



HAL
open science

Accelerated solution of Helmholtz equation with iterative Krylov methods on GPU

Abal-Kassim Cheik Ahamed, F. Magoules

► **To cite this version:**

Abal-Kassim Cheik Ahamed, F. Magoules. Accelerated solution of Helmholtz equation with iterative Krylov methods on GPU. 16th International Conference on High Performance and Communications (HPCC 2014), Aug 2014, Paris, France. <10.1109/hpcc.2014.16>. <hal-01273942>

HAL Id: hal-01273942

<https://hal.science/hal-01273942v1>

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.



HAL Authorization

Accelerated solution of Helmholtz equation with Iterative Krylov Methods on GPU

Abal-Kassim Cheik Ahamed*

Frédéric Magoulès†

Abstract

This paper gives an analysis and an evaluation of linear algebra operations on Graphics Processing Unit (GPU) with complex number arithmetics with double precision. Knowing the performance of these operations, iterative Krylov methods are considered to solve the acoustic problem efficiently. Numerical experiments carried out on a set of acoustic matrices arising from the modelisation of acoustic phenomena within a cylinder and a car compartment are exposed, exhibiting the performance, robustness and efficiency of our algorithms, with a ratio up to 27x for dot product, 10x for sparse matrix-vector product and solvers in complex double precision arithmetics.

Keywords: Krylov method; Linear algebra operation; Sparse matrix; Graphics Processing Unit; CUDA; Parallel computing; Acoustic; Helmholtz equation

1 Introduction

Amongst other industries, acoustic performance is a major concern in automotive companies. To tackle this, several models are used. In this study, we limit ourselves to acoustic models applying to closed cavities where the acoustic problem is independent from the surrounding structure. We are then able to assume that the pressure field does not have any interactions with the enclosed structure.

Over the past decade, Graphics Processing Units (GPUs) have undergone the same rapid evolution as Central Processing Units (CPUs) in the past. Its extent is such as computers with high Floating point Operations Per Second (Flops) are now widely spread. If when they appeared, the usage of GPUs were focused on computations associated with graphic display, they encountered a considerable evolution ten years later when General Purpose GPUs (GPGPUs) started to use graphic hardware for numerical computations. With the computing power of GPUs drastically increasing, science and engineering started to use GPUs massively. This was also due to NVIDIA releasing Compute Unified Device Architecture (CUDA) [45], which offered a high level GPGPU-based programming language.

Before the migration of GPGPU in 2000, most numerical simulations were carried out on CPU clusters. However, the release of CUDA enabled scientists to take advantage of the high computational power of GPUs: the CPU freed itself from heavy calculations by distributing them to the GPU. The workload became shared out as follows: the CPU managed the sequential blocks of the code while the GPU dealt with the parallel ones. This can be explained by their respective architecture. The simplified architecture of a CPU processor is composed of several memories with multiple levels of cache memories associated, a basic unit of computation and a more complex control unit.

In this paper, we use *Alinea* [10, 9], our own research group library, which implements many algorithms in C++. *Alinea* is a scalable library used both for linear algebra operations and more advanced operations, such as iterative Krylov on both CPU and GPU clusters for

*CUDA Research Center, Applied Mathematics and Systems Laboratory, École Centrale Paris, France.

†CUDA Research Center, Applied Mathematics and Systems Laboratory, École Centrale Paris, France (correspondence, frederic.magoules@hotmail.com).

real and complex number arithmetics in single and double precision. In this paper we are concerned only with complex number arithmetics. We propose an evaluation of the solution of sparse and large size linear systems with matrices, described by complex number arithmetics, and arising from the finite [23] or infinite [3, 2, 1] element discretisation of acoustic problems, modeled with Helmholtz equation. Iterative Krylov are well suited for this kind of problems. Nevertheless, these methods require the computation of many linear algebra operations, in particular sparse matrix-vector product (SpMV), which requires high computational time. To accelerate the solution of these sparse linear systems, a computational effort is required to perform linear algebra operations efficiently.

The remainder of the paper is organised as follows. Section 2 briefly gives the background and motivation of this work. Section 3 describes the industrial test cases used to evaluate and analyze our work. Section 5 presents how complex number arithmetics are designed. In Section 6 we give the evaluation of basic (addition of vectors, scale of vectors, etc.) and advanced (sparse matrix-vector multiplication, etc.) linear algebra operations required to perform iterative Krylov methods. Section 7 presents numerical experiments on iterative Krylov methods which clearly exhibit the robustness and effectiveness of our implementation on GPU for solving acoustic problem. Concluding remarks are given in Section 8

2 Background and Motivation

If Graphics Processing Units (GPUs) were first used exclusively for graphics applications like Graphical User Interface, this has evolved in the last few decades and due to their heavy computational power, GPUs are now widely used to accelerate scientific computation by graphics card hardware. GPUs have become a major tool in scientific computation, thanks both to the rapid pace at which they have improved, and the flexibility offered by languages to program on them such as CUDA, an extension from C/C++. This is especially true when the time consumption of a numerical simulation is taken into account, and GPUs have therefore provided a huge boost to a multitude of applications of science [50] and engineering [15].

Unlike Central Processing Units (CPUs), GPUs have an architecture with multiple arithmetic compute units, as well as several levels of memory. This enables them to compute simultaneously the same operation over a million times. For example, the last graphics card of the Kepler family, the K40, has a 4.29 teraflops single-precision and 1.43 teraflops double-precision peak floating point performance. The K40 also has a 12 GB storage memory, and more generally the storage memory of GPUs have also greatly increased with time. The execution time of a GPU algorithm is strongly impacted by both the configuration of the distribution of the threads on the grid [11] and the memory management [18]. For the most time consuming linear algebra operation, the sparse matrix vector product, the structure of the matrix and the type of the matrix storage format are both of most importance, as shown in [5, 9]. This, along with the importance of the distribution of threads, is also proved in [10] on iterative Krylov methods on GPU used to solve linear systems with real number arithmetics.

The acoustic problem, in its simplest linear form, is governed in the frequency domain by the Helmholtz equation with suitable boundary conditions. When high frequency regime is considered, the matrix of the linear system becomes very large. The problem we are aimed to solve arises from the discretization of the Helmholtz equation in a bounded domain Ω , with a boundary condition considered on the outside boundary $\Gamma = \partial\Omega$. The Helmholtz equation is expressed as follows $-\nabla^2 u - k^2 u = g$, where $k = \frac{2\pi F}{c}$ is the wavenumber associated with the frequency $F \in \mathbb{R}$ and $c \in \mathbb{R}$ denotes the velocity of the medium that is different in space. In this paper, Dirichlet boundary conditions are considered along a part of Γ . The frequency domain in which the solution is sought is usually limited, so as to analyze the acoustic response at specific places of the cavity (for instance around the driver's ears). To carry this out, a suitable numerical model has to be used. For complex geometries two

models can be chosen, depending on the boundary conditions. If there are conditions on all boundaries of the domain, then boundary element (BE) methods can be used. Else, finite element (FE) methods, which are methods based on the domain, are used to solve a weak formulation of the problem. When using FE methods, mesh requirements (around 10 nodes per wavelength are necessary) make the mesh sizes gigantic when dealing with high frequencies. This paper focuses on effectively handling large size acoustic problems using FE method.

