

H2020-FETHPC-2014: GA 671633

D4.3

Prototype software, phase 2

October 2017

## DOCUMENT INFORMATION

Scheduled delivery 2017-10-31  
Actual delivery 2017-10-31  
Version 1.0  
Responsible partner INRIA

## DISSEMINATION LEVEL

PU — Public

## REVISION HISTORY

Date	Editor	Status	Ver.	Changes
2017-10-05	Inria members	Draft	0.1	Initial version of document produced
2017-10-27	Inria members	Version	1.0	Revision based on partner comments

## AUTHOR(S)

Laura Grigori, Simplicie Donfack, Olivier Tissot, INRIA

## INTERNAL REVIEWERS

Carl Christian Kjelgaard Mikkelsen, UMU  
Florent Lopez, STFC

## COPYRIGHT

This work is ©by the NLA FET Consortium, 2015–2018. Its duplication is allowed only for personal, educational, or research uses.

## ACKNOWLEDGEMENTS

This project has received funding from the *European Union's Horizon 2020 research and innovation programme* under the grant agreement number 671633.

## Table of Contents

<b>1</b>	<b>Executive Summary</b>	<b>5</b>
<b>2</b>	<b>Introduction</b>	<b>5</b>
<b>3</b>	<b>Using preAlps library</b>	<b>6</b>
3.1	Overview . . . . .	6
3.2	Installation . . . . .	6
3.3	Input data formats . . . . .	8
<b>4</b>	<b>Enlarged Conjugate Gradient</b>	<b>9</b>
4.1	Introduction . . . . .	9
4.2	Routines . . . . .	11
4.3	Example program . . . . .	13
<b>5</b>	<b>LORASC</b>	<b>15</b>
5.1	LORASC preconditioner . . . . .	15
5.2	Routines . . . . .	16
5.3	Example program . . . . .	17
<b>6</b>	<b>Experiments</b>	<b>19</b>
6.1	Enlarged Conjugate Gradient . . . . .	19
6.1.1	Environment . . . . .	19
6.1.2	Description of test matrices . . . . .	20
6.1.3	Performance of ECG . . . . .	20
6.2	LORASC . . . . .	23
6.2.1	Environment . . . . .	23
6.2.2	Performance of LORASC . . . . .	23
<b>7</b>	<b>Routines</b>	<b>25</b>
<b>8</b>	<b>Conclusion</b>	<b>27</b>
<b>9</b>	<b>Acknowledgments</b>	<b>28</b>

## List of Figures

1	Runtimes for Ela400 with a tolerance of $10^{-5}$ on 48 processors. . . . .	21
2	Strong scaling for Ela400 with a tolerance of $10^{-5}$ . . . . .	22
3	Performance of LORASC on 3D linear elasticity problems . . . . .	24
4	Strong scaling performance of LORASC on elasticity 3D problem . . . . .	25

## List of Tables

1	Dependencies of preAlps with external libraries. . . . .	7
2	Test matrix . . . . .	20
3	Runtimes comparison between ECG(12) and PETSc PCG for Ela400 . . .	22
4	Runtime of Block Jacobi and LORASC for a 3D linear elasticity problem .	25

# 1 Executive Summary

This document describes a prototype version of `preAlps`, a library that implements communication avoiding solvers based on enlarged Krylov subspace methods and robust preconditioners. It focuses in particular on enlarged CG (ECG) and LORASC preconditioner described in deliverable D4.2.

This document explains how to install `preAlps` and how to use its routines to solve a sparse linear system of equations. The library can be downloaded from <https://github.com/NLAFET/preAlps>. ECG can be used as an iterative solver and can be combined with block Jacobi, LORASC, or any other efficient preconditioner. It is based on reverse communication so that it can be used also for matrix-free problems. LORASC can be used in combination with ECG but also with any other Krylov solver for symmetric positive definite matrices. Our experiments show that ECG and LORASC can be used efficiently for solving challenging problems.

# 2 Introduction

The *Description of Action* document states for Deliverable 4.3:

*“Prototype software, phase 2  
Prototypes for Krylov-based iterative methods, and multilevel preconditioners.”*

This deliverable is in the context of WP4, Task 4.2 (Iterative methods) and Task 4.3 (Preconditioners).

The focus of WP4 is on solving large sparse linear systems of equations by using preconditioned iterative methods based on Krylov subspace solvers. Each iteration of a typical Krylov subspace solver involves matrix-vector multiplications and several dot products in order to orthogonalize the Krylov subspace. These dot products require collective communication among all processors which does not scale to a very large number of processors. This is the main bottleneck for the scalability of Krylov subspace methods. To increase the scalability of Krylov subspace solvers, our research in NLAFET focuses on enlarged Krylov subspace methods [8] and a robust and algebraic preconditioner LORASC [6].

