



HAL
open science

Parallel sub-structuring methods for solving sparse linear systems on a cluster of gpus

Abal-Kassim Cheik Ahamed, F. Magoules

► **To cite this version:**

Abal-Kassim Cheik Ahamed, F. Magoules. Parallel sub-structuring methods for solving sparse linear systems on a cluster of gpus. 16th International Conference on High Performance and Communications (HPCC 2014), Aug 2014, Paris, France. 10.1109/hpcc.2014.24 . hal-01273941

HAL Id: hal-01273941

<https://hal.science/hal-01273941v1>

Submitted on 12 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Parallel Sub-Structuring Methods for solving Sparse Linear Systems on a cluster of GPU

Abal-Kassim Cheik Ahamed*

Frédéric Magoulès†

Abstract

The main objective of this work consists in analyzing sub-structuring method for the parallel solution of sparse linear systems with matrices arising from the discretization of partial differential equations such as finite element, finite volume and finite difference. With the success encountered by the general-purpose processing on graphics processing units (GPGPU), we develop an hybrid multiGPUs and CPUs sub-structuring algorithm. GPU computing, with CUDA, is used to accelerate the operations performed on each processor. Numerical experiments have been performed on a set of matrices arising from engineering problems. We compare C+MPI implementation on classical CPU cluster with C+MPI+CUDA on a cluster of GPU. The performance comparison shows a speed-up for the sub-structuring method up to 19 times in double precision by using CUDA.

Keywords: Sub-structuring method; Linear algebra; Conjugate Gradient; Parallel and distributed computing; Graphics Processing Unit; GPU Computing; CUDA; Finite element.

1 Introduction

Many engineering problems lead to the computation of large size sparse linear systems arising from the discretization of numerical methods such as finite element, finite volume and finite difference. Iterative Krylov methods are suitable to solve these kind of problems. According to the properties of the matrices involved in the linear systems, Krylov methods do not have the same efficiency [49]. In this paper, we consider problems that leads to symmetric and positive-definite matrices, which therefore places the Conjugate Gradient (CG) method as a suitable and efficient Krylov method. This method requires the computation of linear algebra operations such as scalar product, addition of vectors, matrix-vector multiplication (SpMV) that can be costly in terms of computations on a conventional processor (CPU). The use of GPGPU model can accelerate these operations. GPGPU is more efficient for large size problems. Unfortunately, most graphics cards have very limited memory size. We therefore propose to parallelize the CG algorithm [45, 50, 1] using sub-structuring approach, which corresponds to a natural parallelization methodology and is easy to implement. Then, GPU is used to accelerate the computation of local linear algebra operations. Each sub-structure is associated with a single processor (CPU) and an accelerated device (GPU).

This paper is organized as follows. The first section presents the partitionning of the data. The next section (4) presents the main points of the implementation of the conjugate gradient in parallel for different partitionning, leading to the design of linear algebra operations. Section 5 collects and analyzes the numerical results and analyze them. Matrices and experiments hardware are also presented. Finally, concluding remarks are given in Section 6.

*CUDA Research Center, École Centrale Paris, France.

†CUDA Research Center, École Centrale Paris, France (correspondence, frederic.magoules@hotmail.com).




Table 1: Left (matrix), Right (matrix pattern)

2 GPU: An Effective Accelerator

2.1 GPU computing

Since a few decades, Graphics Processing Units (GPUs) are used to accelerate scientific computation by graphics card hardware. These graphics cards were before used only for graphics applications such as Graphical User Interface. The GPU is a processor with wide computational resources. The rapid improvement of GPU performances has allowed to give to GPU Computing an important place in scientific computing, helped by the flexibility of programming on GPU with language such as CUDA. GPU Computing or GPGPU has become essential in scientific computing when we deal with time consuming of a numerical simulation. Current GPUs are enable to compute simultaneously similar operations by more than one million of threads. The last graphics card of Kepler family, K40, reaches 4.29 teraflops single-precision and 1.43 teraflops double-precision peak floating point performance. The storage memory of the graphics cards have also seen their memory storage evolved in recent years, up to 12 GB for the K40. The performance of GPU algorithms strongly depends on both the configuration of the distribution of the threads on the grid [9] and the memory [23]. As indicated in [9], the performance of sparse matrix-vector multiplication, which is a time consuming linear algebra operation, depends on the structure of the matrix, i.e., the pattern of non-zero value, and the format of the matrix storage. References [24] [10] [11] confirm the influence of the distribution of the threads, the matrix structure and storage format, when solving linear systems with iterative Krylov methods on GPU.

2.2 Sparse matrix formats

Usually, solving partial differential equations by numerical methods such as finite element method (FEM) lead to large and sparse matrix, i.e., only a few elements of the matrix are nonzero. The distribution of non-zero coefficients depends on the features of the original problem. Sparse matrix is called *structured* when the non-zero values form a regular pattern along diagonals, otherwise it is called *unstructured*. The performance of the algorithms strongly depends on the structure of the sparse matrix [5, 4, 7]. In terms of memory storage, sparse matrices are stored in compressed formats, which consists in only allowing memory to their non-zero coefficients. Different data storage structures exist [49] such as Compressed-Sparse Row (CSR), Coordinate (COO), ELLPACK (ELL), Hybrid (HYB), etc. In this work, we consider the CSR format. The CSR format stores the matrix using three one-dimensional arrays, as drawn in Fig. 1. Two arrays of size nnz , AA and JA store respectively the non-zero coefficients of the matrix in consecutive rows and the column indices, i.e., $JA(k)$ is the column index in A matrix of $AA(k)$. The third array, IA , of size $n + 1$, stores pointers to the beginning of each row. $IA(i)$ and $IA(i + 1) - 1$ correspond to the beginning and the end of the i -th row in arrays AA and JA , i.e., $IA(n + 1) = nnz + 1$. An example in CSR format of matrix A (Table 1) is given in Figure 1.

index	1	2	3	4	5	6	7	8	9	10	11
AA (nnz)	-5	14	8	1	2	10	4	2	9	15	7
JA (nnz)	1	2	2	3	1	3	2	4	5	3	5
IA (n+1)	1	3	5	7	10	12	nnz+1				

Figure 1: Compressed sparse row storage format (CSR) of matrix Table 1

3 Graph and Matrix Partitioning

Undoubtedly, parallelism is the future of computing. The main step in parallel processing consists in distributing the data on the cluster processors, what is commonly called parallel distributed computing. In this section we describe how data are distributed among processors for different splitting strategies: band-row, band-column, and sub-structuring splitting. The distribution of data is accomplished as a pre-processing step, independently to the solver code. The data such as matrix, right hand-size, vector solution and local to global, are written into file, and will be input by the solver presented in section 4.

