3.4909e-5,
     . -4.3102e-5, -3.4102e-5, 1.7088e-4, -6.6465e-5, -2.7215e-5,
     . -6.4661e-5, -4.7225e-5, -6.6465e-5, 3.4234e-4, -5.0918e-5, -6.2797e-5, -5.0270e-5,
     . -6.7748e-5, -5.2776e-5, -4.4635e-5, 3.4104e-4, -5.8173e-5, -5.6230e-5, -6.1480e-5,
     . -5.2312e-5, -5.0918e-5, 3.4252e-4, -7.1496e-5, -5.5712e-5, -5.1108e-5, -6.0978e-5,
     . -3.9348e-5, -5.8173e-5, -7.1496e-5, 3.4468e-4, -4.5355e-5, -6.3924e-5, -6.6386e-5,
     . -3.4909e-5, -5.6230e-5, 1.2292e-4, -3.1777e-5,
     . -2.7215e-5, -6.2797e-5, 1.2287e-4, -3.2859e-5,
     . -6.1480e-5, -4.5355e-5, -3.1777e-5, 1.7057e-4, -3.1963e-5,
     . -5.0270e-5, -5.5712e-5, -3.2859e-5, 1.7088e-4, -3.2040e-5,
     . -5.1108e-5, -6.3924e-5, 1.6950e-4, -2.4862e-5, -2.9603e-5,
     . -6.6386e-5, -3.1963e-5, -2.4862e-5, 1.2321e-4,
```

```
#include <stdio.h>

int main()
{
// Commands
  enum Command {
    command_end        = 0,
    command_set_size   = 1,
    command_set_row    = 2,
    command_set_x      = 3,
    command_set_b      = 4,
    command_set_fixed  = 5,
    command_solver_init = 6,
    command_solver_run  = 7};
// Number of equations
  int M = 22;
// Matrix values
  double values[120] = {
    1.2326e-4, -3.6051e-5, -3.3997e-5, -5.3215e-5,
    -3.6051e-5, 1.7051e-4, -4.3495e-5, -6.0878e-5, -3.0087e-5,
    -3.3997e-5, 1.7032e-4, -5.5748e-5, -4.5665e-5, -3.4909e-5,
    -5.3215e-5, -4.3495e-5, -5.5748e-5, 3.4166e-4, -7.1143e-5, -6.8663e-5, -4.9398e-5,
    -6.0878e-5, -7.1143e-5, 3.4267e-4, -5.8604e-5, -4.4280e-5, -4.3102e-5, -6.4661e-5,
    -4.5665e-5, -6.8663e-5, 3.4155e-4, -5.6230e-5, -4.4270e-5, -5.8971e-5, -6.7748e-5,
    -3.0087e-5, -5.8604e-5, 1.2279e-4, -3.4102e-5,
    -3.4909e-5, -5.6230e-5, 1.2292e-4, -3.1777e-5,
    -4.9398e-5, -4.4280e-5, -4.4270e-5, 3.2961e-4, -4.7225e-5, -5.2776e-5, -5.2312e-5, -
    -5.8971e-5, -3.1777e-5, 1.7029e-4, -4.4635e-5, -3.4909e-5,
    -4.3102e-5, -3.4102e-5, 1.7088e-4, -6.6465e-5, -2.7215e-5,
    -6.4661e-5, -4.7225e-5, -6.6465e-5, 3.4234e-4, -5.0918e-5, -6.2797e-5, -5.0270e-5,
    -6.7748e-5, -5.2776e-5, -4.4635e-5, 3.4104e-4, -5.8173e-5, -5.6230e-5, -6.1480e-5,
    -5.2312e-5, -5.0918e-5, 3.4252e-4, -7.1496e-5, -5.5712e-5, -5.1108e-5, -6.0978e-5,
    -3.9348e-5, -5.8173e-5, -7.1496e-5, 3.4468e-4, -4.5355e-5, -6.3924e-5, -6.6386e-5,
    -3.4909e-5, -5.6230e-5, 1.2292e-4, -3.1777e-5,
    -2.7215e-5, -6.2797e-5, 1.2287e-4, -3.2859e-5,
    -6.1480e-5, -4.5355e-5, -3.1777e-5, 1.7057e-4, -3.1963e-5,
    -5.0270e-5, -5.5712e-5, -3.2859e-5, 1.7088e-4, -3.2040e-5,
    -5.1108e-5, -6.3924e-5, 1.6950e-4, -2.4862e-5, -2.9603e-5,
    -6.6386e-5, -3.1963e-5, -2.4862e-5, 1.2321e-4,
```

```
     . -6.0978e-5, -3.2040e-5, -2.9603e-5, 1.2262e-4/
C Matrix indexes
      INTEGER*4 indexes(120) /
     . 1, 2, 3, 4,
     . 1, 2, 4, 5, 7,
     . 1, 3, 4, 6, 8,
     . 1, 2, 3, 4, 5, 6, 9,
     . 2, 4, 5, 7, 9, 11, 12,
     . 3, 4, 6, 8, 9, 10, 13,
     . 2, 5, 7, 11,
     . 3, 6, 8, 10,
     . 4, 5, 6, 9, 12, 13, 14, 15,
     . 6, 8, 10, 13, 16,
     . 5, 7, 11, 12, 17,
     . 5, 9, 11, 12, 14, 17, 19,
     . 6, 9, 10, 13, 15, 16, 18,
     . 9, 12, 14, 15, 19, 20, 22,
     . 9, 13, 14, 15, 18, 20, 21,
     . 10, 13, 16, 18,
     . 11, 12, 17, 19,
     . 13, 15, 16, 18, 21,
     . 12, 14, 17, 19, 22,
     . 14, 15, 20, 21, 22,
     . 15, 18, 20, 21,
     . 14, 19, 20, 22/
C Row sizes
      INTEGER*4 count(22) /4, 5, 5, 7, 7, 7, 4, 4, 8, 5, 5, 7, 7, 7, 7, 4, 4, 5, 5, 5, 4, 4/
C Vector with constrain values
      REAL*8 x(22) /273, 273, 0, 0, 0, 0, 273, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 293, 0/
C Vector of independent terms
      REAL*8 b(22) /0, 0, 0, 0, 0, 0, 0, -4.3388e-4, 0, -8.6776e-4, 0, 0, 0, 0, 0, -4.3388e-4, -2.
C Vector that indicates where the constrains are
      INTEGER*1 fixed(22) /1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0/
C Result vector
      REAL*8 r(22)
C Variable to indicate command
      INTEGER*1 command
C Data pipe
      INTEGER*4 EqnData /100/
```

```
     -6.0978e-5, -3.2040e-5, -2.9603e-5, 1.2262e-4};
// Matrix indexes
  int indexes[120]= {
     1, 2, 3, 4,
     1, 2, 4, 5, 7,
     1, 3, 4, 6, 8,
     1, 2, 3, 4, 5, 6, 9,
     2, 4, 5, 7, 9, 11, 12,
     3, 4, 6, 8, 9, 10, 13,
     2, 5, 7, 11,
     3, 6, 8, 10,
     4, 5, 6, 9, 12, 13, 14, 15,
     6, 8, 10, 13, 16,
     5, 7, 11, 12, 17,
     5, 9, 11, 12, 14, 17, 19,
     6, 9, 10, 13, 15, 16, 18,
     9, 12, 14, 15, 19, 20, 22,
     9, 13, 14, 15, 18, 20, 21,
     10, 13, 16, 18,
     11, 12, 17, 19,
     13, 15, 16, 18, 21,
     12, 14, 17, 19, 22,
     14, 15, 20, 21, 22,
     15, 18, 20, 21,
     14, 19, 20, 22};
// Row sizes
  int count[22] = {4, 5, 5, 7, 7, 7, 4, 4, 8, 5, 5, 7, 7, 7, 7, 4, 4, 5, 5, 5, 4, 4};
// Vector with constrain values
  double x[22] = {273, 273, 0, 0, 0, 0, 273, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 293, 0};
// Vector of independent terms
  double b[22] = {0, 0, 0, 0, 0, 0, 0, -4.3388e-4, 0, -8.6776e-4, 0, 0, 0, 0, 0, -4.3388e-4, -2.
// Vector that indicates where the constrains are
  bool fixed[22] = {1, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0};
// Result vector
  double r[22];
// Variable to indicate command
  char command;
// Data pipe
  FILE* EqnData;
```

```
C Result pipe
      INTEGER*4 EqnResult /101/
C Indexes
      INTEGER*4 i, j, j1, jn
C Names for the pipes
#ifdef WIN32
      CHARACTER*128 dataname /'\\.\pipe\EqnData'/
      CHARACTER*128 resultname /'\\.\pipe\EqnResult'/
#else
      CHARACTER*128 dataname /'/tmp/EqnData'/
      CHARACTER*128 resultname /'/tmp/EqnResult'/
#endif

C Open data and result pipes
      OPEN(EqnData, FILE=dataname, ACCESS='STREAM')
      OPEN(EqnResult, FILE=resultname, ACCESS='STREAM')

C Send number of equations
      command = command_set_size
      WRITE(EqnData) command, M