Enlarged Krylov subspace methods [8] represent a new approach that consists of enlarging the Krylov subspace by a maximum of  $t$  vectors per iteration, based on a domain decomposition of the graph of the input matrix. The solution of the linear system is sought in the enlarged subspace, which is a superset of the classical subspace. The enlarged Krylov projection subspace methods lead to faster convergence in terms of number iterations and parallelizable algorithms with less communication, compared to classical Krylov methods. In addition, each iteration of such a method relies on multiplying a matrix with multiple vectors at once, and this typically allows for better resource utilization.

LORASC [6] is a robust algebraic preconditioner that can be built and applied in parallel. The graph of the matrix is partitioned by using nested dissection partitioning into  $N$  disjoint domains and a separator formed by the vertices connecting the  $N$  domains. The permuted matrix has an arrow head structure with  $N+1$  blocks on the main diagonal. The preconditioner relies on the Cholesky factorization of the first  $N$  diagonal blocks

and on approximating the Schur complement corresponding to the separator block. The approximation of the Schur complement involves the factorization of the last diagonal block and a low rank correction obtained by solving a generalized eigenvalue problem.

This document describes a prototype of the `preAlps` library developed as part of NLAFET. This library aims at solving large scale sparse linear systems on massively parallel architectures. In its current state the library contains routines for solving symmetric positive definite (SPD) linear systems with enlarged Conjugate Gradient (ECG) in parallel. It also provides routines for using LORASC as a preconditioner, which can be used in combination with ECG or with another Krylov method for SPD matrices as for example provided in PETSC[4]. We also provide preliminary results that show the parallel performance of the library.

## 3 Using `preAlps` library

### 3.1 Overview

`preAlps` is a parallel library that implements several preconditioners and iterative solvers based on communication avoiding principle. It focuses on two main aspects: the reformulation of Krylov-based iterative methods to allow a drastic reduction in the number of global communications with respect to classical formulations, and the design of robust preconditioners to accelerate the convergence of iterative methods. The current first prototype of the library provides an efficient implementation of ECG and LORASC preconditioner. ECG is based on enriching the Krylov subspace used in classical methods that allows to reduce drastically the communication cost of the iterative solver. LORASC is based on a low rank approximation of the Schur complement. These algorithms are well suited for parallelism and are described in details in [8] and [6]. In our experiments, by increasing the numbers of vectors added to the Krylov subspace by 12 instead of 1 as in the traditional conjugate gradient (CG) approach, ECG is up to 3x faster than the equivalent solver routine in PETSC, while LORASC is up to 7x faster than Block Jacobi preconditioner.

### 3.2 Installation

`preAlps` is developed at INRIA in the context of the NLAFET project, and it can be downloaded from [https://github.com/NLAFET/preAlps](https://github.com/NLA<span>F</span>ET/preAlps). The library depends on few external libraries that need to be installed and linked with `preAlps`:

- BLAS and LAPACK [3]: BLAS is a standard library for performing basic vector and matrix operations. LAPACK is a standard software for numerical linear algebra. Although any library providing BLAS and LAPACK can be used, we recommend MKL [17].
- METIS[11] and ParMETIS [12]: the sequential and parallel graph partitioning tools. METIS is required in order to use ECG, while ParMETIS is required in order to use LORASC. We recommend to install ParMETIS as it already contains all METIS routines. `preAlps` were tested with METIS 5.1.0 and ParMETIS

	BLAS	METIS	ParMETIS	CPaLAMeM	PARPACK	MUMPS	PARDISO
ECG Solver	×	×		×			MKL
LORASC preconditioner	×		×	×	×	×	ANY
Full installation	×	×	×	×	×	×	MKL

Table 1: Dependencies of preAlps with external librairies.

4.0.3. These partitioning tools can be downloaded from <http://glaros.dtc.umn.edu/gkhome/metis/metis/overview> and <http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview>.