3.1 Band-row splitting

The partition of the set of equations leads to allocate to each processor a band of rows, corresponding to the block of the vectors it treats. In Table 2, these terms are located on a colorful area. The band-row splitting approach consists in partitioning the matrix A of size $n \times n$ into horizontal band matrices. Each processor is in charge of the management of a band-row matrix of size $N_p \times n$ and the associated unknown vector x of size $N_p \times 1$, as drawn in Table 2. This method of partitioning by band-row allows to exhibit a sufficient degree of parallelism properly balanced. For this, it implies to assign to all processors, a block of rows of same size, containing approximately the same number of non-zero coefficients. It unfortunately suffers from a major lack of granularity for implementation on a distributed memory system. The algorithm of band-row partitioning of a given matrix for a given




Table 2: A band-row splitting of matrix A

processor p is presented in Algorithm 1.

Algorithm 1 Band-row partitioning for the p -th band

Require: $A(n \times n, \text{rowptr}, \text{colidx}, \text{val})$: csr matrix
Require: $0 < N_b \leq n$: number of row-bands
Require: $1 < p \leq N_b$: processor number (or band number)
 1: $\text{numb_rows_per_band} \leftarrow n/N_b$
 2: $N_p \leftarrow \text{numb_rows_per_band}$
 3: **if** $p = N_b - 1$ **then**
 4: $N_p \leftarrow n - p \times \text{numb_rows_per_band}$
 5: **end if**
 6: $\text{row_displ} \leftarrow i \times \text{numb_rows_per_band}$
 7: $\text{nnz_displ} \leftarrow A.\text{rowptr}[\text{row_displ}]$
 8: $\text{nnz_count} \leftarrow A.\text{rowptr}[\text{row_displ} + N_p] - A.\text{rowptr}[\text{row_displ}]$
 9: **{Allocate band-row matrix } $A_{\text{band}}(N_p \times n, \text{nnz_count})$ **}**
 10: **{// - copy row indices}**
 11: $A_{\text{band}}.\text{rowptr}[:] \leftarrow A.\text{rowptr}[\text{row_displ} + :]$
 12: **{// - shift row indices to local}**
 13: $A_{\text{band}}.\text{rowptr}[:] \leftarrow A_{\text{band}}.\text{rowptr}[:] - A_{\text{band}}.\text{rowptr}[1]$
 14: $A_{\text{band}}.\text{rowptr}[N_p + 1] \leftarrow \text{nnz_count}$
 15: **{// - copy cols numb and coef}**
 16: $A_{\text{band}}.\text{rowptr}[:] \leftarrow A.[\text{nnz_displ} + :]$
 17: $A_{\text{band}}.\text{val}[:] \leftarrow A.[\text{nnz_displ} + :]$**

3.2 Band-column splitting

As for the band-row splitting, the band-column approach consists in partitioning the matrix A into vertical band matrices. Each processor is in charge of the management of a band-column matrix of size $n \times N_p$. The associated unknown vector x of size $N_p \times 1$ is splitting into horizontal band vectors as in band-row splitting as described in Table 3. Algorithm 2




Table 3: A band-column splitting of matrix A

describes the column-row partitioning procedure of a given matrix for a given processor p . According to the structure of CSR format, the computation of the number of non-zero values of each band requires a particular calculation, unlike the row partitioning. At line 8 of the Algorithm 2, we recover the number of non-zero values computed outside the routine. All non-zero values of all processors are stored into an independent array, which is built using the same test process described at line 14 of Algorithm 2.

3.3 Block-diagonal splitting

Table 4 gives the schema of block-diagonal partitioning. In this section we briefly highlight the particularity of the diagonal block and his computation in parallel. When the product is performed by the matrix, the product of the diagonal block requires only local terms of the

Algorithm 2 Band-column partitionning for the p -th band

Require: $A(n \times n, \text{rowptr}, \text{colidx}, \text{val})$: csr matrix
Require: $0 < N_b \leq n$: number of row-bands
Require: $1 < p \leq N_b$: processor number (or band number)
 1: $\text{numb_cols_per_band} \leftarrow n/N_b$
 2: $N_p \leftarrow \text{numb_cols_per_band}$
 3: **if** $p = N_b - 1$ **then**
 4: $N_p \leftarrow n - p \times \text{numb_cols_per_band}$
 5: **end if**
 6: $\text{col_displ} \leftarrow i \times \text{numb_cols_per_band}$
 7: $\text{nnz_displ} \leftarrow A.\text{rowptr}[\text{row_displ}]$
 8: **{Recover nnzcount of the band}**
 9: **{Allocate band-row matrix } $A_{\text{band}}(n \times N_p, \text{nnzcount})$ **}**
 10: $c \leftarrow 1$
 11: **for** $i = 1$ **to** n **do**
 12: $A_{\text{band}}.\text{rowptr}[i] \leftarrow c$
 13: **for** $k = A_{\text{band}}.\text{rowptr}[i]$ **to** $A_{\text{band}}.\text{rowptr}[i + 1]$ **do**
 14: **if** $\text{nnz_displ} \leq A.\text{colidx}[k] < \text{nnz_displ} + N_p$ **then**
 15: $A_{\text{band}}.\text{colidx}[c] \leftarrow A.\text{colidx}[k] - \text{nnz_displ}$
 16: $A_{\text{band}}.\text{val}[c] \leftarrow A.\text{val}[k]$
 17: $c \leftarrow c + 1$;
 18: **end if**
 19: **end for**
 20: **end for****




Table 4: Block-diagonal partitionning

vector x . In contrast, off-diagonal coefficients require the corresponding terms of the vector x . The diagonal blocks are thick black lines in Figure 4.

The optimal splitting is the one that partitions the mesh into sub-structures of the same size, in order to balance the load with a smallest possible boundary to limit data transfers. Sub-structures should be as spherical as possible topologically, since it is the sphere which has the smaller outer surface.