C Send rows
      command = command_set_row
      j1 = 1
      DO i = 1, M
        jn = j1 + count(i) - 1
        WRITE(EqnData) command, i, count(i), (indexes(j), j = j1, jn), (values(j), j = j1, jn)
        j1 = jn + 1
      ENDDO



C Send vector with constrain values
      command = command_set_x
      WRITE(EqnData) command, (x(j), j = 1, M)
```

```
// Result pipe
  FILE* EqnResult;
// Indexes
  int i, j1;
// Names for the pipes
#ifdef WIN32
  const char* dataname = "\\\\.\\pipe\\EqnData";
  const char* resultname = "\\\\.\\pipe\\EqnResult";
#else
  const char* dataname = "/tmp/EqnData";
  const char* resultname = "/tmp/EqnResult";
#endif

// Open data and result pipes
  EqnData = fopen(dataname, "wb");
  EqnResult = fopen(resultname, "rb");

// Send number of equations
  command = command_set_size;
  fwrite(&command, 1, 1, EqnData);
  fwrite(&M, sizeof(int), 1, EqnData);

// Send rows
  command = command_set_row;
  j1 = 0;
  for (i = 0; i < M; ++i) {
    int row = i + 1;
    fwrite(&command, 1, 1, EqnData);
    fwrite(&row, sizeof(int), 1, EqnData);
    fwrite(&count[i], sizeof(int), 1, EqnData);
    fwrite(indexes + j1, sizeof(int), count[i], EqnData);
    fwrite(values + j1, sizeof(double), count[i], EqnData);
    j1 += count[i];
  }

  // Send vector with constrain values
  command = command_set_x;
  fwrite(&command, 1, 1, EqnData);
  fwrite(x, sizeof(double), M, EqnData);
```

```
C Send vector of independent terms
    command = command_set_b
    WRITE(EqnData) command, (b(j), j = 1, M)


C Send vector that indicates where the constrains are
    command = command_set_fixed
    WRITE(EqnData) command, (fixed(j), j = 1, M)


C Initialize solver
    command = command_solver_init
    WRITE(EqnData) command

C Run solver and read result
    command = command_solver_run
    WRITE(EqnData) command
    FLUSH(EqnData)
    READ(EqnResult) (r(j), j = 1, M)

C Display result
    DO i = 1, M
     WRITE(6, *) r(i)
    ENDDO

C Send end session command
    command = command_end
    WRITE(EqnData) command
    FLUSH(EqnData)

C Close pipes
    CLOSE(EqnResult)
    CLOSE(EqnData)

    END
```

```
  // Send vector of independent terms
  command = command_set_b;
  fwrite(&command, 1, 1, EqnData);
  fwrite(b, sizeof(double), M, EqnData);

  // Send vector that indicates where the constrains are
  command = command_set_fixed;
  fwrite(&command, 1, 1, EqnData);
  fwrite(fixed, sizeof(bool), M, EqnData);

  // Initialize solver
  command = command_solver_init;
  fwrite(&command, 1, 1, EqnData);

// Run solver and read result
  command = command_solver_run;
  fwrite(&command, 1, 1, EqnData);
  fflush(EqnData);
  fread(r, sizeof(double), M, EqnResult);

// Display result
  for (i = 0; i < M; ++i) {
    printf("%f\n", r[i]);
  }

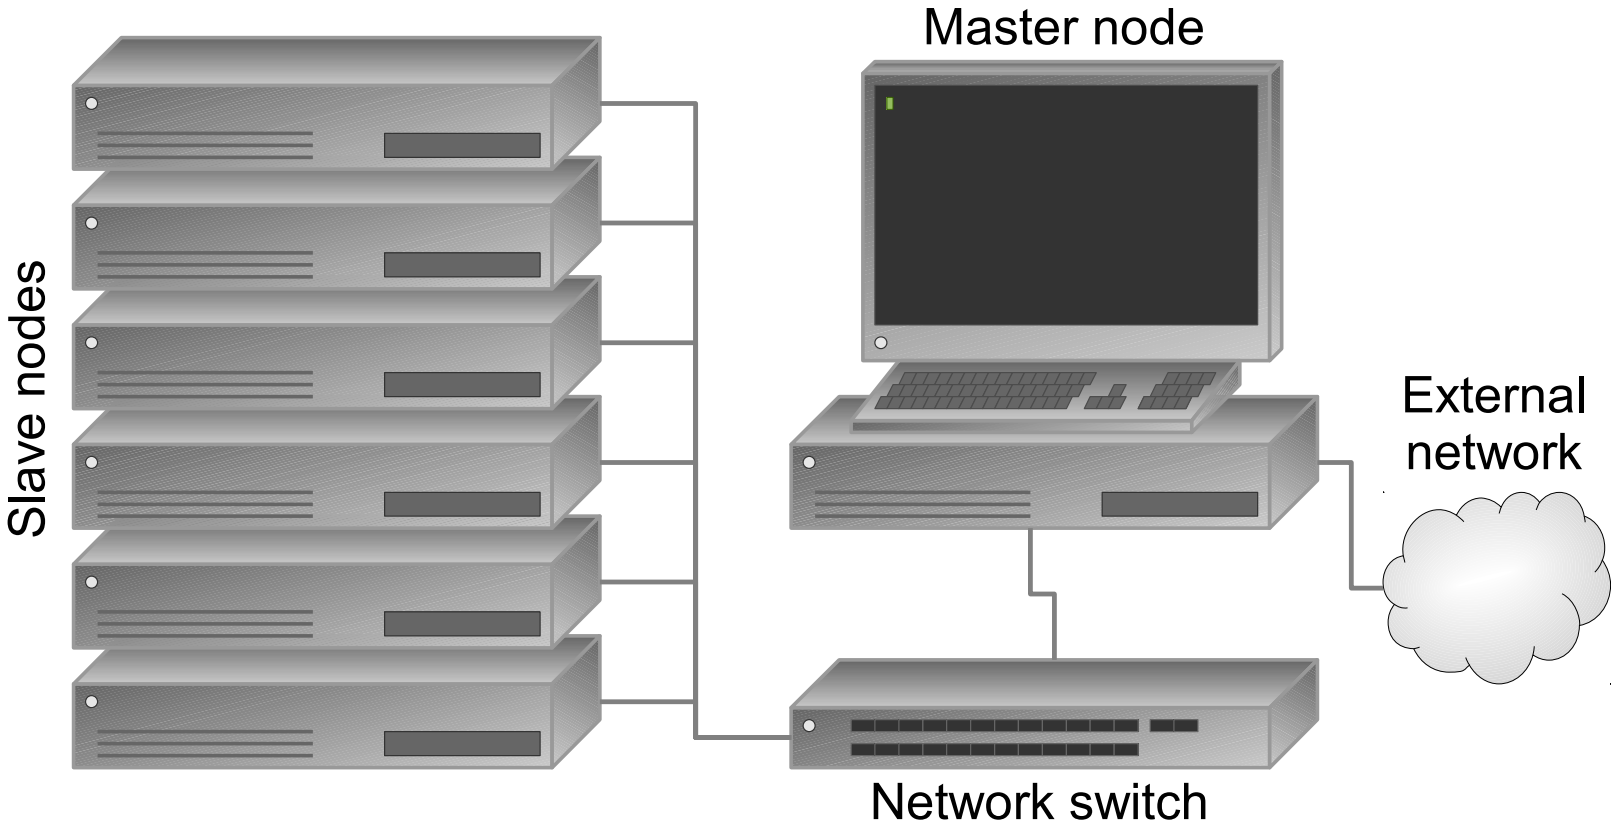
// Send end session command
  command = command_end;
  fwrite(&command, 1, 1, EqnData);
  fflush(EqnData);

// Close pipes
  fclose(EqnResult);
  fclose(EqnData);