3 Case Study: automotive and cylinder acoustic

In this part, we present the finite element meshes used for solving the acoustic problems arising from the automotive industry [35].

We now focus on two numerical examples that enable us to evaluate the performances of our procedures. We consider the study in a car compartment with Audi (Audi3D) and Twingo (Twingo3D), and in a cylinder (3D cylinder). Let's look at the example of a *car compartment*. The goal is to construct the frequency response function at driver's ear from the velocity boundary conditions along the firewall. Understanding this problem can help solve similar problems where the evaluation of the acoustic response to vibrating panels inside a cavity is at stake. Several sources can explain such mechanical vibrations. The vibrations can indeed be air-borne or structural-borne. And the prediction of these vibrations can be a difficult task. In the case of automotive applications, the higher the frequency is, the worst the quality of numerical predictions for mechanical vibrations is. Acoustic predictions depending on the treatment of these mechanical vibrations, precise acoustic predictions are possible only if correct vibration profiles along the car body are provided. According to advanced FE methodologies used on car bodies, computing accurate results becomes difficult when the frequencies are higher than 2500 Hz. Such difficulty to produce correct results at high frequencies can be explained by the complex mechanical structure of a car body. Usual models do not consider parameters which are essential to understand the behavior of a car body at high frequencies, such as the characteristics of the connections, the damping properties, etc. We can note that modifying the models to make them take such parameters into account is not easy. At these high frequencies, some variability effects become important and complicate the predictions. For instance two car bodies produced in the same way, and that could be considered identical, may present drastically different vibro-acoustic behaviors at such frequencies. The meshes of Audi 3D with different refinements (size h) are presented in Figure 1. Figure 2 describes the meshes of Twingo 3D with different refinements (size h). In Figure 3 are illustrated the Cylinder 3D meshes for three refinements (size h).

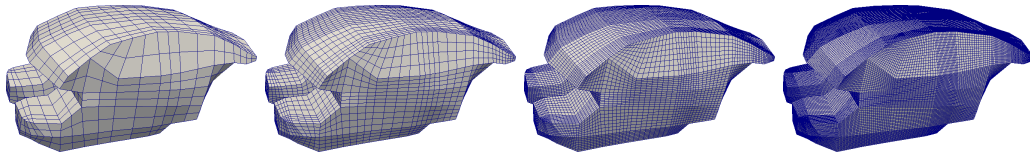


Figure 1: Audi 3D, $h = (0.133425, 0.066604, 0.033289, 0.016643)$

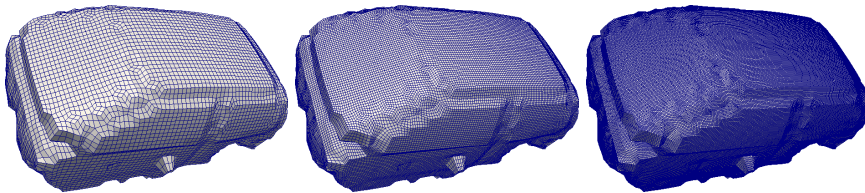


Figure 2: Twingo 3D, $h = (0.077866, 0.038791, 0.019379)$




Figure 3: Cylinder 3D, $h = (0.033973, 0.016949, 0.008342)$

3.1 Matrices tested

For the analysis proposed in this paper, we use matrices coming from the FE discretization of the Helmholtz equation for the study of car compartment and cylinder acoustic problems. These matrices are in stored Compressed Sparse Row (CSR). Table 1 collect a set of matrices respectively arising from Audi3D, Twingo3D and Cylinder3D. These tables give the properties of each considered matrix where h is the size of the square matrix, nz the number of nonzero elements, *density* the density, i.e., the number of nonzero elements divided by the total number of matrix elements, *bandwidth* the upper bandwidth which is equal to lower bandwidth for symmetric matrix, *maxrow* the maximum row density, nz/h the mean row density and $nz/hstddev$ the standard deviation of nz/h . The sparse matrix pattern and the coefficients distribution are respectively reported in the first column and second column.

3.2 Sparse matrix formats

The matrices of the acoustic problem are of large size and are sparse. Most of the elements are zero. The distribution of non-zero values depends on the properties of the initial problem. To take the best advantage of the memory storage, sparse matrices are stored in compressed formats, i.e., only non-zero elements are allocated. It exists divers storage formats [49] such as Compressed-Sparse Row (CSR) [5], Coordinate (COO), ELLPACK (ELL), Hybrid (HYB), etc. In this paper, we have chosen CSR format to store matrices. The matrix is stored in three one-dimensional arrays, as drawn in Figure 4 and corresponds o the matrix shown Table 2. The first two arrays of size nz , AA and JA contain respectively the non-zero coefficients of the matrix in row major order and the column indices, hence $JA(j)$ corresponds to the column index in dense matrix A of the coefficient $AA(j)$. The last array, IA , of size $n + 1$, contains pointers to the start of each row. $IA(i)$ and $IA(i - 1) - 1$ match to the start and the end of the $i - th$ row in arrays AA and JA , i.e., $IA(n + 1) = nz + 1$.

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

Figure 4: Compressed sparse row storage format (CSR) of matrix, TABLE 2

3.3 Hardware configuration

The experiments have been performed on workstation based on an Intel Core i7 920 2.67Ghz, which has 4 physical cores and 4 logical cores, 12GB RAM, and two nVidia graphics card: a Tesla K20c (device #0) with 4799GB memory and GeForce GTX 570 with 1279MB memory (device #1). The cards are double precision compatible. In this paper the devices Tesla K20c and GTX570 are respectively denoted $gpu\#0$ and $gpu\#1$.