3.4 Sub-structuring splitting

In order to illustrate the sub-structuring method we consider a problem coming from the finite element discretization of an elliptic partial differential problem. To simplify the analysis, we consider the Laplace equation. However, the analysis can be carried out to any coercive elliptic problem. The model problem for the unknown u , in a bounded domain Ω with homogeneous Dirichlet boundary conditions on the boundary $\partial\Omega = \Gamma$ can be expressed as: for $f \in L^2(\Omega)$, find $u \in H^1(\Omega)$ such that

$$-\nabla^2 u = f \quad \text{in } \Omega \tag{1}$$

$$u = 0 \quad \text{on } \Gamma \tag{2}$$

An equivalent variational formulation of this problem can be formulated as: for $f \in L^2(\Omega)$, find $u \in H_0^1(\Omega)$ such that

$$\forall v \in H_0^1(\Omega), \quad \int_{\Omega} \nabla u \nabla v = \int_{\Omega} f v \quad (3)$$

This problem is well posed, i.e., has one and only one solution. After a Galerkin discretization with finite elements and a choice of nodal basis, the linear system is obtained

$$Su = f \quad (4)$$

where f denotes the right hand side, x the unknown and S the stiffness matrix which is a sparse, symmetric and positive-definite matrix. Conjugate gradient detailed in section 4 is used to solve this linear system.

In practice, mesh partitionning is a crucial step of finite element method. A finite element matrix is associated with a finite element mesh and the elements of the matrix are correlated with the interaction of the basis functions defined in the elements of the mesh. The total matrix is calculated as an assembly of elementary matrices. Let's consider a global domain Ω partitionned into two sub-domains without overlap Ω_1 and Ω_2 , with a shared interface Γ as drawn in Fig. 2. When a suitable numerotation of the degrees of freedom is harnessed,




Figure 2: Two sub-domains

the stiffness matrix of the initial considered model problem can be written as the following matrix:

$$S = \begin{pmatrix} S_{11} & 0 & S_{13} \\ 0 & S_{22} & S_{23} \\ S_{31} & S_{32} & S_{33} \end{pmatrix} \quad (5)$$

It is formulated considering the case where the set of nodes numbered 1 and 2 are respectively associated to the sub-domains Ω_1 and Ω_2 . The last set of nodes numbered 3 corresponds to the interface nodes of both sub-domains. The corresponding linear system for previous matrix (5) can be reformulated as follows

$$\begin{pmatrix} S_{11} & 0 & S_{13} \\ 0 & S_{22} & S_{23} \\ S_{31} & S_{32} & S_{33} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} f_1 \\ f_2 \\ f_3 \end{pmatrix} \quad (6)$$

where $x = (x_1, x_2, x_3)^t$ is the unknown vector and $f = (f_1, f_2, f_3)^t$ is the right hand side. The blocks S_{13} and S_{23} are respectively the transpose matrix of S_{31} and S_{32} , and the blocks S_{11} and S_{22} are symmetric positive-definite whether S was symmetric positive-definite. By assigning the different sub-domains at distinct processors, the local matrices can be formulated in parallel as follows:

$$S_1 = \begin{pmatrix} S_{11} & S_{13} \\ S_{31} & S_{33}^{(1)} \end{pmatrix}, \quad S_2 = \begin{pmatrix} S_{22} & S_{23} \\ S_{32} & S_{33}^{(2)} \end{pmatrix} \quad (7)$$

The blocks $S_{33}^{(1)}$ and $S_{33}^{(2)}$ denotes the interaction between the nodes on the interface Γ , respectively integrated in sub-domains Ω_1 and on Ω_2 , i.e.,

$$S_{33} = S_{33}^{(1)} + S_{33}^{(2)} \quad (8)$$

In practice, the sub-domains Ω_1 and Ω_2 respectively knows the set of nodes (1, 3) and (2,3)). In order to avoid deadlock, the list of neighboring interface is re-ordering using the Welsh-Powell algorithm [53, 8] for graph coloring.

4 Conjugate Gradient (CG) for solving $Ax = b$

Among all iterative methods, the Conjugate Gradient method (CG) is very efficient for solving the linear system $Ax = b$ when A is a symmetric positive-definite matrix. CG algorithm has the advantage of being effective and easy to implement [49]. The algorithm is based on the minimization of the distance of the iterative solution ($\|Ax - b\|_2$) to the exact solution on Krylov subspaces.

If we multiply S by v to obtain Sv , we can then multiply S by the last obtained vector to get S^2v , and etc., so it is trivial to construct a Krylov sequences

$$v, Sv, S^2v, S^3v, S^4v, S^5v, \dots$$

Krylov subspace methods are a large category of iterative methods used to solve linear systems from a wide range of applications. As introduced, in each iteration one or more sparse matrix-vector products is used, and Krylov subspace methods add vector(s) to a basis for one or more Krylov subspace(s)

$$V^{(s)} = \text{Span}\{v_1, Sv_1, \dots, S^{s-1}v_1\} \quad (9)$$

where v_1 is the initial vector solution, which is equal to the initial residual $r_1 = f - Sx_1$.

4.1 General conjugate gradient algorithm

Knowing how data are distributed on each processors, we now present the conjugate gradient method for each type of distribution. Algorithm 3 gives the basic conjugate gradient algorithm with preconditioner. Diagonal preconditioner is considered in this paper. We denote x_1 the initial vector solution, r_k the residual, x_k the solution, p_k the descent direction vector at the iteration k and (\cdot, \cdot) is the Euclidean scalar product. At each iteration, a sparse matrix-vector product, which is the most time consuming operation, has to be performed. The other operations are only scalar products and linear combinations of vectors. The CPU

Algorithm 3 Conjugate Gradient method

- 1: Compute $r_1 = f - Sx_1$, $z_1 = M^{-1}r_1$, set $p_1 = z_1$
 - 2: **for** $k = 1, 2, \dots$ **do**
 - 3: $\alpha_k = (z_k, r_k) / (Sp_k, p_k)$
 - 4: $x_{k+1} = x_k + \alpha_k p_k$
 - 5: $r_{k+1} = r_k - \alpha_k Sp_k$
 - 6: $z_{k+1} = M^{-1}r_{k+1}$
 - 7: $\beta_k = (r_{k+1}, z_{k+1}) / (r_k, z_k)$
 - 8: $p_{k+1} = z_{k+1} + \beta_k p_k$
 - 9: **end for**
-

and GPU code are similar, as described in Algorithm 3, except that in GPU version the linear algebra operations are performed on GPU. Many studies [21, 2, 41, 10, 12, 13, 26, 25] in sparse linear systems demonstrated the effectiveness of GPU Computing compared to sequential and parallel CPU code for large size matrices. The performance are more better when memory [19, 18, 15] is well managed and threading distribution is well tuned [16, 9].