   return 0;
}
```

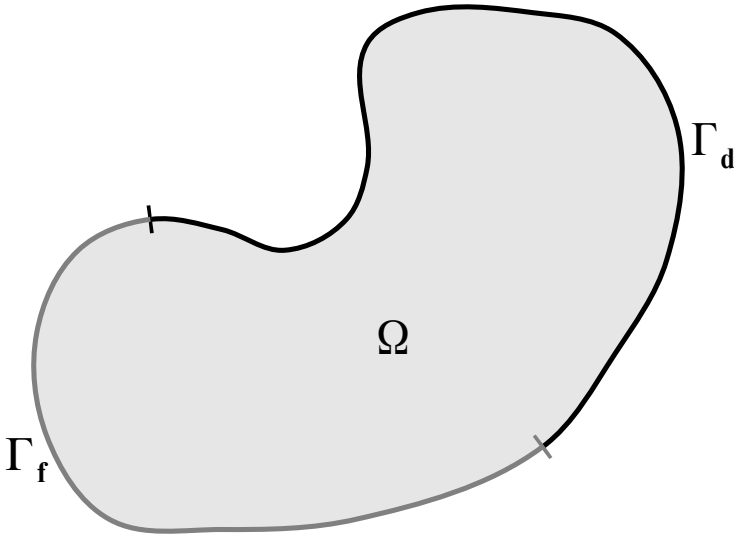
Parallelization in computer clusters

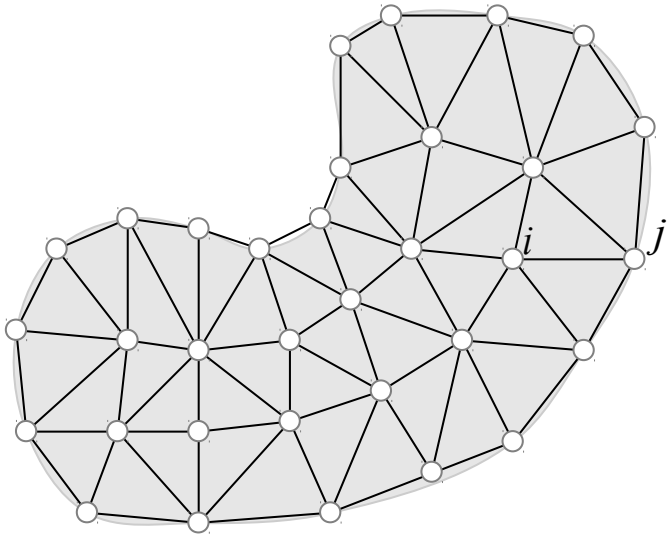
We developed a software program that runs in parallel in a Beowulf cluster [Ster95]. A Beowulf cluster consists of several multi-core computers (nodes) connected with a high speed network.

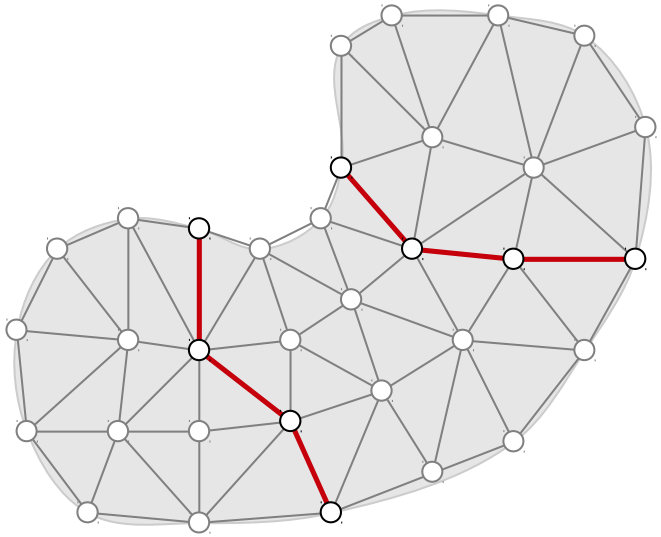


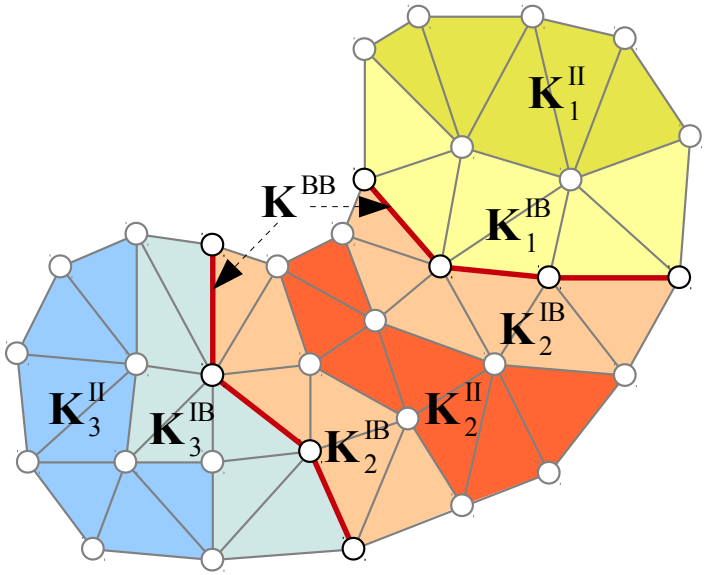
Schur substructuring method

This is a domain decomposition method without overlapping [Krui04].









$$\begin{pmatrix} \mathbf{K}_1^{II} & \mathbf{0} & \mathbf{0} & \mathbf{K}_1^{IB} \\ \mathbf{0} & \mathbf{K}_2^{II} & \mathbf{0} & \mathbf{K}_2^{IB} \\ \mathbf{0} & \mathbf{0} & \mathbf{K}_3^{II} & \mathbf{K}_3^{IB} \\ \mathbf{K}_1^{BI} & \mathbf{K}_2^{BI} & \mathbf{K}_3^{BI} & \mathbf{K}^{BB} \end{pmatrix}$$

We can arrange (reorder variables) of the system of equations to have the following form

$$\begin{pmatrix} \mathbf{K}_1^{\mathrm{II}} & & \mathbf{0} & & & \mathbf{K}_1^{\mathrm{IB}} \\ & \mathbf{K}_2^{\mathrm{II}} & & & & \mathbf{K}_2^{\mathrm{IB}} \\ \mathbf{0} & & \mathbf{K}_3^{\mathrm{II}} & & & \mathbf{K}_3^{\mathrm{IB}} \\ \vdots & & & \ddots & & \vdots \\ & & & & \mathbf{K}_p^{\mathrm{II}} & \mathbf{K}_p^{\mathrm{IB}} \\ \mathbf{K}_1^{\mathrm{BI}} & \mathbf{K}_2^{\mathrm{BI}} & \mathbf{K}_3^{\mathrm{BI}} & \cdots & \mathbf{K}_p^{\mathrm{BI}} & \mathbf{K}^{\mathrm{BB}} \end{pmatrix} \begin{pmatrix} \mathbf{d}_1^{\mathrm{I}} \\ \mathbf{d}_2^{\mathrm{I}} \\ \mathbf{d}_3^{\mathrm{I}} \\ \vdots \\ \mathbf{d}_p^{\mathrm{I}} \\ \mathbf{d}^{\mathrm{B}} \end{pmatrix} = \begin{pmatrix} \mathbf{f}_1^{\mathrm{I}} \\ \mathbf{f}_2^{\mathrm{I}} \\ \mathbf{f}_3^{\mathrm{I}} \\ \vdots \\ \mathbf{f}_p^{\mathrm{I}} \\ \mathbf{f}^{\mathrm{B}} \end{pmatrix}.$$