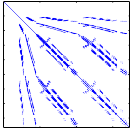
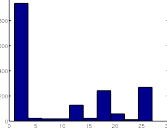
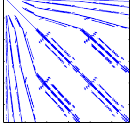
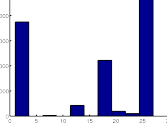
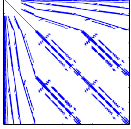
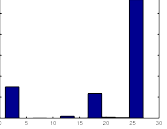
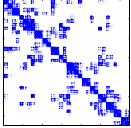
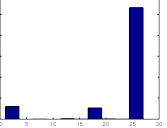
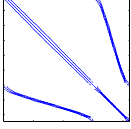
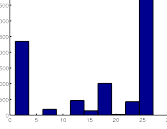
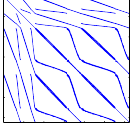
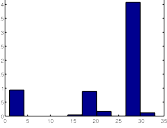
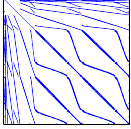
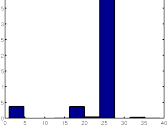
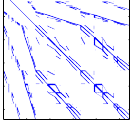
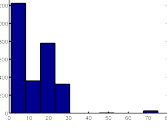
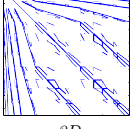
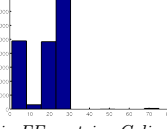
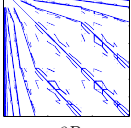
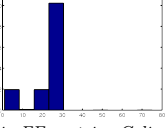
		<p>Audi3D-1 h = size = 1727 density = 0.550 bandwidth = 1436 nonzero = 16393 nonzero/h = 9.492 max row = 27 nonzero/h stddev = 10.205</p>
<i>3D acoustic FE matrix. Audi car (mesh size = 0.133425, length wave = 3.5).</i>		
		<p>Audi3D-2 h = size = 11637 density = 0.139 bandwidth = 11237 nonzero = 188455 nonzero/h = 16.194 max row = 27 nonzero/h stddev = 11.223</p>
<i>3D acoustic FE matrix. Audi car (mesh size = 0.066604, length wave = 3.5).</i>		
		<p>Audi3D-3 h = size = 85001 density = 0.025 bandwidth = 84474 nonzero = 1781707 nonzero/h = 20.961 max row = 27 nonzero/h stddev = 9.832</p>
<i>3D acoustic FE matrix. Audi car (mesh size = 0.033289, length wave = 3.5).</i>		
		<p>Audi3D-4 h = size = 648849 density = 0.004 bandwidth = 520461 nonzero = 15444211 nonzero/h = 23.802 max row = 27 nonzero/h stddev = 7.720</p>
<i>3D acoustic FE matrix. Audi car (mesh size = 0.016643, length wave = 3.5).</i>		
		<p>Twingo3D-0 h = size = 8439 density = 0.202 bandwidth = 6268 nonzero = 143889 nonzero/h = 17.050 max row = 27 nonzero/h stddev = 11.047</p>
<i>3D acoustic FE matrix. Twingo car (mesh size = 0.077866, length wave = 9.5).</i>		
		<p>Twingo3D-1 h = size = 62357 density = 0.035 bandwidth = 53935 nonzero = 1351521 nonzero/h = 21.674 max row = 33 nonzero/h stddev = 9.364</p>
<i>3D acoustic FE matrix. Twingo car (mesh size = 0.038791, length wave = 9.5).</i>		
		<p>Twingo3D-2 h = size = 479169 density = 0.005 bandwidth = 470625 nonzero = 11616477 nonzero/h = 24.243 max row = 39 nonzero/h stddev = 7.233</p>
<i>3D acoustic FE matrix. Twingo car (mesh size = 0.019379, length wave = 9.5).</i>		
		<p>Cylinder3D-0 h = size = 2717 density = 0.420 bandwidth = 2361 nonzero = 30969 nonzero/h = 11.398 max row = 75 nonzero/h stddev = 11.453</p>
<i>3D acoustic FE matrix. Cylinder (mesh size = 0.033973, length wave = 9.5).</i>		
		<p>Cylinder3D-1 h = size = 19041 density = 0.095 bandwidth = 18629 nonzero = 343677 nonzero/h = 18.049 max row = 75 nonzero/h stddev = 11.051</p>
<i>3D acoustic FE matrix. Cylinder (mesh size = 0.016949, length wave = 9.5).</i>		
		<p>Cylinder3D-2 h = size = 142049 density = 0.016 bandwidth = 141289 nonzero = 3151773 nonzero/h = 22.188 max row = 75 nonzero/h stddev = 9.125</p>
<i>3D acoustic FE matrix. Cylinder (mesh size = 0.008342, length wave = 9.5).</i>		

Table 1: Sketches of Audi, Twingo, Cylinder FE matrices

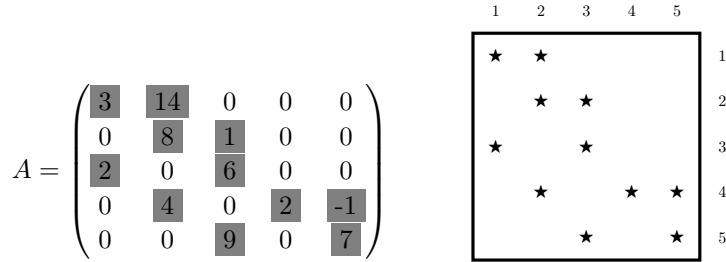


Table 2: Left (matrix), Right (matrix pattern)

The fact that the accuracy of the native clocks of the GPU and the host are respectively a few nanoseconds and a few milliseconds can generate parasites on the measuring of execution time of our programs. To overcome it, we execute the same operations of our benchmark at least 10 times and at least enough times for the total measured time to exceed by at least 100 times the clock accuracy. For the sake of accuracy, we in fact run every operation 100 times, and the reported times correspond to the average time.

4 Handling Complex Numbers on CUDA

The finite element discretization of the Helmholtz equations for acoustic problems leads to complex number arithmetics matrices. GPU have originally been proposed for integer arithmetics. Most numerical simulations need floating point number arithmetics, which decrease the performance of computations. It is even worse in terms of floating point operations when we use numbers with double precision. Since acoustic problems involve complex number arithmetics with double precision floating point operations, the awaited efficiencies are exceedingly low. In order to get the best advantage of GPU architecture, we aim to best optimize the usage of complex number on CUDA. From the viewpoint of numerical calculation, a complex number is defined as a set of two real numbers, which corresponding to real and imaginary part. A natural way to represent a complex number in memory through the following structure

```

struct complex {
    double x; // real part
    double y; // imaginary part
};

```

Listing 1: Complex number in memory

Considering that the previous structure has not padding, both real numbers x and y are contiguous on memory, i.e., the offset of both is null. In our CUDA implementation, we keep the same design as previously. The native library of CUDA, `libcudart.so`, offers `cuDoubleComplex` a `double2` structure (typedef `double2 cuDoubleComplex`), which consists of the same procedure given in Listing 1. For performance outcomes, we prefer to design a complex template class `complex<T>`, which redefines all the operations available in standard `std::complex`. The functions of the standard `complex` are called `host`, i.e., only executable by the CPU. To be executable on GPU, the functions are redefined with `__host__ __device__`, which consists of both CPU and GPU code. We place our complex class template structure into a namespace `stdmrg` in order to avoid confusing it with the standard. This class can be used in both host and device platform. As in all CUDA code, in our library, there exists a copy function that allows to transfer data from host to device, and from device to host. This function offers the possibilities to copy from CPU with `complex<T>` or `stdmrg::complex<T>` to GPU with `stdmrg::complex<T>`. Similarly, from GPU `stdmrg::complex<T>` to `complex<T>` or `stdmrg::complex<T>`.