In this paper, we use Alinea, our research group library that offers linear algebra operations in both CPU and GPU CUDA 4.0 [44, 3]. The implementation of this library are analyzed in [10, 9, 11, 12] and have shown its effectiveness and robustness compared to Cusp[6], CUBLAS [42], CUSPARSE [43] for double precision arithmetics.

4.2 Band-row algorithm

4.2.1 Sparse matrix-vector product

The processor that will perform the matrix-vector product for a band-row has only the corresponding terms of the vector x , colored area in Table 2. In order to carry out the sparse matrix-vector, this process needs all the terms of the vector x . The first step consists therefore to collect the terms that lacks, located out of the colored area in Figure 2. As it is the same for all processors, it will therefore be necessary to reconstruct the full vector x on each processor. This operation corresponds to a classic collective exchange, where everyone is both transmitter and receiver.

In this work, instead of using the collective operation, *MPI_Allgather*, including in message passing library (MPI), we use the equivalent *Send/Recv*, with a *left-right* ordering of sending and receiving. For the processor p , the *left-right* ordering consists in respectively sending and receiving to and from $k = p-1, k = p+1, k = p-2, k = p+2, k = p-3, k = p+3, \dots$, if $k > 0$. This process is described in Fig. 3. The number of arithmetical operations

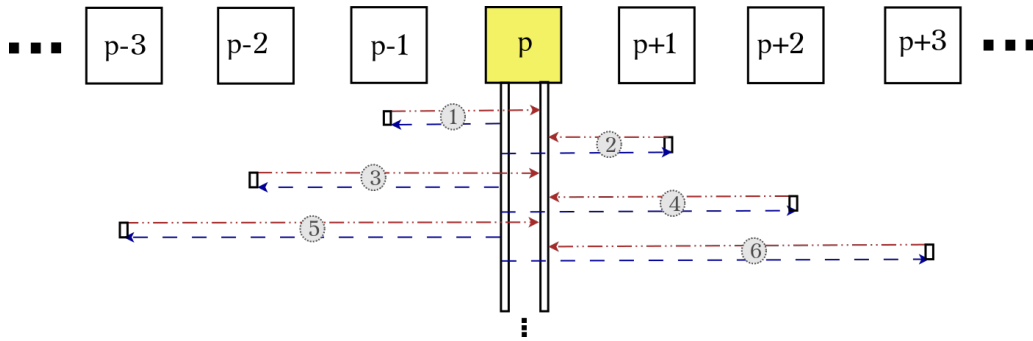


Figure 3: Send/Recv ordering f processor number p

requires to perform the local sparse matrix-vector multiplication, which is approximately $\frac{K \times n}{s}$, where s is the number of processors, n the dimension of the matrix, and K the average number of non-zero coefficients per row. On the other hand, the total number of terms of the vector x to recover before performing the product is approximately $\frac{(s-1) \cdot n}{s}$, if the local matrix has non-zero values in almost all columns. The amount of data is not small compared with the number of arithmetic operations. To optimize communications, it consists in finding a way to limit drastically the number of external values of vector x , located on the others processors, and necessary to compute the product by the matrix.

4.2.2 Basic linear algebra operations

The computation of dot product is a relatively simple operation. Each processor performs local dot product, i.e., multiplies its elements and sums them, from their two local vectors. Finally, the local sums are added using *MPI_Allreduce* with *MPI_SUM* operation. Then, each processor has the global dot product. The operations such as addition of vectors, element wise product, etc. do not change compared to the sequential code. For GPU version, local operations have been performed on graphics card.

4.3 Band-column algorithm

4.3.1 Sparse matrix-vector product

Unlike the band-row splitting sparse matrix-vector multiplication, the SpMV for band-column splitting avoid the exchange of the vector x . However, an $MPI_Allreduce(MPI_SUM)$ is required, in order to assemble the vector $y = Ax$. Note that basic operations are the same as for band-row splitting. For GPU version, local operations have been performed on graphics card. For the matrix-vector product, the results are first send to CPU before applying the same procedure as in CPU.

4.4 Sub-structuring method

4.4.1 Sparse matrix-vector product

Knowing how matrix is partitionned into sub-structures, Section 3.4 now focus on the analysis of the sub-structuring method based on CG algorithm described in Algorithm 3. This algorithm requires to perform a multiplication of the matrix S by a descent direction vector $x = (x_1, x_2, x_3)^t$ at each iteration. With the splitting into two sub-domains, the global matrix-vector multiplication can be written as follows:

$$\begin{aligned} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} &= \begin{pmatrix} S_{11} & 0 & S_{13} \\ 0 & S_{22} & S_{23} \\ S_{31} & S_{32} & S_{33} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} \\ &= \begin{pmatrix} S_{11}x_1 + S_{13}x_3 \\ S_{22}x_2 + S_{23}x_3 \\ S_{31}x_1 + S_{32}x_2 + S_{33}x_3 \end{pmatrix} \end{aligned}$$

Considering the local matrices described equation (7), we can independently compute the both local matrix-vector products as follows

$$\begin{aligned} \begin{pmatrix} y_1 \\ y_3^{(1)} \end{pmatrix} &= \begin{pmatrix} S_{11}x_1 + S_{13}x_3 \\ S_{31}x_1 + S_{33}^{(1)}x_3 \end{pmatrix} \\ \begin{pmatrix} y_2 \\ y_3^{(2)} \end{pmatrix} &= \begin{pmatrix} S_{22}x_2 + S_{23}x_3 \\ S_{32}x_2 + S_{33}^{(2)}x_3 \end{pmatrix} \end{aligned}$$

Since $S_{33} = S_{33}^{(1)} + S_{33}^{(2)}$, and $y_3 = y_3^{(1)} + y_3^{(2)}$. According to this last remark, SpMV can be calculated in two steps:

- calculate the local matrix-vector multiplication in each sub-domain
- assemble on the interface, the local contributions

The first step involves only local data. The second requires the exchange of data between processes dealing with sub-domains with a common interface. In order to assemble interface values of neighboring sub-domains, each processor responsible to a sub-domain must know the description of its interfaces.