The superscript II denotes entries that capture the relationship between nodes inside a partition. BB is used to indicate entries in the matrix that relate nodes on the boundary. Finally IB and BI are used for entries with values dependent of nodes in the boundary and nodes inside the partition.

Thus, the system can be separated in p different systems,

$$\begin{pmatrix} \mathbf{K}_i^{\mathrm{II}} & \mathbf{K}_i^{\mathrm{IB}} \\ \mathbf{K}_i^{\mathrm{BI}} & \mathbf{K}^{\mathrm{BB}} \end{pmatrix} \begin{pmatrix} \mathbf{d}_i^{\mathrm{I}} \\ \mathbf{d}^{\mathrm{B}} \end{pmatrix} = \begin{pmatrix} \mathbf{f}_i^{\mathrm{I}} \\ \mathbf{f}^{\mathrm{B}} \end{pmatrix}, \; i = 1 \ldots p.$$

For each partition i the vector of unknowns $\mathbf{d}_i^{\mathrm{I}}$ as

$$\mathbf{d}_i^{\mathrm{I}} = \left(\mathbf{K}_i^{\mathrm{II}} \right)^{-1} \left(\mathbf{f}_i^{\mathrm{I}} - \mathbf{K}_i^{\mathrm{IB}} \mathbf{d}^{\mathrm{B}} \right).$$

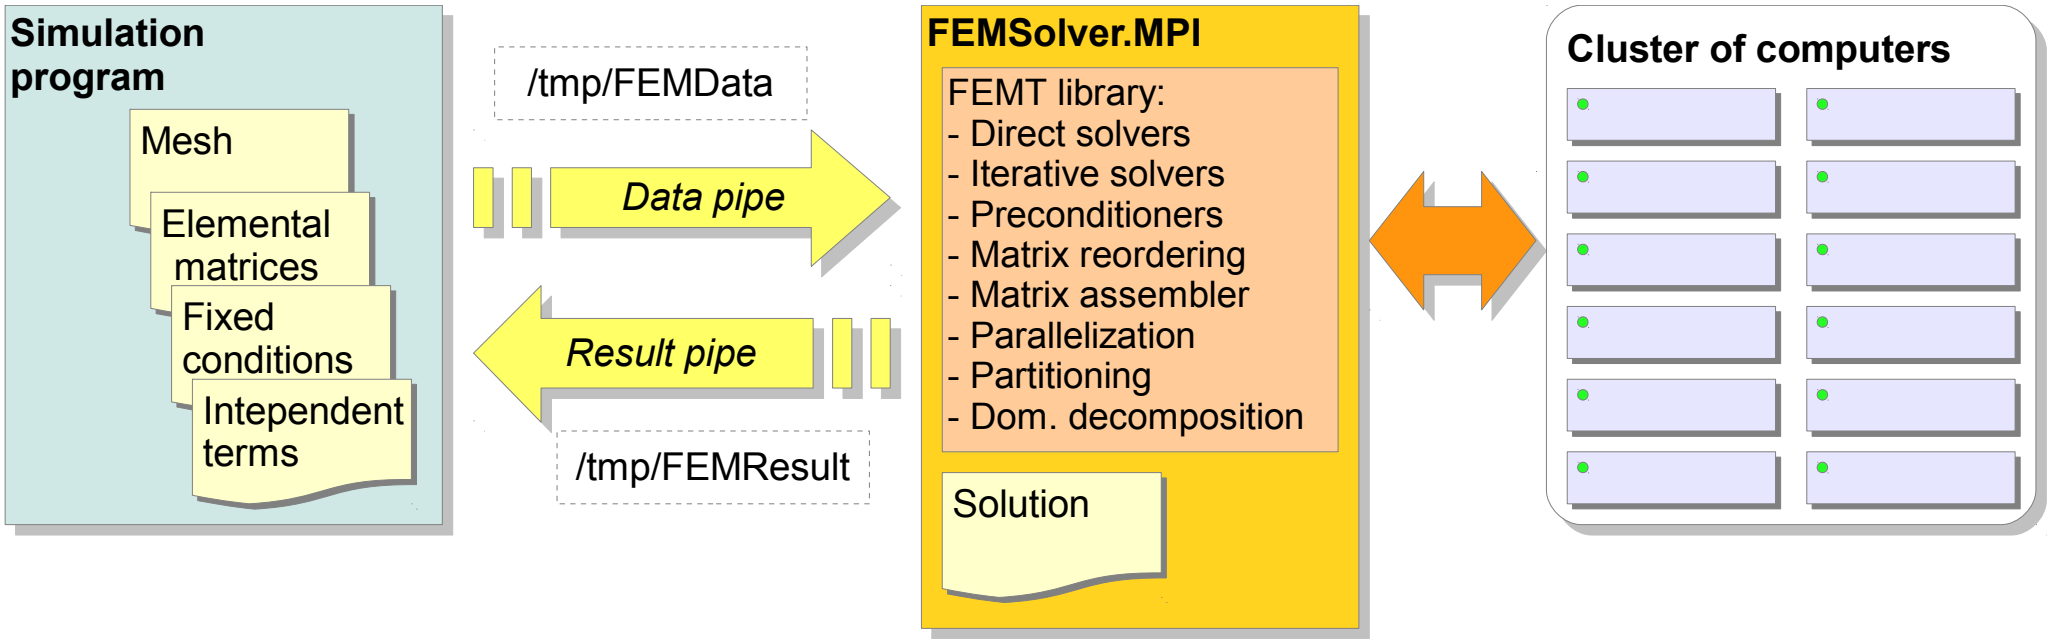
After applying Gaussian elimination by blocks the reduced system of equations becomes

$$\left(\mathbf{K}^{\mathrm{BB}}-\sum_{i=1}^{p}\mathbf{K}_i^{\mathrm{BI}}\left(\mathbf{K}_i^{\mathrm{II}}\right)^{-1}\mathbf{K}_i^{\mathrm{IB}}\right)\mathbf{d}^{\mathrm{B}}=\mathbf{f}^{\mathrm{B}}-\sum_{i=1}^{p}\mathbf{K}_i^{\mathrm{BI}}\left(\mathbf{K}_i^{\mathrm{II}}\right)^{-1}\mathbf{f}_i^{\mathrm{I}}.$$

Once the vector $\mathbf{d}^{\mathrm{B}}$ is computed, we can calculate the internal unknowns $\mathbf{d}_i^{\mathrm{I}}$.

FEMSolver.MPI

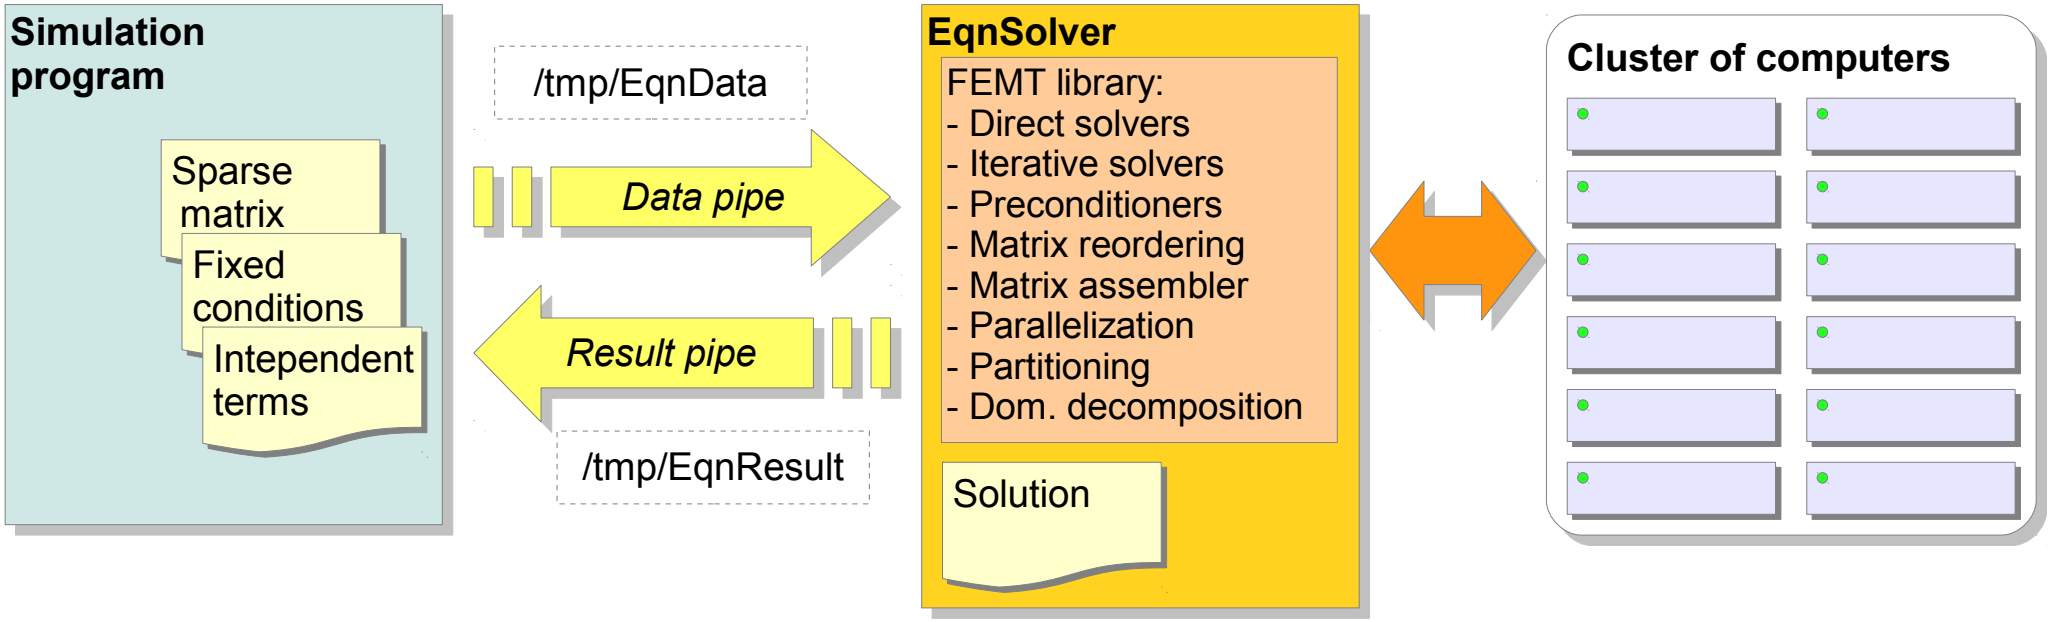
FEMSolver.MPI is a program similar to FEMSolver, but instead of solving the system of equations using a single computer, it can use a cluster of computers to distribute the workload and solve even larger systems of equations.