Distribution of threads is not an automated process. References [16] [9] proved that the threading organization strongly impacts the performance of the numerical algorithm. GPU implementations using advanced gridification auto-tuning techniques, as developed in [10, 9] are thus used in the following.

5 Basic Linear Algebra Operations

This section presents basic linear algebra algorithms, including assign of a vector, scale of vectors, element wise product, addition of vectors and dot product. We also collect the numerical experiments results we have performed to analyze the speed-up of the GPU code compared to the CPU code for complex number arithmetics with double precision.

Table 3 and Figure 5 give the complex double precision execution time in milliseconds (ms) of the *assign operation*, where h represents the size of the vector.

h	cpu	cpu	gpu#0	gpu#0	gpu#1	gpu#1	ratio#0	ratio#1
	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>cpu/#0</i>	<i>cpu/#1</i>
100,000	0.10	0.96	0.07	1.47	0.07	1.41	1.53	1.47
500,000	0.75	0.67	0.13	3.75	0.14	3.60	5.59	5.37
1,000,000	2.04	0.49	0.22	4.65	0.26	3.91	9.49	7.98
8,000,000	15.71	0.51	1.61	4.96	1.56	5.12	9.74	10.06
10,000,000	20.00	0.50	1.75	5.70	1.92	5.20	11.40	10.40
15,000,000	27.50	0.55	2.56	5.85	3.03	4.95	10.73	9.08

Table 3: Complex double precision Assign of vector (ZASSIGN)

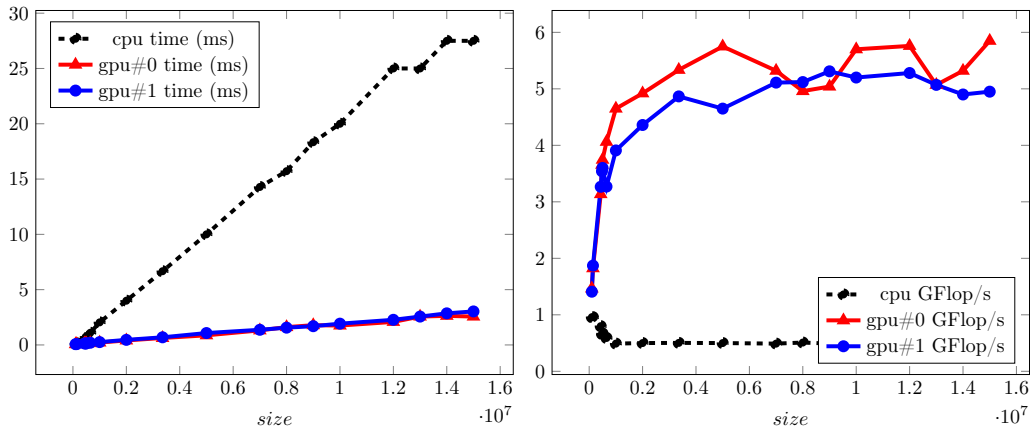


Figure 5: ZASSIGN [left: time in ms, right: GFlops]

The scale *scale operation* kernel is described as follows

```

__global__ void Scal(
    stdmrg::complex<double> alpha,
    const stdmrg::complex<double>* d_x, int size) {
    unsigned int x = blockDim.x * blockIdx.x + threadIdx.x;
    unsigned int y = threadIdx.y + blockDim.y * blockIdx.y;
    int pitch = blockDim.x * gridDim.x;
    int idx = x + y * pitch;
    if ( idx < size ) {
        d_x[idx] = alpha * d_x[idx];
    }
}

```

The running times of the scale *scale* operation for different size of vectors are given in Table 4 and drawn in Figure 6, where h represents the size of the vectors.

h	cpu <i>time (ms)</i>	cpu <i>Gflops</i>	gpu#0 <i>time (ms)</i>	gpu#0 <i>Gflops</i>	gpu#1 <i>time (ms)</i>	gpu#1 <i>Gflops</i>	ratio#0 <i>cpu/#0</i>	ratio#1 <i>cpu/#1</i>
100,000	0.81	0.74	0.07	8.60	0.07	8.26	11.56	11.10
500,000	4.17	0.72	0.15	19.50	0.17	17.16	27.08	23.83
1,000,000	8.33	0.72	0.29	20.70	0.30	20.04	28.75	27.83
8,000,000	65.00	0.74	1.75	27.36	2.22	21.60	37.05	29.25
10,000,000	85.00	0.71	2.38	25.20	2.56	23.40	35.70	33.15
15,000,000	120.00	0.75	3.13	28.80	3.70	24.30	38.40	32.40

Table 4: Complex double precision Scale of vectors (ZSCAL)

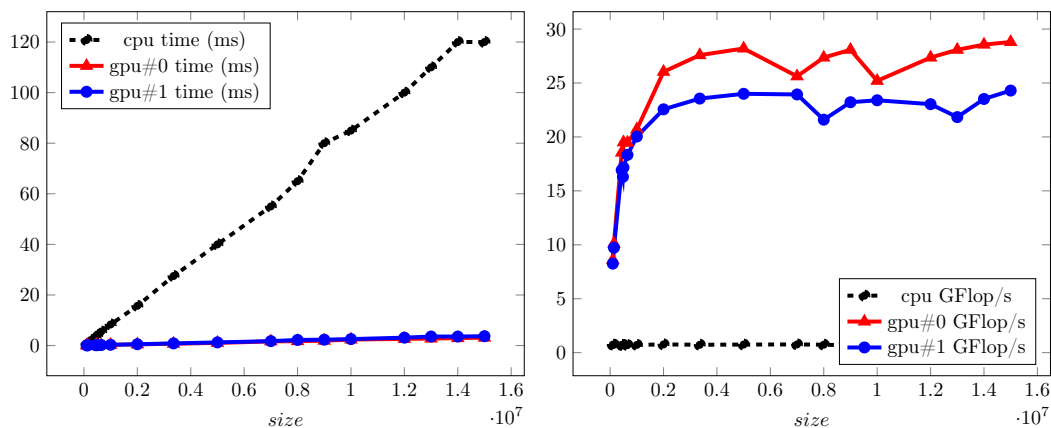


Figure 6: ZSCAL [left: time in ms, right: GFlops]

Double-precision complex Alpha X Plus Y (Zaxpy), i.e., $y[i] = \alpha \times x[i] + y[i]$, is a level one (vector) operation between two complex number arithmetics vectors in the Basic Linear Algebra Subprograms (BLAS) package. The addition of vectors is performed by a simple GPU kernel of the form:

```

__global__ void Daxpy(stdmrg::complex<double> alpha,
                    const stdmrg::complex<double>* d_x,
                    stdmrg::complex<double>* d_y, int size) {
    unsigned int x = blockIdx.x * blockDim.x + threadIdx.x;
    unsigned int y = threadIdx.y + blockIdx.y * blockDim.y;
    int pitch = blockDim.x * gridDim.x;
    int idx = x + y * pitch;
    if ( idx < size ) {
        d_y[idx] = alpha * d_x[idx] + d_y[idx];
    }
}

```

The vector containing the final result of the addition overwrites the contents of the second vector operand d_y .