4.4.2 Exchange at the interfaces

When a sub-domain Ω_i has several neighboring sub-domains, we denote Γ_{ij} the interface between Ω_i and Ω_j as described in Fig. 4. An interface is identified by its neighboring sub-domains and the equations associated with its nodes. The interface is evaluated from its sub-domains using sparse matrix-vector product. This computation is in two steps for each neighboring sub-domain (Algorithm 4): *collect* the values of the local vector $y = Ax$




Figure 4: Interface description

Algorithm 4 Construct inner buffer and send to neighboring

```

for  $s = 1$  to  $number\_of\_neighboring$  do
  for  $i = 1$  to  $n_s$  do
     $buffer_s(i) = y(list_s(i))$ 
  end for
  Send  $buffer_s$  to neighbour(s)
end for

```

for all interfaces nodes, and then *send* this list to the y vector of the interface equation. The next step consists in updating these changes to all neighboring sub-domains at interfaces equations. First, the contributions of the array containing the result of the matrix-vector product at the interface are received and then values on corresponding interface nodes are updated. This process is described in Algorithm 5. In GPU code, the construction of the inner buffer is carried out on CPU before sending it to the neighboring sub-domain. When

Algorithm 5 Receiving interface results and updating interface equations

```

for  $s = 1$  to  $number\_of\_neighboring$  do
  Receive  $buffer_s$  from neighbour(s)
  for  $i = 1$  to  $n_s$  do
     $y(list_s(i)) = y(list_s(i)) + buffer_s(i)$ 
  end for
end for

```

an equation is shared by several interfaces, the node value of the local vector $y = Ax$ in question is sent to all interfaces to which it belongs. For any number of sub-domains, the mechanism of interface exchange and update is similar to those previously presented. In GPU code, the procedure of exchange has been performed on CPU and then the assembled vector is copied back to GPU, before continuing the algorithm.

The use of the sub-structuring approach in the iterative GC algorithm is inherently parallel and making it excellent candidate for implementation on parallel computers. Indeed, we can distribute the sub-domains over all the available processors and thus compute locally the matrix-vector products, independently in parallel, and use distributed memory in order to limit the memory usage. As explained previously, after the computing of local matrix-vector multiplications, they required to be assembled along the interface. The key ingredient of the data is the local matrix C_i that arise from the finite element discretization, with this approach, each node i only needs to store C_i , the corresponding local matrix to the sub-domain Ω_i , which is only a fraction of the original matrix. The dot product requires that each processor compute a weighted combination of the interface contributions in order to update

their own data. After that, an *MPI_Allreduce* is required to compute the global inner product. Transfers between CPU and GPU at each iteration can decrease the performance of the exchanges algorithms.

This new version of the CG algorithm based on sub-structuring method introduces only two new steps, which reside on data exchange. They consist firstly to share the contributions of local computed SpMV at the interface. Each machine requires to know the list of nodes along the interface and the number of neighboring sub-domains. Secondly, to assemble results over the cluster in order to piece together the local scalar product. This action, realized with MPI, is independant to the splitting. Finally, another advantage of this algorithm is that it can easily be generalised for n sub-domains. This approach presents however two disadvantages. The first drawback arises from an informatic point of view. The granularity, i.e., the number of operations to be performed by the processors compared to the amount of data received or send by the processors may be weak. Indeed, here the granularity is proportional to the number of nodes in the sub-domains compared to the number of nodes on the interface. The number of operations depend on the first parameter and the data transfer depends on the second parameter. If a lot of sub-domains are used, the interface size will not be small compared to the local sub-problem size. This means that the processors realize few computing operations (a local matrix-vector product) and a lot of communications. The second and more important drawback is an algorithmic one. The classical parallel preconditioners per sub-domain are based on incomplete factorization of the local matrices. Such preconditioners are less and less efficient when the number of sub-domain increases. The efficiency due to the parallelism is balanced by a slow convergence of the iterative CG method.

REMARK: A way to avoid this problem, is to use some preconditioners based on domain decomposition methods [20, 47, 52, 51, 28, 35]. Optimized transmissions conditions between the sub-domains are usually defined like in [14, 22, 32, 33, 34, 27, 29, 31] for the Schwarz method, or in [30] for the FETI method. Optimization of the transmissions conditions based on small patch also exists as first introduced in [48, 36, 37, 40] for acoustics and in [38, 39] for linear elasticity. The problem on the interface, is then solved on the CPU with an iterative algorithm. At each iteration, each sub-problem defined in each sub-domain is solved on the GPU with the CG method. This hybrid CPU-GPU domain decomposition method implementation was first proposed in [46] for the FETI method and in [12], for the optimized Schwarz method, but is outside the scope of this paper. In the following we present the results obtained with the CG algorithm issued from a sub-structuring approach as described in the previous paragraphs.

5 Experimental Evaluation

This section reports and presents the evaluation of the a set of numerical experiments.

5.1 Hardware platform

For our performance analysis we considered a machine based on an Intel Core i7 920 2.67Ghz, which has 4 physical cores and 4 logical cores, 12GB RAM, and two different system accelerated generations of nVidia graphics card: GTX275 with 895MB memory, which is double precision compatible.

5.2 Matrices tested

In order to evaluate our analysis on large scale engineering problems, we use a set of matrices from the University of Florida repository [17]. The properties of these matrices are reported in Table 5 where h , nz , $density$, $bandwidth$, $maxrow$, nz/h and $nz/hstddev$ present the size of the matrix, the number of non-zero coefficients, the density corresponding to the number of non-zero values divided by the total number of matrix elements, the upper bandwidth,

the maximum row density, the mean row density and the standard deviation of nz/h . A pattern of the structure matrix is collected in the first column, and the histogram row density is described in the second column.

5.3 Numerical results

The time in seconds (s) of partitioning in row-bands, column-bands and sub-structures are reported in Table 6. The first column gives the name of the matrix. In the second column are collected sub-structuring splitting using the metis software at graph coloring step. The band column splitting are given in third column. The last column gives the band-row splitting. Table 7 reports respectively the running times of sub-structuring CG in seconds (s) (CSR) on CPU and GPU. The experiments have been performed on a single node. Numerical results presented in Table 7 clearly highlight the efficiency of GPU device compared to CPU parallel computation for solving linear systems with sub-structuring methods. The speed up to 19.9x for 4GPUs and 3.3x for 4CPUs (af_shell8). The ratio between sequential CPU and 1GPU reaches 25.7x for the same matrix.