It uses the MPI technology to handle communication between nodes in the cluster. It makes high performance computing (HPC) easy to use.

EqnSolver.MPI (alpha stage)

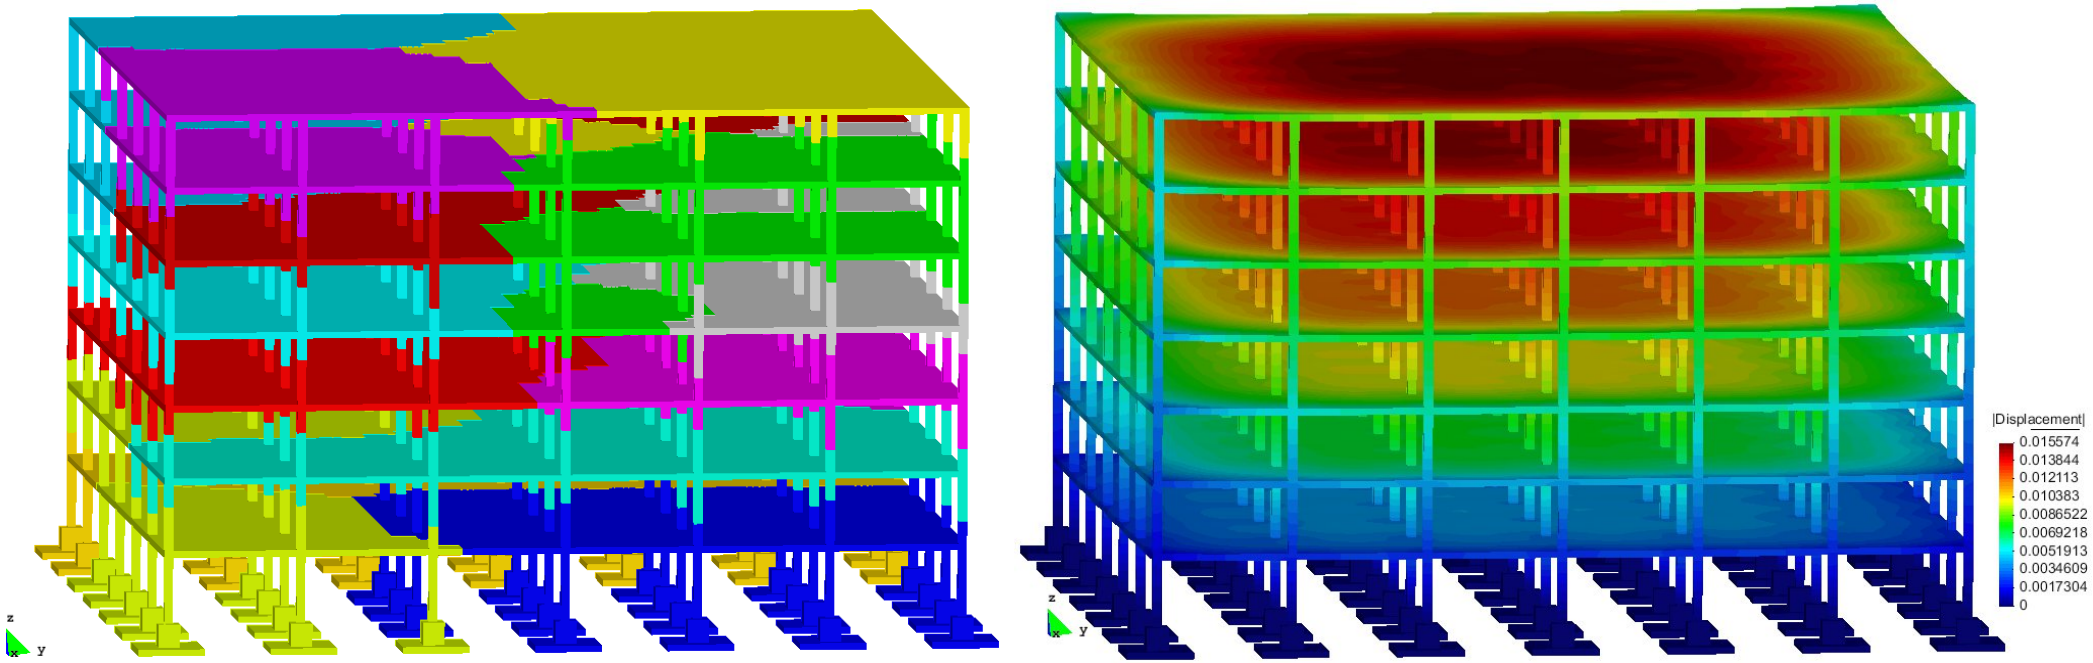
EqnSolver.MPI is a program similar to EqnSolver, but instead of solving the system of equations using a single computer, it can use a cluster of computers to distribute the workload and solve even larger systems of equations.



Building deformation

We used a cluster with 15 nodes, each one with two dual core Intel Xeon E5502 (1.87GHz) processors, a total of 60 cores.

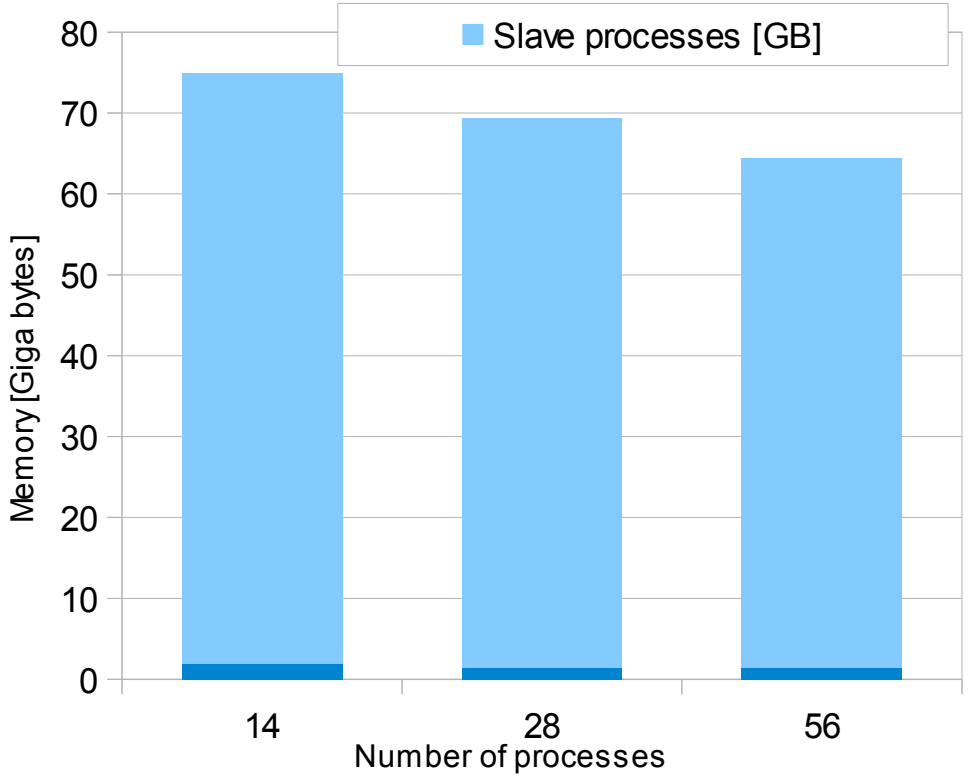
The problem tested is a 3D solid model of a building that is deformed due to self weight. The geometry is divided in 1’336,832 elements, with 1’708,273 nodes, with three degrees of freedom per node the resulting system of equations has 5’124,819 unknowns. Tolerance used is $1x10^{-10}$.



Substructuration of the domain (left) resulting deformation (right)

Number of processes	Partitioning time [s]	Inversion time (Cholesky) [s]	Schur complement time (CG) [s]	CG steps	Total time [s]
14	47.6	18520.8	4444.5	6927	23025.0
28	45.7	6269.5	2444.5	8119	8771.6
56	44.1	2257.1	2296.3	9627	4608.9

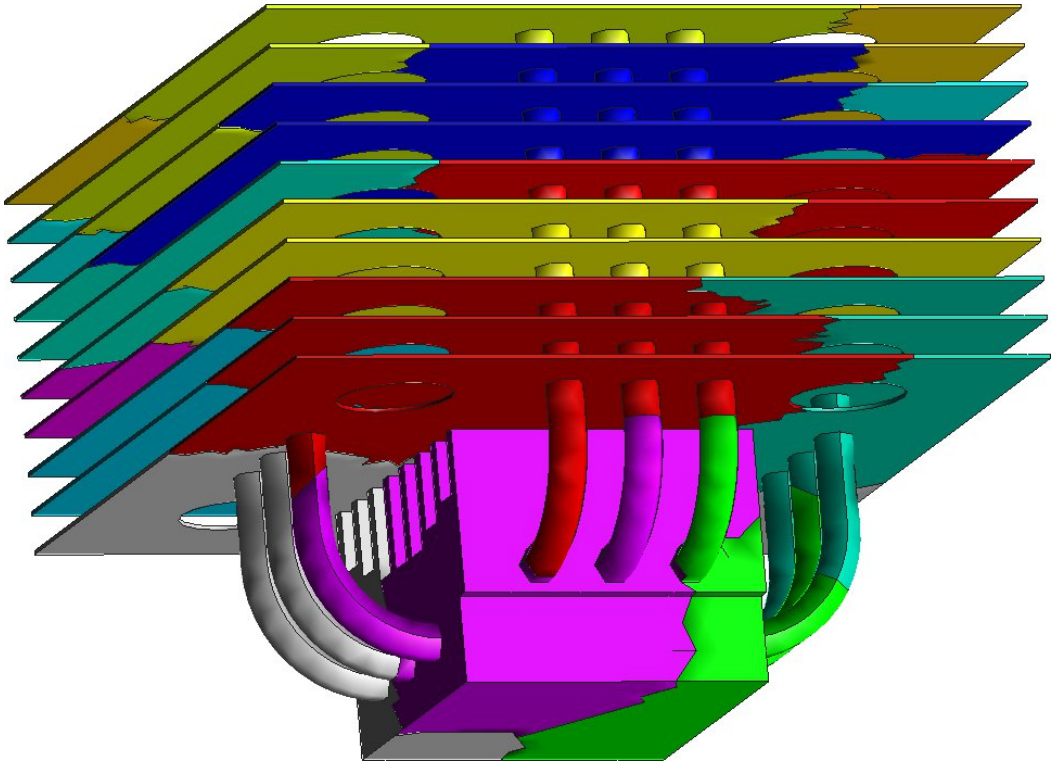


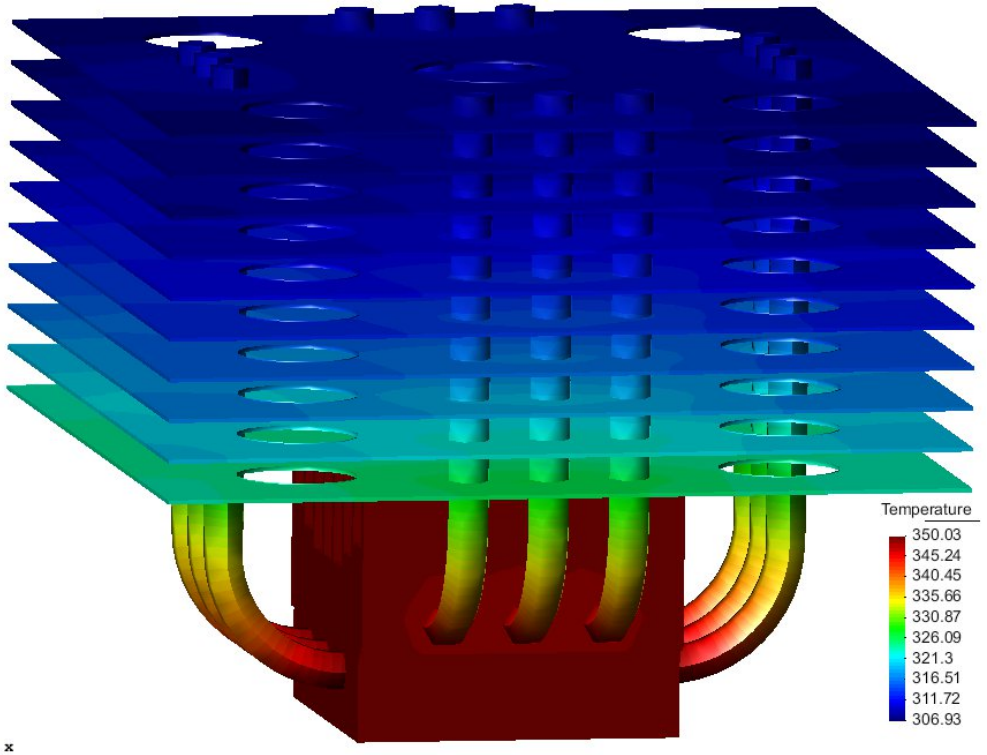


Number of processes	Master process [GB]	Slave processes [GB]	Total memory [GB]
14	1.89	73.00	74.89
28	1.43	67.88	69.32
56	1.43	62.97	64.41

Heat diffusion

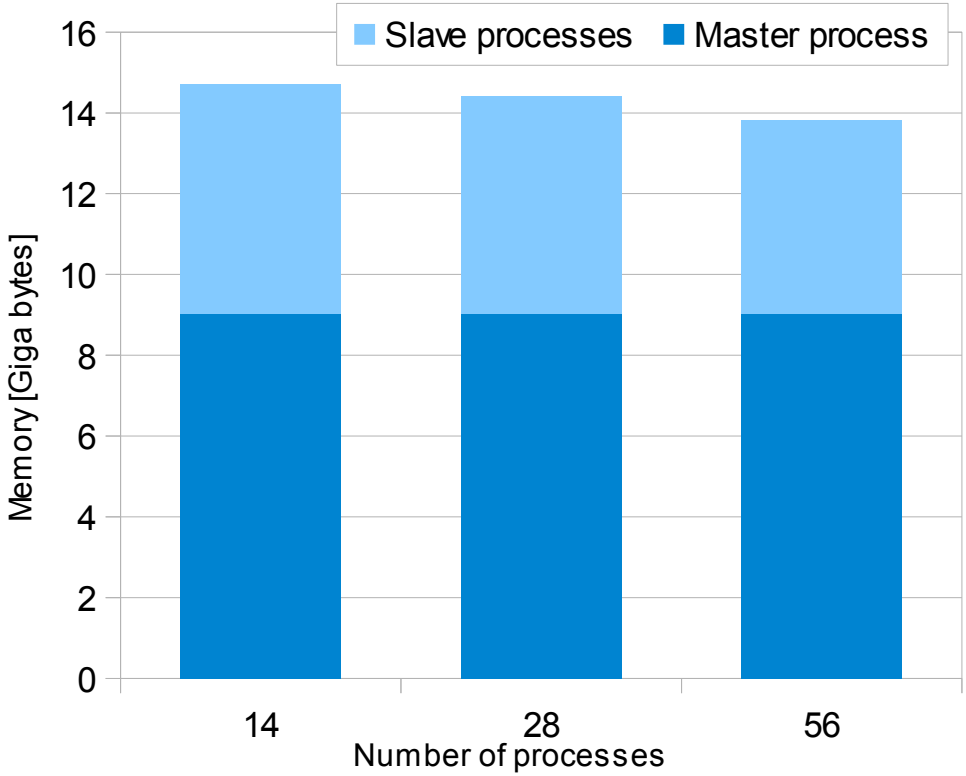
This is a 3D model of a heat sink, in this problem the base of the heat sink is set to a certain temperature and heat is lost in all the surfaces at a fixed rate. The geometry is divided in 4’493,232 elements, with 1’084,185 nodes. The system of equations solved had 1’084,185 unknowns.