Table 5 and Figure 7 report the complex number arithmetics with double precision execution time of our implementation for the *Zaxpy* operation.

The *element wise product* or *element by element product*, i.e., $y[i] = x[i] \times y[i]$, is intuitively parallel placing it as an excellent candidate for application on GPU. The product result of the elements of both vectors overwrites the corresponding elements of the second vector by a simple GPU kernel, presented in the following:

h	cpu	cpu	gpu#0	gpu#0	gpu#1	gpu#1	ratio#0	ratio#1
	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>cpu/#0</i>	<i>cpu/#1</i>
100,000	0.83	0.97	0.08	9.46	0.09	9.18	9.77	9.49
500,000	4.17	0.96	0.23	17.48	0.23	17.16	18.21	17.88
1,000,000	8.33	0.96	0.42	19.12	0.38	20.88	19.92	21.75
8,000,000	65.00	0.98	2.70	23.68	2.86	22.40	24.05	22.75
10,000,000	85.00	0.94	3.23	24.80	3.70	21.60	26.35	22.95
15,000,000	130.00	0.92	5.00	24.00	5.56	21.60	26.00	23.40

Table 5: Complex double precision Addition of vectors (ZAXPY)

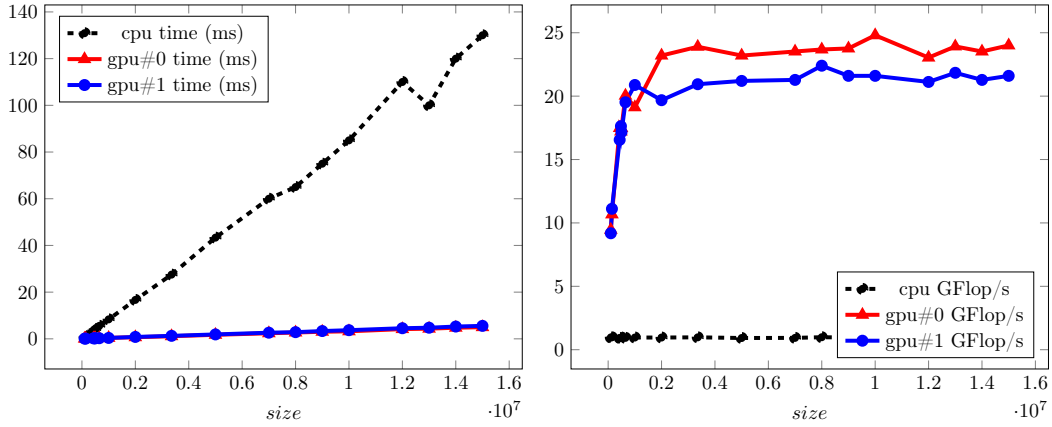


Figure 7: ZAXPY [left: time in ms, right: GFlops]

```

__global__ void ElementWiseProduct(
    stdmrg::complex<double> alpha,
    const stdmrg::complex<double>* d_x,
    stdmrg::complex<double>* d_y, int size) {
    unsigned int x = blockIdx.x * blockDim.x + threadIdx.x;
    unsigned int y = threadIdx.y + blockIdx.y * blockDim.y;
    int pitch = blockDim.x * gridDim.x;
    int idx = x + y * pitch;
    if ( idx < size ) {
        d_y[idx] = d_x[idx] * d_y[idx];
    }
}

```

Table 6 and Figure 8 shows the double precision execution time of our implementation for the *element wise product* operation.

Table 6: Complex double precision Element wise product (ZAXMY)

h	cpu	cpu	gpu#0	gpu#0	gpu#1	gpu#1	ratio#0	ratio#1
	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>cpu/#0</i>	<i>cpu/#1</i>
100,000	1.37	0.44	0.09	6.53	0.09	6.66	14.92	15.21
500,000	6.25	0.48	0.21	14.04	0.22	13.38	29.25	27.88
1,000,000	13.75	0.44	0.37	16.02	0.42	14.34	36.71	32.86
8,000,000	100.00	0.48	3.13	15.36	3.13	15.36	32.00	32.00
10,000,000	130.00	0.46	3.45	17.40	3.70	16.20	37.70	35.10
15,000,000	190.00	0.47	5.00	18.00	5.26	17.10	38.00	36.10

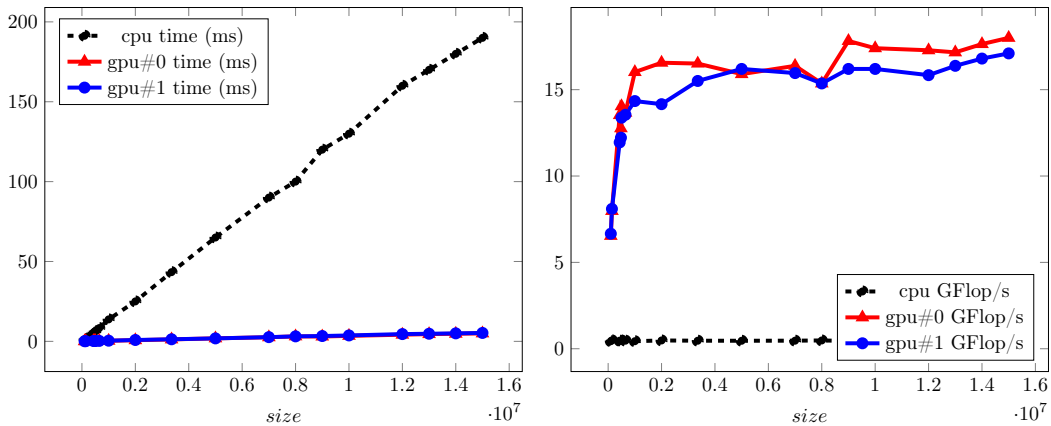


Figure 8: ZAXMY [left: time in ms, right: GFlops]