- CPaLAMeM (C Parallel Linear Algebra Memory Management): a library developed at INRIA (Alpines group) and part in the context of NLAFFET. It provides several highly efficient routines for manipulating sparse and dense matrices. It also provides several tools for memory management and debugging purpose. CPaLAMeM library is distributed with `preAlps` and will be installed with the library.
- PARPACK [13]: a parallel library used to solve eigenvalue problems. PARPACK is required in order to use LORASC. A latest version can be downloaded from <http://www.caam.rice.edu/software/ARPACK/>. At the moment, only PARPACK is supported in `preAlps`, but we plan to use other eigenvalue solvers as the one developed in the NLAFFET project (see Deliverable 2.5 for dense matrices).
- MUMPS [2]: a distributed parallel sparse direct solver. `preAlps` were tested with MUMPS 5.1.2. It can be downloaded from <http://mumps.enseeiht.fr/>
- PARDISO: a sequential and multithreaded sparse direct solver. This library is optional if MKL is already provided. If MKL is not provided, PARDISO from <http://pardiso-project.org/> should be installed.

In addition, the example program `test_ecgsolve.c` allows the end user to call CG and Block Jacobi from PETSC and compare its performance with ECG from `preAlps`. PETSC [4] provides a suite of routines for scientific applications including solvers and preconditioners.

The complete installation of `preAlps` could be summarized as follows:

1. Install all the dependencies libraries.

Table 1 shows the list of dependencies of the main components of `preAlps` with external librairies. In order to use ECG Solver, BLAS, METIS and MKL (with PARDISO) are required. In order to use LORASC, BLAS, ParMETIS, CPaLAMeM, PARPACK and MUMPS are required. For the full installation of `preAlps`, BLAS, ParMETIS, PARPACK, MUMPS and MKL (with PARDISO) are required. In all the cases, CPaLAMeM library is provided with `preAlps`.



2. At the top of the `preAlps` folder, edit the `make.inc` file in order to set the corresponding path for the external libraries.
3. Next, type `make install_cpalamem` to install CPaLAMeM.  
We note that the git repository of `preAlps` contains the latest stable version of CPaLAMeM. If a new version of `preAlps` is available, it will be necessary to type `make remove_cpalamem` before `make install_cpalamem` in order to reinstall an updated version of CPaLAMeM.
4. Then, type `make` to compile `preAlps` and create the library `lib/libpreAlps.a` and the example programs.
5. Now, the functions from the lib `preAlps` can be called by a program by including their corresponding header file. The folder `example` provide few stand-alone programs to test the library.

### 3.3 Input data formats

Most routines in `preAlps` require matrices stored into a compressed sparse row format (CSR). CPaLAMeM provides a routine to read such matrices as `CPLM_MatCSR_t` objects. The `CPLM_MatCSR_t` structure are presented as follows:

- `m, n` : the size of the matrix.
- `rowPtr`: the beginning index of each rows as described in the CSR format.
- `colInd, val`: the column indices and the corresponding values of each non-zeros elements of the matrix.

The following block of code presents an example for reading an external matrix.

```
1  /* Including headers */
2  #include <CPLM_mat_csr.h>
3
4  int main(int argc, char **argv){
5
6      /* The matrix file to load */
7      char matrixFileName []="cage4.mtx";
8
9      /* Create an empty CSR Matrix */
10     CPLM_Mat_CSR_t A = CPLM_MatCSRNULL();
11
12     /* Load the matrix file */
13     CPLM_LoadMatrixMarket(matrixFileName, &A);
14     ...
```

## 4 Enlarged Conjugate Gradient

### 4.1 Introduction

We briefly recall the definition of Enlarged Conjugate Gradient (ECG) [8] for solving the linear system  $Ax = b$ . The method has been previously described in *Deliverable 4.2*.

We use the following notations:  $U^\top$  is the transpose of a matrix  $U$ ,  $A$  is a symmetric ( $A^\top = A$ ) positive definite ( $x^\top Ax > 0, \forall x \neq 0$ ) real matrix of size  $n \times n$ ,  $B$  is a real matrix of size  $n \times t$  where  $t$  is the number of vectors used in Enlarged CG,  $B^{(i)}$  is the  $i$ -th column of a matrix  $B$ ,  $X_0$  is an initial guess for the linear system  $AX = B$ , *i.e.* it is a real matrix of size  $n \times t$ . We denote the initial residual matrix  $R_0 = B - AX_0$ . We call  $t$  the initial block size. Unless otherwise stated,  $\|\cdot\|$  denotes the usual euclidean norm both for vectors and matrices.