Number of processes	Partitioning time [s]	Inversion time (Cholesky) [s]	Schur complement time (CG) [s]	CG steps	Total time [s]
14	144.9	798.5	68.1	307	1020.5
28	146.6	242.0	52.1	348	467.1
56	144.2	82.8	27.6	391	264.0

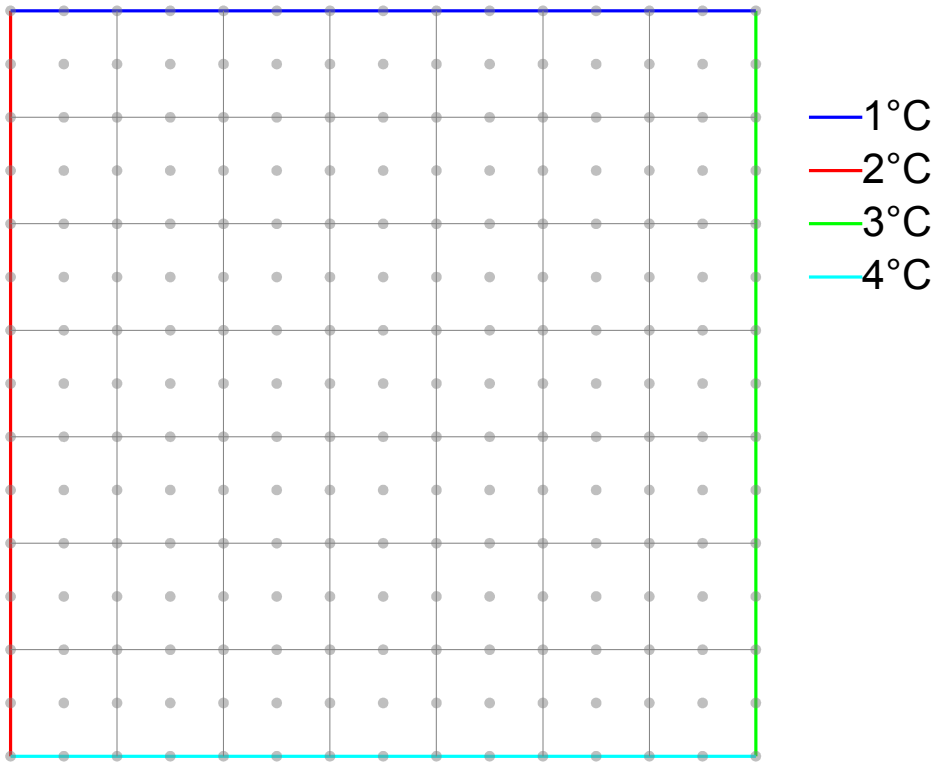


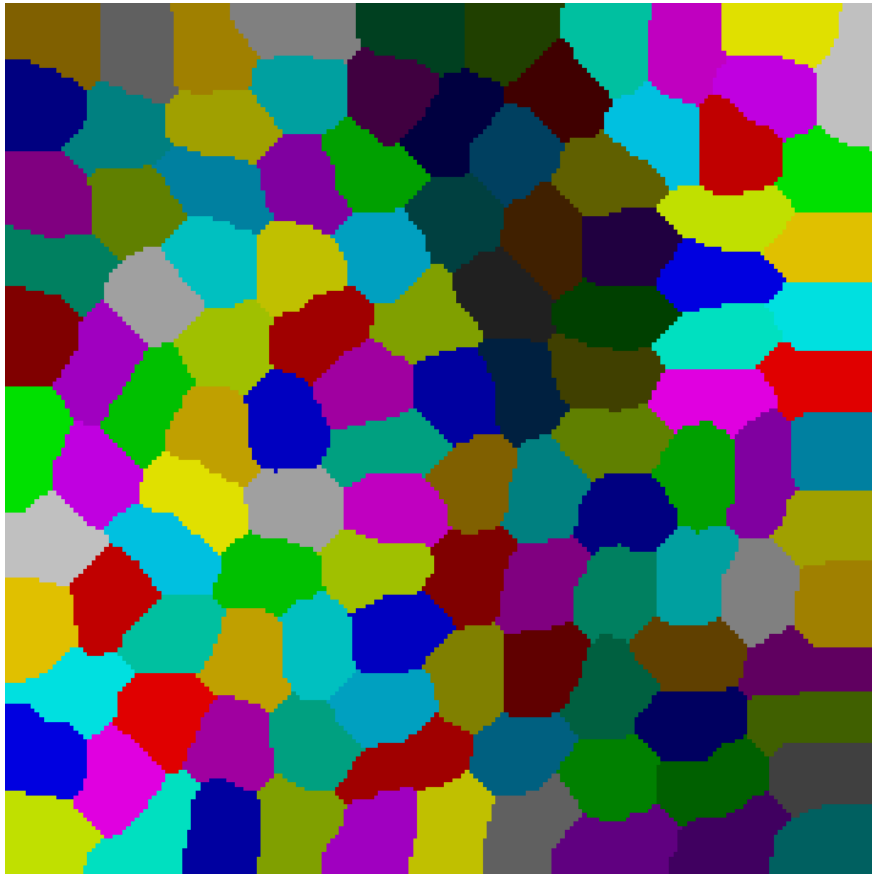


Number of processes	Master process [GB]	Slave processes [GB]	Total memory [GB]
14	9.03	5.67	14.70
28	9.03	5.38	14.41
56	9.03	4.80	13.82

Larger systems of equations

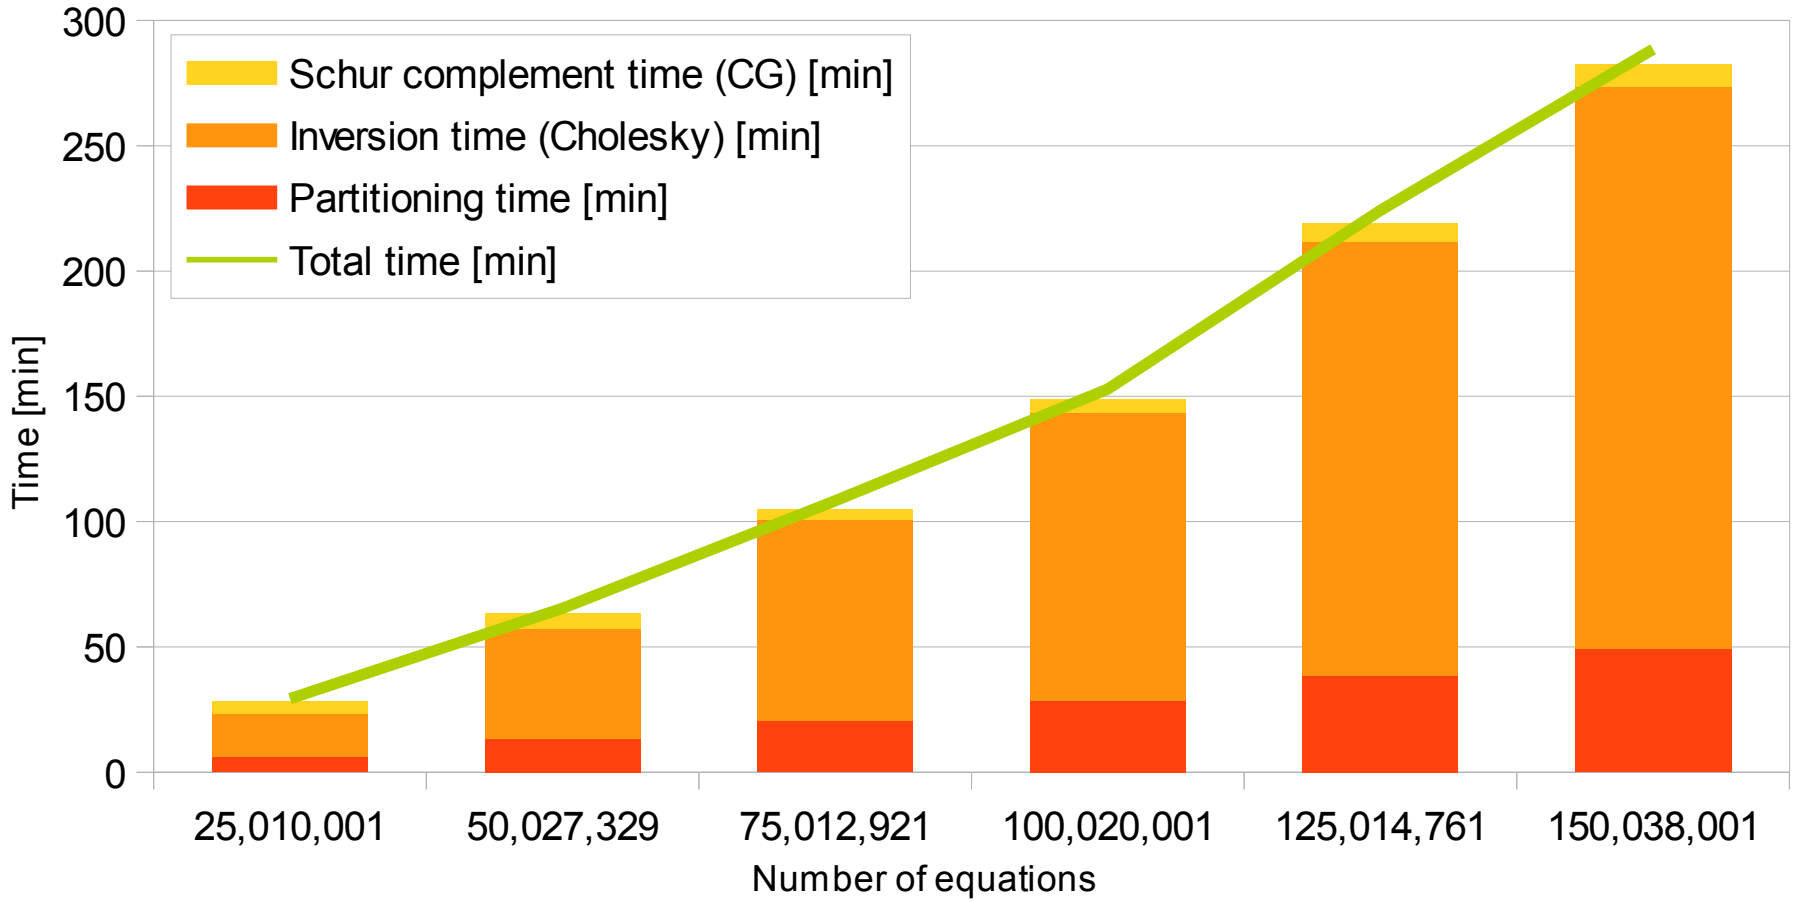
To test solution times in larger systems of equations we set a simple geometry. We calculated the temperature distribution of a metalic square with Dirichlet conditions on all boundaries.



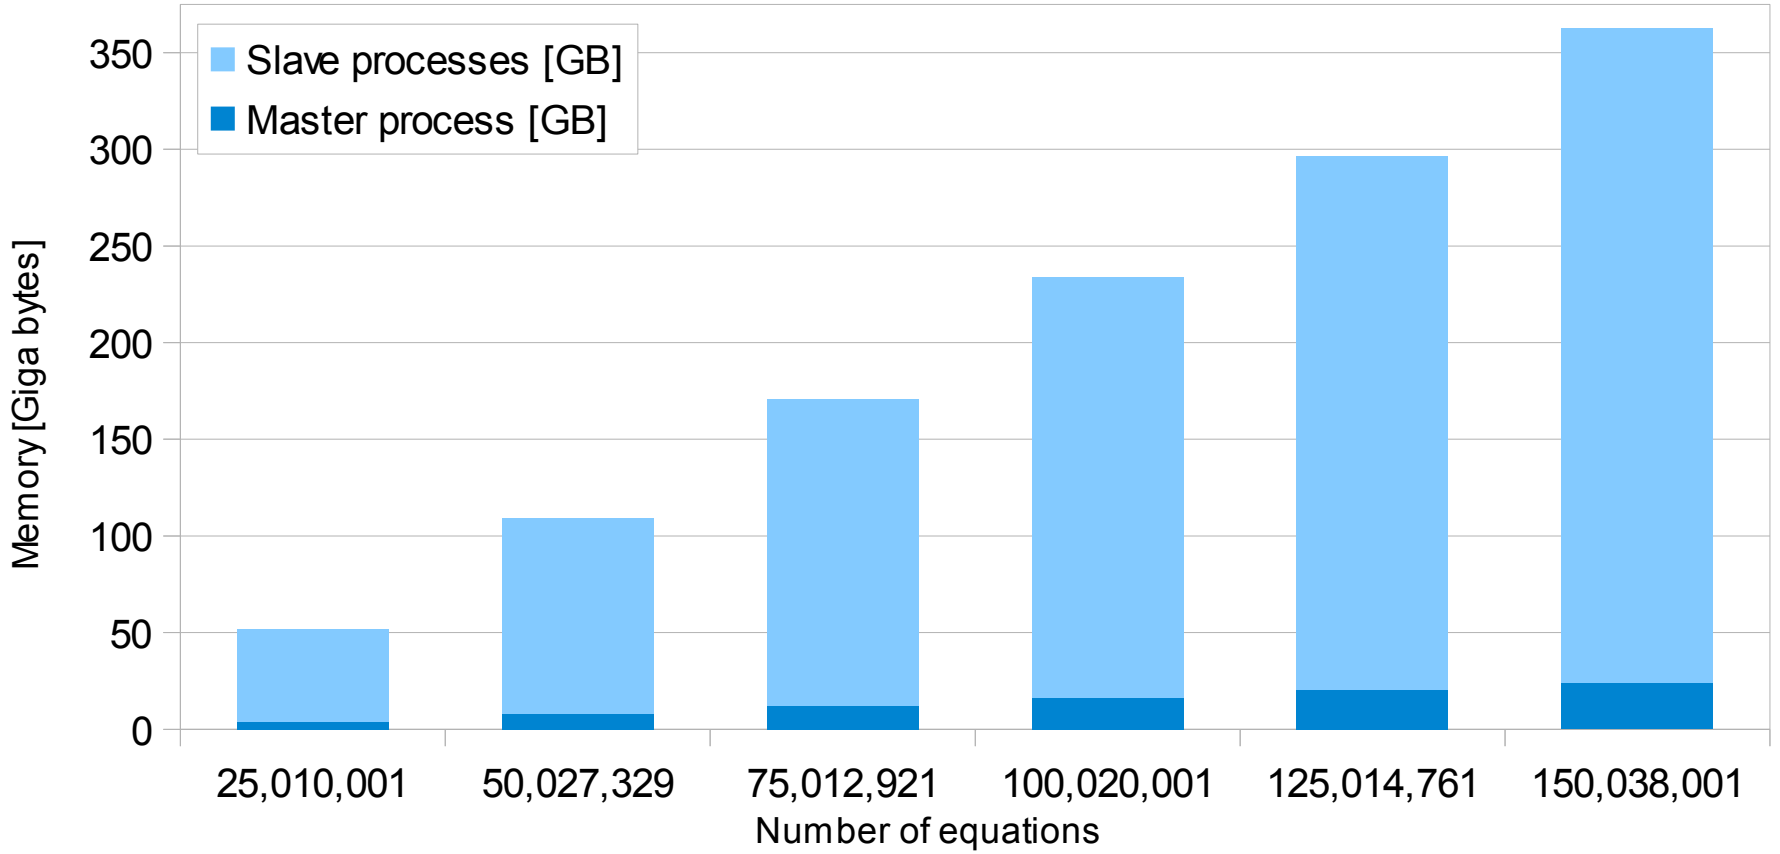


The domain was discretized using quadrilaterals with nine nodes, the discretization made was from 25 million nodes up to 150 million nodes.

In all cases we divided the domain into 116 partitions, each partition is solved in one core.



Equations	Partitioning time [min]	Inversion time (Cholesky) [min]	Schur complement time (CG) [min]	CG steps	Total time [min]
25,010,001	6.2	17.3	4.7	872.0	29.4
50,027,329	13.3	43.7	6.3	1012.0	65.4
75,012,921	20.6	80.2	4.3	1136.0	108.3
100,020,001	28.5	115.1	5.4	1225.0	152.9
125,014,761	38.3	173.5	7.5	1329.0	224.2
150,038,001	49.3	224.1	8.9	1362.0	288.5



Equations	Master process [GB]	Average slave processes [GB]	Slave processes [GB]	Total memory [GB]
25,010,001	4.05	0.41	47.74	51.79
50,027,329	8.10	0.87	101.21	109.31
75,012,921	12.15	1.37	158.54	170.68
100,020,001	16.20	1.88	217.51	233.71
125,014,761	20.25	2.38	276.04	296.29
150,038,001	24.30	2.92	338.29	362.60

Thank you.
Questions?

FEMT:

http://www.cimat.mx/~miguelvargas/FEMT/

Contact:
miguelvargas@cimat.mx
http://www.cimat.mx/~miguelvargas

References