One very expensive operation to do on CPUs for large vectors is the dot product. However, the most basic implementation of this operation, which uses a simple loop with simultaneous sums, is not very efficient on GPUs. This is why we split the dot product algorithm into two separate parts. The first one consists of the element by element parallel multiplication of the vectors, while the second one consists of the summation of all the results given by the first part. On a sequential processor, to implement the second part, a simple loop where a single variable is incremented is enough. Each element of the input data is handled by a thread, and the current sum (the partial result which is the sum of the n^{th} first elements) is stored in the first thread of this block at the end of the reduction. To obtain the result of the dot product, the algorithm then returns the sum of all the partial sums of the different blocks. In Table 7 and Figure 9, we compare the double precision execution time of our implementation for the *dot product* on both CPU and GPU.

h	cpu	cpu	gpu#0	gpu#0	gpu#1	gpu#1	ratio#0	ratio#1
	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>cpu/#0</i>	<i>cpu/#1</i>
100,000	0.88	0.91	0.12	6.63	0.12	6.43	7.27	7.05
500,000	4.55	0.88	0.27	14.68	0.29	13.92	16.68	15.82
1,000,000	9.09	0.88	0.42	19.04	0.43	18.40	21.64	20.91
8,000,000	70.00	0.91	2.70	23.68	2.86	22.40	25.90	24.50
10,000,000	90.00	0.89	3.45	23.20	3.57	22.40	26.10	25.20
15,000,000	130.00	0.92	5.26	22.80	5.56	21.60	24.70	23.40

Table 7: Complex double precision Dot product (ZDOT)

The results of the norm operation are given in Table 8 and Figure 10.

As we can see in the presented results, GPU is clearly more effective than CPU with complex number arithmetics in double precision. In addition, for the operations considered on two different GPUs, the results are slightly similar, even if the first GPU is a little bit more efficient than the second. In the following, all experiments are performed on the second device, i.e., `gpu#0`.

6 Advanced Linear Algebra Operations

Many methods, such as finite element analysis, require handling widely large size sparse matrices, i.e., only few elements are non-zero values. To store efficiently these matrices on GPU memory, different structures exist, e.g., Compressed-sparse Row format. These structures

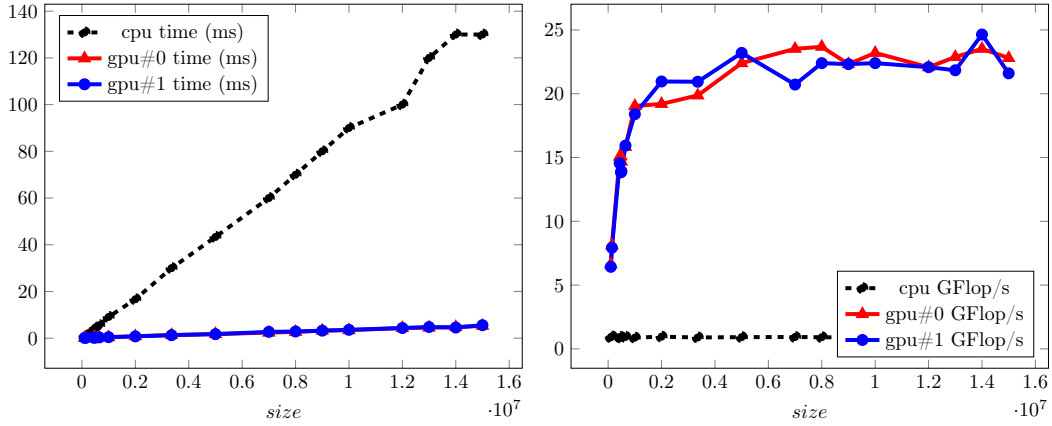


Figure 9: ZDOT [left: time in ms, right: GFlops]

h	cpu	cpu	gpu#0	gpu#0	gpu#1	gpu#1	ratio#0	ratio#1
	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>cpu/#0</i>	<i>cpu/#1</i>
100,000	1.72	0.29	0.13	3.85	0.13	3.97	13.28	13.69
500,000	7.69	0.33	0.25	9.98	0.23	10.95	30.69	33.69
1,000,000	16.67	0.30	0.43	11.65	0.36	14.05	38.83	46.83
8,000,000	140.00	0.29	2.78	14.40	2.13	18.80	50.40	65.80
10,000,000	170.00	0.29	3.45	14.50	2.78	18.00	49.30	61.20
15,000,000	260.00	0.29	5.00	15.00	4.00	18.75	52.00	65.00

Table 8: Complex double precision NormL2 (ZNORM)

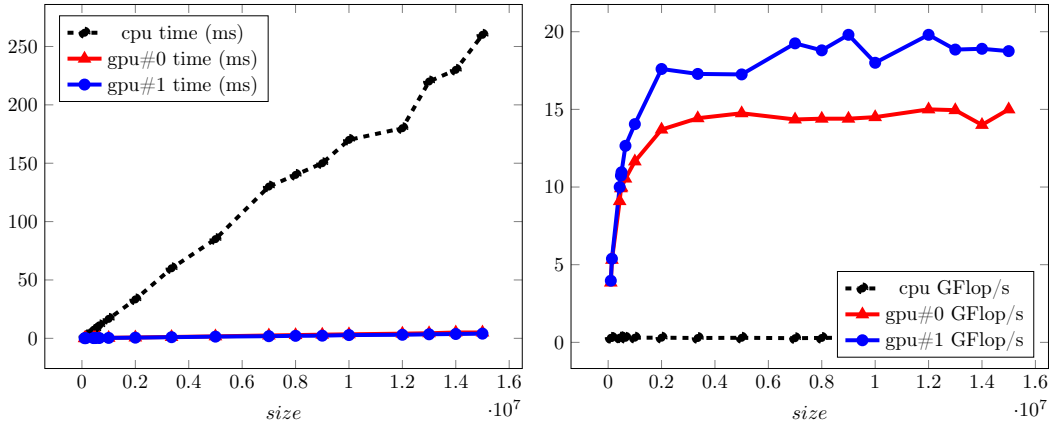


Figure 10: ZNORM [left: time in ms, right: GFlops]

are very important for Sparse matrix-vector product (SpMV) operation for instance, which is one of the most time consuming operation in sparse matrix computation. As explained before, the matrix-vector product proposed in this paper uses advanced auto-tuned techniques to organize threads on the CUDA grid. References [42, 54, 53, 24, 46, 17] clearly demonstrated the efficiency of SpMV on GPU compared to CPU for real number arithmetics. Complex number arithmetics with double precision still remains a challenge.