Given a splitting decomposition represented by the operator  $T$ ,

$$T(x) = \begin{pmatrix} * \\ \vdots \\ * \\ & * \\ & \vdots \\ & * \\ & & \ddots \\ & & & * \\ & & & \vdots \\ & & & * \\ & & & & * \\ & & & & \vdots \\ & & & & * \end{pmatrix}$$

where  $x$  is a vector of size  $n$  and  $T(x)$  is of size  $n \times t$ . The first columns of  $T(x)$  contains the first components of  $x$  and so on for the following columns of  $T(x)$ . The initial residual denoted  $r_0$ , the corresponding enlarged Krylov subspace is defined as

$$\mathcal{K}_{k,t} = \text{span}^\square\{T(r_0), AT(r_0), \dots, A^{k-1}T(r_0)\}. \tag{4.1}$$

Using this definition, *Grigori, Moufawad and Nataf* [8] derive a *Short Recurrence Enlarged CG*. The stopping criterion is the EUCLIDEAN norm of  $r_k = \sum_i R_k^{(i)}$  and the approximate solution is  $x_k = \sum_i X_k^{(i)}$ . In [7], we use the framework of Block Conjugate Gradient [16] to derive two versions of ECG: Orthodir (Algorithm 1) which corresponds to the SRE-CG in [8] and Orthomin (Algorithm 2) which corresponds to the Block Conjugate Gradient of O’Leary.

---

**Algorithm 1** Preconditioned Orthodir Enlarged CG

---

**Require:**  $A, M, b, x_0, k_{\max}, \varepsilon_{\text{solver}}$ **Ensure:**  $\|b - Ax_k\| < \varepsilon_{\text{solver}}$  or  $k = k_{\max}$ 

```

1:  $R_0 = T(b - Ax_0)$ 
2:  $r_0 = b - Ax_0$ 
3:  $P_0 = 0$ 
4:  $P_1 = M^{-1}R_0$ 
5:  $k = 1$ 
6: while  $\|r_{k-1}\| > \varepsilon_{\text{solver}}\|b\|$  and  $k < k_{\max}$  do
7:    $Q_k = AP_k$ 
8:   A-orthonormalize( $P_k, Q_k$ )
9:    $\alpha_k = P_k^\top R_{k-1}$ 
10:   $X_k = X_{k-1} + P_k \alpha_k$ 
11:   $R_k = R_{k-1} - Q_k \alpha_k$ 
12:   $Z_k = M^{-1}Q_k$ 
13:   $P_{k+1} = Z_k - P_k Q_k^\top Z_k - P_{k-1} Q_{k-1}^\top Z_k$ 
14:   $r_k = \sum_{i=1}^t R_k^{(i)}$ 
15:   $k = k + 1$ 
16: end while
17:  $x_k = \sum_{i=1}^t X_k^{(i)}$ 

```

---



---

**Algorithm 2** Preconditioned Orthomin Enlarged CG

---

**Require:**  $A, M, b, x_0, k_{\max}, \varepsilon_{\text{solver}}$ **Ensure:**  $\|b - Ax_k\| < \varepsilon_{\text{solver}}$  or  $k = k_{\max}$ 

```

1:  $R_0 = B - AX_0$ 
2:  $r_0 = b - Ax_0$ 
3:  $P_1 = M^{-1}R_0$ 
4:  $k = 1$ 
5: while  $\|r_{k-1}\| > \varepsilon_{\text{solver}}\|b\|$  and  $k < k_{\max}$  do
6:    $Q_k = AP_k$ 
7:   A-orthonormalize( $P_k, Q_k$ )
8:    $\alpha_k = P_k^\top R_{k-1}$ 
9:    $X_k = X_{k-1} + P_k \alpha_k$ 
10:   $R_k = R_{k-1} - Q_k \alpha_k$ 
11:   $Z_k = M^{-1}R_k$ 
12:   $P_{k+1} = Z_k - P_k Q_k^\top Z_k$ 
13:   $r_k = \sum_{i=1}^t R_k^{(i)}$ 
14:   $k = k + 1$ 
15: end while
16:  $x_k = \sum_{i=1}^t X_k^{(i)}$ 

```

---

## 4.2 Routines

Our implementation of ECG is based on Reverse Communication Interface [10] and written in C and MPI. Following this scheme we provide 4 routines:

- `preAlps_ECGInitialize(ECG_t* ecg, double* rhs, int* rci_request)`,
- `preAlps_ECGIterate(ECG_t* ecg, int* rci_request)`,
- `preAlps_ECGStoppingCriterion(ECG_t* ecg, int* stop)`
- `preAlps_ECGFinalize(ECG_t* ecg, double* solution)`.