[Chow98]E. Chow, Y. Saad. Approximate Inverse Preconditioners via Sparse-Sparse Iterations. SIAM Journal on Scientific Computing. Vol. 19-3, pp. 995-1023. 1998.

[Chow01]E. Chow. Parallel implementation and practical use of sparse approximate inverse preconditioners with a priori sparsity patterns. International Journal of High Performance Computing, Vol 15. pp 56-74, 2001.

[DAze93]E. F. D'Azevedo, V. L. Eijkhout, C. H. Romine. Conjugate Gradient Algorithms with Reduced Synchronization Overhead on Distributed Memory Multiprocessors. Lapack Working Note 56. 1993.

[Drep07] U. Drepper. What Every Programmer Should Know About Memory. Red Hat, Inc. 2007.

[Farh91] C. Farhat and F. X. Roux, A method of finite element tearing and interconnecting and its parallel solution algorithm, Internat. J. Numer. Meths. Engrg. 32, 1205-1227 (1991)

[Gall90] K. A. Gallivan, M. T. Heath, E. Ng, J. M. Ortega, B. W. Peyton, R. J. Plemmons, C. H. Romine, A. H. Sameh, R. G. Voigt, Parallel Algorithms for Matrix Computations, SIAM, 1990.

[Geor81] A. George, J. W. H. Liu. Computer solution of large sparse positive definite systems. Prentice-Hall, 1981.

[Geor89] A. George, J. W. H. Liu. The evolution of the minimum degree ordering algorithm. SIAM Review Vol 31-1, pp 1-19, 1989.

[Golu96] G. H. Golub, C. F. Van Loan. Matrix Computations. Third edition. The Johns Hopkins University Press, 1996.

[Heat91] M T. Heath, E. Ng, B. W. Peyton. Parallel Algorithms for Sparse Linear Systems. SIAM Review, Vol. 33, No. 3, pp. 420-460, 1991.

[Hilb77] H. M. Hilber, T. J. R. Hughes, and R. L. Taylor. Improved numerical dissipation for time integration algorithms in structural dynamics. Earthquake Eng. and Struct. Dynamics, 5:283–292, 1977.

[Kary99] G. Karypis, V. Kumar. A Fast and Highly Quality Multilevel Scheme for Partitioning Irregular Graphs. SIAM Journal on Scientific Computing, Vol. 20-1, pp. 359-392, 1999.