6 Conclusion

This paper gives an analysis of a parallel sub-structuring method based on conjugate gradient method for solving large and sparse linear systems on a cluster of GPU Computing. We have harnessed the efficiency of parallel algorithms, coupled with the high power computing of GPU. We have evaluated a parallelized conjugate gradient algorithm using sub-structuring method, which has a natural approach of parallelization. The experiments have been performed on large sparse matrices arising from large scale engineering problems. The results clearly show the interest of sub-structuring accelerated with GPU Computing to solve linear systems for symmetric positive-definite matrices. The relative gains of the GPU cluster reaches 19x for 4GPUs compared to a 4 CPUs, and up to 27x compared to sequential CPU.

Acknowledgment

The authors acknowledge the CUDA Research Center at Ecole Centrale Paris (France) for its support and for providing the computing facilities.

References

- [1] H. Anzt, V. Heuveline, and B. Rucker. Mixed precision iterative refinement methods for linear systems: Convergence analysis based on Krylov subspace methods. In K. Jónasson, editor, *PARA (2)*, volume 7134 of *Lecture notes in Computer Science*, pages 237–247. Springer, 2010.
- [2] J. M. Bahi, R. Couturier, and L. Z. Khodja. Parallel GMRES implementation for solving sparse linear systems on GPU clusters. In *Proceedings of the 19th High Performance Computing Symposia*, pages 12–19, San Diego, CA, USA, 2011. Society for Computer Simulation International.
- [3] A. Bakhoda, G. Yuan, W. Fung, H. Wong, and T. Aamodt. Analyzing CUDA workloads using a detailed GPU simulator. In *Performance Analysis of Systems and Software, 2009. ISPASS 2009. IEEE International Symposium on*, pages 163–174, April 2009.
- [4] N. Bell and M. Garland. Efficient sparse matrix-vector multiplication on CUDA. Nvidia Technical Report NVR-2008-004, Nvidia Corporation, 2008.

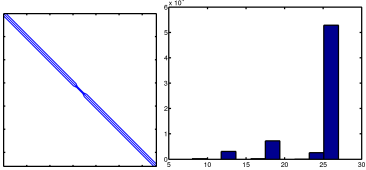
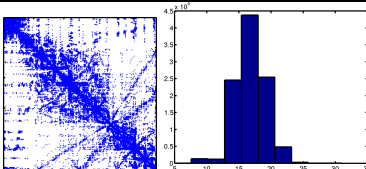
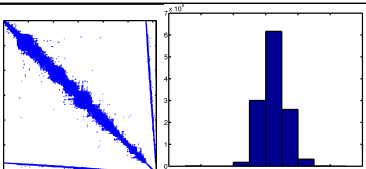
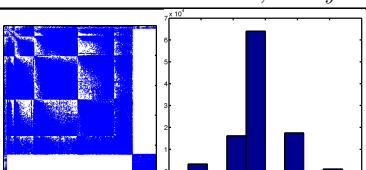
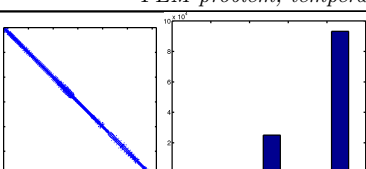
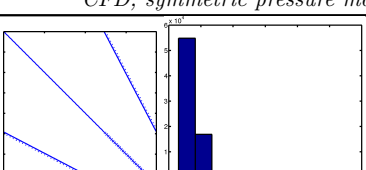
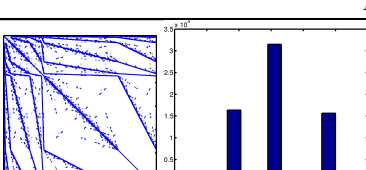
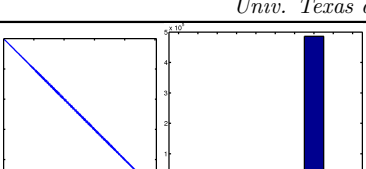
	<p>qa8fm</p> <p>h = size = 66127 density = 0.038 bandwidth = 1048 nonzero = 1660579 nonzero/h = 25.112 max row = 27 nonzero/h stdev = 4.183</p>
<i>3D acoustic FE mass matrix. A. Cunningham, Vibro-Acoustic Sciences Inc.</i>	
	<p>2c_sphere</p> <p>h = size = 101492 density = 0.016 bandwidth = 100407 nonzero = 1647264 nonzero/h = 16.230 max row = 31 nonzero/h stdev = 2.654</p>
<i>FEM, electromagnetics, 2cubes in a sphere. Evan Um, Geophysics, Stanford.</i>	
	<p>thermal2</p> <p>h = size = 1228045 density = 0.001 bandwidth = 1226000 nonzero = 8580313 nonzero/h = 6.987 max row = 11 nonzero/h stdev = 0.811</p>
<i>Unstructured FEM, steady state thermal problem. Dani Schmid, Univ. Oslo.</i>	
	<p>thermo_TK</p> <p>h = size = 102158 density = 0.007 bandwidth = 102138 nonzero = 711558 nonzero/h = 6.965 max row = 10 nonzero/h stdev = 0.715</p>
<i>FEM problem, temperature and deformation of a steel cylinder.</i>	
	<p>cf2</p> <p>h = size = 123440 density = 0.020 bandwidth = 4332 nonzero = 3085406 nonzero/h = 24.995 max row = 30 nonzero/h stdev = 3.888</p>
<i>CFD, symmetric pressure matrix, from Ed Rothberg, Silicon Graphics, Inc.</i>	
	<p>finan512</p> <p>h = size = 74752 density = 0.011 bandwidth = 74724 nonzero = 596992 nonzero/h = 7.986 max row = 55 nonzero/h stdev = 6.278</p>
<i>Economic problem</i>	
	<p>Dubcova2</p> <p>h = size = 65025 density = 0.024 bandwidth = 64820 nonzero = 1030225 nonzero/h = 15.844 max row = 25 nonzero/h stdev = 5.762</p>
<i>Univ. Texas at El Paso, from a PDE solver.</i>	
	<p>af_shell8</p> <p>h = size = 504855 density = 0.007 bandwidth = 4909 nonzero = 17579155 nonzero/h = 34.820 max row = 40 nonzero/h stdev = 1.285</p>
<i>Olaf Schenk, Univ. Basel: AutoForm Eng. GmbH, Zurich. sheet metal forming.</i>	