We now reports in Table 9 the SpMV execution time and the number of floating operations per second when using the CSR format within *Alinea* for complex number arithmetics

with double precision. Numerical experiments clearly show that GPU operations are efficient

problem	cpu	cpu	gpu#0	gpu#0	gpu#1	gpu#1	ratio#0	ratio#1
	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>time (ms)</i>	<i>Gflops</i>	<i>cpu/#0</i>	<i>cpu/#1</i>
Audi3D-0	0.01	0.61	0.07	0.12	0.06	0.13	0.19	0.21
Audi3D-1	0.20	0.67	0.11	1.23	0.12	1.07	1.84	1.60
Audi3D-2	2.22	0.68	0.37	4.03	0.42	3.56	5.93	5.24
Audi3D-3	20.00	0.71	2.22	6.41	3.03	4.70	9.00	6.60
Audi3D-4	180.00	0.69	18.33	6.74	24.00	5.15	9.82	7.50
Twingo3D-0	1.67	0.69	0.28	4.06	0.33	3.45	5.88	5.00
Twingo3D-1	15.71	0.69	1.79	6.05	2.44	4.43	8.80	6.44
Twingo3D-2	140.00	0.66	14.29	6.51	16.67	5.58	9.80	8.40
Cylinder3D-0	0.37	0.67	0.13	1.90	0.15	1.67	2.84	2.51
Cylinder3D-1	3.70	0.74	0.53	5.20	0.69	3.96	7.00	5.33
Cylinder3D-2	36.67	0.69	4.00	6.30	5.26	4.79	9.17	6.97

Table 9: Double precision CSR matrix-vector multiplication

than CPU operations.

7 Iterative Krylov methods

Knowing the efficiency of GPU architecture to perform linear algebra operations for complex number arithmetics with double precision, we now extend our analysis to Krylov iterative methods [25, 42, 4, 54]. We have thus aimed to develop a preconditioned bi-conjugate gradient stabilized method (P-Bi-CGSTAB), a preconditioned P-BiCGSTAB parametered (l) and a preconditioned transpose-free quasi-minimal residual method (P-tfQMR) [49], with optimized CUDA and dynamic auto-tuning on GPU. Reference [9] exhibits the effectiveness of our template implementation for real number arithmetics compared to Cusp [6], CUBLAS [43], CUSPARSE [44].

In general, the transfer of data between the host (CPU) and the device (GPU), is one of the most time consuming aspects of the algorithm [8]. In our Krylov algorithms implementation, all inputs are sent once from the CPU to the GPU before the iterative part starts. However, at each iteration, there has to be more than one dot product (or norm), which means that data must be copied from GPU to CPU. The Krylov algorithm we have developed is launched on the host (CPU) but all the computing steps such as Zdot, Znorm, Zaxpy, or SpMV take place on the device (GPU). The implementations both on CPU and GPU are rigorously the same. All the results of the iterative Krylov methods presented in this paper are achieved with a residual tolerance threshold of 1×10^{-9} , an initial guess of zero and a maximum number of iterations equals to 1000 for all of the presented methods. The aim of this experiments consists in comparing the analysis both on CPU and GPU, with the same code, for complex number arithmetics with double precision. Table 10 shows the speed-up obtained for Audi3D. In Table 11 the speed-up for Twingo3D are collected. The CPU and GPU times in second, and the corresponding ratio are reported in Table 12. The numerical results of the considered iterative Krylov methods confirm the efficiency of GPU computations compared to CPU for solving sparse linear systems. In addition, the ratio increases when the size of the problem increases for all cases, i.e., when the mesh size is more fine, GPU is more efficient compared to CPU. Unfortunately, when the mesh is too fine, the corresponding assembled matrix becomes too large for GPU memory. . . In this case, domain decomposition method [51], [47], [52], [19], [37], [30] based on iterative methods is an

problem	#iter	CPU time (s)	GPU time (s)	speed-up
<i>P-BiCGSTAB</i>				
Audi3D-1	21	0.01	0.030	0.33
Audi3D-2	53	0.24	0.106	2.26
Audi3D-3	94	4.01	0.703	5.71
Audi3D-4	183	85.70	9.209	9.31
<i>P-BiCGSTAB(8)</i>				
Audi3D-1	6	0.03	0.110	0.27
Audi3D-2	12	0.52	0.286	1.82
Audi3D-3	31	12.47	2.162	5.77
Audi3D-4	70	266.26	30.100	8.85
<i>P-TFQMR</i>				
Audi3D-1	24	0.02	0.040	0.50
Audi3D-2	52	0.27	0.113	2.40
Audi3D-3	99	4.71	0.755	6.24
Audi3D-4	214	102.17	10.786	9.47

Table 10: Speed-up of Audi3D

problem	#iter	CPU time (s)	GPU time (s)	speed-up
<i>P-BiCGSTAB</i>				
Twingo3D-0	563	1.85	1.008	1.84
Twingo3D-1	1000	29.45	5.730	5.14
Twingo3D-2	1000	295.66	37.670	7.85
<i>P-BiCGSTAB(8)</i>				
Twingo3D-0	1000	31.2	20.970	1.49
Twingo3D-1	1000	273.81	54.630	5.01
Twingo3D-2	1000	2559.67	324.500	7.89
<i>P-TFQMR</i>				
Twingo3D-0	366	1.34	0.626	2.14
Twingo3D-1	954	30.4	5.438	5.59
Twingo3D-2	1000	318.93	38.090	8.37

Table 11: Speed-up of Twingo3D

problem	#iter	CPU time (s)	GPU time (s)	speed-up
<i>P-BiCGSTAB</i>				
Cylinder3D-0	49	0.040	0.060	0.67
Cylinder3D-1	86	0.680	0.216	3.15
Cylinder3D-2	162	12.290	1.886	6.52
<i>P-BiCGSTAB(8)</i>				
Cylinder3D-0	10	0.090	0.182	0.50
Cylinder3D-1	19	1.450	0.521	2.79
Cylinder3D-2	71	47.670	7.460	6.39
<i>P-QMR</i>				
Cylinder3D-0	85	0.090	0.106	0.85
Cylinder3D-1	196	2.050	0.482	4.25
Cylinder3D-2	282	23.870	3.302	7.23

Table 12: Speed-up of Cylinder3D

issue. The Schwarz method [26], [27], [28], [7] have encountered a success for solving large size problem. To faster the convergence, many references [14, 22, 32, 31, 29] exhibit the interest of optimizing the interface conditions between the subdomains. In order to use this approach for acoustic problems modeled by the Helmholtz equation, continuous optimized interface conditions between the subdomains must be developed as in [33], [34], [36], [20]. Alternative discrete optimization techniques as introduced in [48], [41], [38], [21], [40], [39].

Authors present how domain decomposition method is designed efficiently on GPU in [12, 13] and proved in [13] the interest of Schwarz methods on a cluster of GPUs.