[Krui04] J. Kruis. "Domain Decomposition Methods on Parallel Computers". Progress in Engineering Computational Technology, pp 299-322. Saxe-Coburg Publications. Stirling, Scotland, UK. 2004.

[MPIF08] Message Passing Interface Forum. MPI: A Message-Passing Interface Standard, Version 2.1. University of Tennessee, 2008.

[Saad03] Y. Saad. Iterative Methods for Sparse Linear Systems. SIAM, 2003.

[Smit96] B. F. Smith, P. E. Bjorstad, W. D. Gropp. Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations. Cambridge University Press, 1996.

[Sori00] M. Soria-Guerrero. Parallel multigrid algorithms for computational fluid dynamics and heat transfer. Universitat Politècnica de Catalunya. Departament de Màquines i Motors Tèrmics. 2000. http://www.tesisenred.net/handle/10803/6678

[Ster95] T. Sterling, D. J. Becker, D. Savarese, J. E. Dorband, U. A. Ranawake, C. V. Packer. BEOWULF: A Parallel Workstation For Scientific Computation. Proceedings of the 24th International Conference on Parallel Processing, 1995.

[Tose05] A. Toselli, O. Widlund. Domain Decomposition Methods - Algorithms and Theory. Springer, 2005.

[Varg10] J. M. Vargas-Felix, S. Botello-Rionda. “Parallel Direct Solvers for Finite Element Problems”. Comunicaciones del CIMAT, I-10-08 (CC), 2010. http://www.cimat.mx/reportes/enlinea/I-10-08.pdf

[Wulf95] W. A. Wulf , S. A. Mckee. Hitting the Memory Wall: Implications of the Obvious. Computer Architecture News, 23(1):20-24, March 1995.

[Yann81] M. Yannakakis. Computing the minimum fill-in is NP-complete. SIAM Journal on Algebraic Discrete Methods, Volume 2, Issue 1, pp 77-79, March, 1981.

[Zien05] O.C. Zienkiewicz, R.L. Taylor, J.Z. Zhu, The Finite Element Method: Its Basis and Fundamentals. Sixth edition, 2005.