Table 5: Sketches of finite element matrices

Matrix	metis	band-c	band-r
<i>2-partitionning</i>			
qa8fm	3.16	1.15	1.20
2c_sphere	5.40	2.49	2.49
thermo_TK	1.58	1.03	1.0
cfid2	6.74	3.72	2.06
thermal2	12.53	6.98	6.53
af_shell18	68.05	26.93	27.65
finan512	0.83	0.44	0.42
<i>4-partitionning</i>			
qa8fm	4.60	1.19	1.11
2c_sphere	5.14	2.52	2.49
thermo_TK	2.83	1.03	1.0
cfid2	8.28	2.52	2.04
thermal2	15.72	7.37	7.63
af_shell18	86.18	28.32	51.36
finan512	1.01	0.45	0.42
<i>8-partitionning</i>			
qa8fm	6.64	1.28	1.51
2c_sphere	6.08	2.58	2.49
thermo_TK	1.88	1.07	1.0
cfid2	10.84	2.25	2.61
thermal2	17.33	8.01	7.75
af_shell18	104.41	28.72	27.38
finan512	1.21	0.48	0.42

Table 6: Execution time of partitionning (s)

Matrix	#iter.	1CPU	2CPU	2CPUs	2GPUs	4CPUs	4GPUs	8CPUs	8GPUs
2c_sphere	24	0.386	0.026	0.209	0.047	0.124	0.065	0.125	0.113
af_shell18	2815	374.356	14.549	198.998	23.422	110.662	18.814	107.668	21.385
cfid2	2818	71.078	3.73	38.114	5.904	21.657	8.28	22.111	12.186
Dubcova2	168	1.776	0.128	0.926	0.364	0.54	0.405	0.85	0.64
finan512	15	0.117	0.017	0.121	0.063	0.067	0.071	0.145	0.12
qa8fm	29	0.418	0.023	0.235	0.099	0.198	0.137	0.168	0.115
thermo_TK	13226	141.359	13.214	74.423	22.877	41.979	32.587	40.888	51.528

Table 7: Execution time for parallel sub-structuring CG (s) for CSR format

- [5] N. Bell and M. Garland. Implementing sparse matrix-vector multiplication on throughput-oriented processors. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis (SC'09)*, pages 1–11, New York, NY, USA, 2009. ACM.
- [6] N. Bell and M. Garland. Cusp: Generic parallel algorithms for sparse matrix and graph computations, 2012. Available on line at: <http://cusplibrary.github.io/> (accessed on August 31, 2021).
- [7] J. Bolz, I. Farmer, E. Grinspun, and P. Schröder. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM Trans. Graph.*, 22(3):917–924, July 2003.
- [8] D. Brélaz. New methods to color the vertices of a Graph. *Commun. ACM*, 22(4):251–256, Apr. 1979.
- [9] A.-K. Cheik Ahamed and F. Magoulès. Fast sparse matrix-vector multiplication on gpu for finite element analysis. In *High Performance Computing and Communication 2012 IEEE 9th International Conference on Embedded Software and Systems (HPCC-ICISS), 2012 IEEE 14th International Conference on*, pages 1307–1314. IEEE Computer Society, 2012.
- [10] A.-K. Cheik Ahamed and F. Magoulès. Iterative methods for sparse linear systems on gpu. In *High Performance Computing and Communication 2012 IEEE 9th International Conference on Embedded Software and Systems (HPCC-ICISS), 2012 IEEE 14th International Conference on*, pages 836–842. IEEE Computer Society, june 2012.
- [11] A.-K. Cheik Ahamed and F. Magoulès. Iterative Krylov methods for gravity problems on Graphics Processing Unit. In *Distributed Computing and Applications to Business, Engineering Science (DCABES), 2013 12th International Symposium on*, pages 16–20. IEEE Computer Society, 2013.
- [12] A.-K. Cheik Ahamed and F. Magoulès. Schwarz method with two-sided transmission conditions for the gravity equations on Graphics Processing Unit. In *Distributed Computing and Applications to Business, Engineering Science (DCABES), 2013 12th International Symposium on*, pages 105–109. IEEE Computer Society, 2013.
- [13] A.-K. Cheik Ahamed and F. Magoulès. A stochastic-based optimized Schwarz method for the gravimetry equations on GPU clusters. In *Domain Decomposition Methods in Science and Engineering XXI*. Springer, 2014.
- [14] P. Chevalier and F. Nataf. Symmetrized method with optimized second-order conditions for the Helmholtz equation. In *Domain decomposition methods, 10 (Boulder, CO, 1997)*, pages 400–407. Amer. Math. Soc., Providence, RI, 1998.
- [15] M. Creel and M. Zubair. High performance implementation of an econometrics and financial application on GPUs. In *High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion.;*, pages 1147–1153, Nov. 2012.
- [16] A. Davidson, Y. Zhang, and J. D. Owens. An auto-tuned method for solving large tridiagonal systems on the gpu. In *Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium*, pages 956–965. IEEE, IEEE, May 2011.
- [17] T. A. Davis and Y. Hu. The university of florida sparse matrix collection. *ACM Trans. Math. Softw.*, 38(1):1–25, 2011.
- [18] L. Djinevski, S. Arsenovski, S. Ristov, and M. Gusev. Optimal configuration of GPU cache memory to maximize the performance. In *ICT Innovations 2013 Web Proceedings*, 2013.