8 Concluding remarks

This paper presented the performance evaluation and analysis of linear algebra operations together with their uses within Krylov methods for solving acoustic problem on Graphics Processing Unit (GPU) for complex number arithmetics with double precision. The numerical experiments have been performed in two different system accelerated generations of nVidia graphics card: GTX570 and Tesla K20c. We have used matrices arising from the finite element modeling of the acoustic within a cylinder and a car compartment. The presented results clearly demonstrate the interest of the use of GPU device to compute linear algebra operations, and outline the robustness, performance and efficiency of solve the Helmholtz equations for acoustic problems.

References

- [1] J.-C. Autrique and F. Magoulès. Numerical analysis of a coupled finite-infinite element method for exterior Helmholtz problems. *Journal of Computational Acoustics*, 14(1):21–43, 2006.
- [2] J.-C. Autrique and F. Magoulès. Studies of an infinite element method for acoustical radiation. *Applied Mathematical Modelling*, 30(7):641–655, 2006.
- [3] J.-C. Autrique and F. Magoulès. Analysis of a conjugated infinite element method for acoustic scattering. *Computers and Structures*, 85(9):518–525, 2007.
- [4] 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.
- [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. Library cusp website, 2010. Available on line at: <http://cusplibrary.github.io/> (accessed on December 14, 2021).
- [7] X.-C. Cai, M. A. Casarin, F. W. E. Jr, and O. B. Widlund. Overlapping schwarz algorithms for solving helmholtz’s equation. In *Domain decomposition methods, 10 (Boulder, CO, 1997)*, page 391–399. Amer. Math. Soc., Providence, RI, 1998.
- [8] A. F. Camargos, V. C. Silva, J.-M. Guichon, and G. Meunier. Iterative solution on gpu of linear systems arising from the a-v edge-fea of time-harmonic electromagnetic phenomena. In *Parallel, Distributed and Network-Based Processing (PDP), 2014 22nd Euromicro International Conference on*, pages 365–371, Feb 2014.

- [9] A.-K. Cheik Ahamed and F. Magoulès. Fast sparse matrix-vector multiplication on graphics processing unit 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 graphics processing unit. 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] M. M. Dehnavi, D. M. Fernandez, and D. Giannacopoulos. Finite-element sparse matrix vector multiplication on graphic processing units. *IEEE*, 2010.
- [18] 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.
- [19] 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.
- [20] M. Gander, L. Halpern, and F. Magoulès. An optimized Schwarz method with two-sided Robin transmission conditions for the Helmholtz equation. *International Journal for Numerical Methods in Fluids*, 55(2):163–175, 2007.
- [21] M. Gander, L. Halpern, F. Magoulès, and F.-X. Roux. Analysis of patch substructuring methods. *International Journal of Applied Mathematics and Computer Science*, 17(3):395–402, 2007.

- [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] I. Harari and F. Magoulès. Numerical investigations of stabilized finite element computations for acoustics. *Wave Motion*, 39(4):339–349, 2004.
- [24] M. Kreutzer, G. Hager, G. Wellein, H. Fehske, A. Basermann, and A. R. Bishop. Sparse matrix-vector multiplication on gpgpu clusters: A new storage format and a scalable implementation. *CoRR*, abs/1112.5588, 2011.
- [25] R. Li and Y. Saad. GPU-accelerated preconditioned iterative linear solvers, 2010.
- [26] P.-L. Lions. On the Schwarz alternating method. I. In R. Glowinski, G. H. Golub, G. A. Meurant, and J. Périaux, editors, *First International Symposium on Domain Decomposition Methods for Partial Differential Equations*, pages 1–42, Philadelphia, PA, 1988. SIAM.
- [27] P.-L. Lions. On the Schwarz alternating method. II. In T. Chan, R. Glowinski, J. Périaux, and O. Widlund, editors, *Domain Decomposition Methods*, pages 47–70, Philadelphia, PA, 1989. SIAM.
- [28] P.-L. Lions. On the Schwarz alternating method. III: a variant for nonoverlapping subdomains. In T. F. Chan, R. Glowinski, J. Périaux, and O. Widlund, editors, *Third International Symposium on Domain Decomposition Methods for Partial Differential Equations, held in Houston, Texas, March 20-22, 1989*, Philadelphia, PA, 1990. SIAM.
- [29] 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.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] F. Magoulès, K. Meerbergen, and J.-P. Coyette. Application of a domain decomposition method with Lagrange multipliers to acoustic problems arising from the automotive industry. *Journal of Computational Acoustics*, 8(3):503–521, 2000.
- [36] 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.

- [37] F. Magoulès and F.-X. Roux. Lagrangian formulation of domain decomposition methods: a unified theory. *Applied Mathematical Modelling*, 30(7):593–615, 2006.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] K. K. Matam and K. Kothapalli. Accelerating sparse matrix vector multiplication in iterative methods using GPU. In G. R. Gao and Y.-C. Tseng, editors, *ICPP*, pages 612–621. IEEE, 2011.
- [43] Nvidia Corporation. Nvidia - library cublas. Available on line at: http://www.nvidia.com/object/cuda_home_new.html (accessed on December 14, 2021).
- [44] Nvidia Corporation. *CUDA Toolkit 4.0, CUSPARSE Library*, 2011. Available on line at: <http://developer.nvidia.com/cuda-toolkit-40> (accessed on December 14, 2021).
- [45] Nvidia Corporation. *CUDA Toolkit Reference MANUAL*, 4.0 edition, 2011. Available on line at: <http://developer.nvidia.com/cuda-toolkit-40> (accessed on December 14, 2021).
- [46] T. Oberhuber, A. Suzuki, and J. Vacata. New row-grouped csr format for storing the sparse matrices on gpu with implementation in cuda. *CoRR*, abs/1012.2270, 2010.
- [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. K. et al, editor, *Proceedings of the 15th Int. Conf. on Domain Decomposition Methods, Berlin, Germany, Jul.21-15, 2003*, Lecture Notes in Computational Science and Engineering (LNCSE). Springer-Verlag, Haidelberg, 2005.
- [49] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition, 2003.
- [50] H. Schive, Y. Tsai, and T. Chiueh. Gamer: A graphic processing unit accelerated adaptive-mesh-refinement code for astrophysics. *The Astrophysical Journal Supplement Series*, 186:457, 2010.
- [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] A. H. E. Zein and A. P. Rendell. From sparse matrix to optimal gpu cuda sparse matrix vector product implementation. In *CCGRID*, pages 808–813. IEEE, 2010.
- [54] A. H. E. Zein and A. P. Rendell. Generating optimal CUDA sparse matrix-vector product implementations for evolving GPU hardware. *Concurrency and Computation: Practice and Experience*, 24(1):3–13, 2012.