- [19] L. Djinevski, S. Arsenovski, S. Ristov, and M. Gusev. Performance drawbacks for matrix multiplication using set associative cache in GPU devices. In *Information & Communication Technology Electronics & Microelectronics (MIPRO), 2013 36th International Convention on*, pages 193–198. IEEE, 2013.
- [20] C. Farhat and F.-X. Roux. A method of finite element tearing and interconnecting and its parallel solution algorithm. *International Journal for Numerical Methods in Engineering*, 32(6):1205–1227, 1991.
- [21] A. Gaikwad and I. Toke. Parallel iterative linear solvers on GPU: A financial engineering case. In *2010 18th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, pages 607–614, Feb. 2010.
- [22] M. J. Gander, L. Halpern, and F. Nataf. Optimized Schwarz methods. In T. Chan, T. Kako, H. Kawarada, and O. Pironneau, editors, *Twelfth International Conference on Domain Decomposition Methods, Chiba, Japan*, pages 15–28, Bergen, 2001. Domain Decomposition Press.
- [23] S. Hong and H. Kim. An analytical model for a GPU architecture with memory-level and thread-level parallelism awareness. *SIGARCH Comput. Archit. News*, 37(3):152–163, June 2009.
- [24] C. Lee, W. W. Ro, and J.-L. Gaudiot. Boosting CUDA applications with CPU-GPU hybrid computing. *Int. J. Parallel Program.*, 42(2):384–404, Apr. 2014.
- [25] N. Li, B. Suchomel, D. Osei-Kuffuor, R. Li, and Y. Saad. Library ITSOL website, November 2010. Available on line at: www-users.cs.umn.edu/~saad/software/ITSOL/index.html (accessed on August 31, 2021).
- [26] R. Li and Y. Saad. GPU-accelerated preconditioned iterative linear solvers, 2010.
- [27] Y. Maday and F. Magoulès. Non-overlapping additive Schwarz methods tuned to highly heterogeneous media. *Comptes Rendus à l’Académie des Sciences*, 341(11):701–705, 2005.
- [28] Y. Maday and F. Magoulès. Absorbing interface conditions for domain decomposition methods: a general presentation. *Computer Methods in Applied Mechanics and Engineering*, 195(29–32):3880–3900, 2006.
- [29] Y. Maday and F. Magoulès. Improved ad hoc interface conditions for Schwarz solution procedure tuned to highly heterogeneous media. *Applied Mathematical Modelling*, 30(8):731–743, 2006.
- [30] Y. Maday and F. Magoulès. Optimal convergence properties of the FETI domain decomposition method. *International Journal for Numerical Methods in Fluids*, 55(1):1–14, 2007.
- [31] Y. Maday and F. Magoulès. Optimized Schwarz methods without overlap for highly heterogeneous media. *Computer Methods in Applied Mechanics and Engineering*, 196(8):1541–1553, 2007.
- [32] F. Magoulès, P. Iványi, and B. Topping. Convergence analysis of Schwarz methods without overlap for the Helmholtz equation. *Computers and Structures*, 82(22):1835–1847, 2004.
- [33] F. Magoulès, P. Iványi, and B. Topping. Non-overlapping Schwarz methods with optimized transmission conditions for the Helmholtz equation. *Computer Methods in Applied Mechanics and Engineering*, 193(45–47):4797–4818, 2004.

- [34] F. Magoulès and R. Putanowicz. Optimal convergence of non-overlapping Schwarz methods for the Helmholtz equation. *Journal of Computational Acoustics*, 13(3):525–545, 2005.
- [35] F. Magoulès and F.-X. Roux. Lagrangian formulation of domain decomposition methods: a unified theory. *Applied Mathematical Modelling*, 30(7):593–615, 2006.
- [36] F. Magoulès, F.-X. Roux, and S. Salmon. Optimal discrete transmission conditions for a non-overlapping domain decomposition method for the Helmholtz equation. *SIAM Journal on Scientific Computing*, 25(5):1497–1515, 2004.
- [37] F. Magoulès, F.-X. Roux, and L. Series. Algebraic way to derive absorbing boundary conditions for the Helmholtz equation. *Journal of Computational Acoustics*, 13(3):433–454, 2005.
- [38] F. Magoulès, F.-X. Roux, and L. Series. Algebraic approximation of Dirichlet-to-Neumann maps for the equations of linear elasticity. *Computer Methods in Applied Mechanics and Engineering*, 195(29–32):3742–3759, 2006.
- [39] F. Magoulès, F.-X. Roux, and L. Series. Algebraic Dirichlet-to-Neumann mapping for linear elasticity problems with extreme contrasts in the coefficients. *Applied Mathematical Modelling*, 30(8):702–713, 2006.
- [40] F. Magoulès, F.-X. Roux, and L. Series. Algebraic approach to absorbing boundary conditions for the Helmholtz equation. *International Journal of Computer Mathematics*, 84(2):231–240, 2007.
- [41] R. Martin, V. Monteiller, D. Komatitsch, S. Perrouty, M. Jessell, S. Bonvalot, and M. Lindsay. Gravity inversion using wavelet-based compression on parallel hybrid CPU/GPU systems: application to southwest ghana. *Geophysical Journal International*, 2013.
- [42] Nvidia Corporation. CUDA toolkit 4.0, CUBLAS Library, 2011. Available on line at: <http://developer.nvidia.com/cuda-toolkit-40> (accessed on August 31, 2021).
- [43] Nvidia Corporation. *CUDA Toolkit 4.0, CUSPARSE Library*, 2011. Available on line at: <http://developer.nvidia.com/cuda-toolkit-40> (accessed on August 31, 2021).
- [44] Nvidia Corporation. *CUDA Toolkit Reference MANUAL*, 4.0 edition, 2011. Available on line at: <http://developer.nvidia.com/cuda-toolkit-40> (accessed on August 31, 2021).
- [45] M. Pakzad, J. L. Lloyd, and C. Phillips. Independent columns: A new parallel ILU preconditioner for the PCG method. *Parallel Computing*, 23(6):637–647, 1997.
- [46] M. Papadrakakis, G. Stavroulakis, and A. Karatarakis. A new era in scientific computing: Domain decomposition methods in hybrid CPU–GPU architectures. *Computer Methods in Applied Mechanics and Engineering*, 200(13-16):1490–1508, Mar. 2011.
- [47] A. Quarteroni and A. Valli. *Domain Decomposition Methods for Partial Differential Equations*. Oxford University Press, Oxford, UK, 1999.
- [48] F.-X. Roux, F. Magoulès, L. Series, and Y. Boubendir. Approximation of optimal interface boundary conditions for two-Lagrange multiplier FETI method. In R. Kornhuber, R. Hoppe, J. Périaux, O. Pironneau, O. Widlund, and J. Xu, editors, *Proceedings of the 15th International Conference on Domain Decomposition Methods, Berlin, Germany, July 21-15, 2003*, Lecture Notes in Computational Science and Engineering (LNCSE). Springer-Verlag, Heidelberg, 2005.

- [49] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition, 2003.
- [50] C. Shen and J. Z. 0001. Robust parallel ILU preconditioning techniques for solving large sparse matrices. In *IPDPS*. IEEE Computer Society, 2002.
- [51] B. Smith, P. Bjorstad, and W. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, UK, 1996.
- [52] A. Toselli and O. Widlund. Domain decomposition methods. *Computational Mathematics*, 34, 2004.
- [53] D. J. A. Welsh and M. B. Powell. An upper bound for the chromatic number of a graph and its application to timetabling problems. *The Computer Journal*, 10(1):85–86, Jan. 1967.