In order to ease its usage, we encapsulate all the required information by ECG in the structure `preAlps_ECG_t`. This structure is defined as,

```

1  typedef struct {
2      /* Array type variables */
3      double*      b;          /* Right hand side */
4      CPLM_Mat_Dense_t* X;     /* Approximated solution */
5      CPLM_Mat_Dense_t* R;     /* Residual */
6      CPLM_Mat_Dense_t* P;     /* Descent direction */
7      CPLM_Mat_Dense_t* AP;    /* A*P */
8      CPLM_Mat_Dense_t* P_prev; /* Previous descent direction */
9      CPLM_Mat_Dense_t* AP_prev; /* A*P_prev */
10     CPLM_Mat_Dense_t* alpha;  /* Descent step */
11     CPLM_Mat_Dense_t* beta;   /* Step to construt search directions */
12     CPLM_Mat_Dense_t* gamma;  /* Step to construct odir search directions */
13     CPLM_Mat_Dense_t* Z;      /* Extra memory */
14     CPLM_Mat_Dense_t* H;      /* Descent directions needed to reduce block size */
15     CPLM_Mat_Dense_t* AH;     /* A*H */
16     double*      work;        /* working array */
17     int*         iwork;       /* working array */
18
19     /* Single value variables */
20     double      normb;        /* norm_2(b) */
21     double      res;         /* norm_2 of the residual */
22     int         iter;        /* Iteration */
23
24     /* Options and parameters */
25     int         globPbSize;   /* Size of the global problem */
26     int         locPbSize;   /* Size of the local problem */
27     int         maxIter;     /* Maximum number of iterations */
28     int         enlFac;      /* Enlarging factor */
29     double      tol;         /* Tolerance */
30     ECG_Ortho_Algo_t ortho_alg; /* A-orthonormalization algorithm */
31     ECG_Block_Size_Red_t bs_red; /* Block size reduction */
32     MPI_Comm    comm;        /* MPI communicator */
33 } ECG_t;

```

where `Ortho_Algo_t` and `Block_Size_Red_t` are enum. The structure `CPLM_Mat_Dense_t` represents local dense matrices, it contains the pointer to the data (`val`), the number of rows (`info.m`), the number of columns (`info.n`) and an enum to specify the matrix storage

type (ROW\_MAJOR or COL\_MAJOR). As it is now, the reduction of the search directions during the iterations (see *Deliverable 4.2*) is not implemented.

First, the user has to declare a variable of type `ECG_t` and set values for the parameters: `comm`, `globPbSize`, `locPbSize`, `maxIter`, `enlFac`, `tol`, `ortho_alg` (`bs_red` is not used for the moment). The variable `globPbSize` corresponds to the number of unknowns of the global solution  $x$  (the dimension of  $A$ ) and `locPbSize` corresponds to the number of unknown owned locally (the number of rows of  $A$  that are stored locally).

Then, in order to allocate memory and initialize the structure, the user has to call `preAlps_ECGInitialize(&ecg, rhs, &rci_request)` where `rhs` is an array (`double *`) representing the right hand side and `rci_request` is an integer. After this call, he has to apply the preconditioner to `ecg.R` and put the result into `ecg.P`. And then, he has to apply the operator  $A$  to `ecg.P` and put the result into `ecg.AP`. These two operations has to be executed in parallel assuming that `ecg.R`, `ecg.P`, `ecg.AP` contains local rows of  $R$ ,  $P$  and  $AP$  which are distributed as row panel over the processors.

Afterwards, the user has to call `preAlps_ECGIterate(&ecg, &rci_request)` until convergence of the method. Following RCI scheme, after this call the user has to check the value of `rci_request`. If `rci_request = 0`, then the user is requested to apply  $A$  on `ecg.P` and put the result into `ecg.AP`. If `rci_request = 1`, then the user can check for convergence of the method. If the method did not converge, then depending of the choice of the orthogonalization algorithm (ORTHODIR or ORTHOMIN) he has to apply the preconditioner to `ecg.R` (ORTHOMIN) or `ecg.AP` (ORTHODIR) and put the result into `ecg.Z`.

When convergence is reached, it is possible to recover the solution and free the structure `ecg` by calling `preAlps_ECGFinalize(&ecg, sol)` where `sol` is an array (`double*`) already allocated.

To sum up, for solving a linear system with a block Jacobi preconditioner, the general calling sequence will be

```

1 // Set parameters
2 ecg.comm = MPI_COMM_WORLD; /* MPI Communicator */
3 ecg.globPbSize = M;        /* Size of the global problem */
4 ecg.locPbSize = m;         /* Size of the local problem */
5 ecg.maxIter = maxIter;    /* Maximum number of iterations */
6 ecg.enlFac = 2;           /* Enlarging factor */
7 ecg.tol = tol;            /* Tolerance of the method */
8 ecg.ortho_alg = ORTHODIR; /* Orthogonalization algorithm */
9 // Allocate memory and initialize variables
10 preAlps_ECGInitialize(&ecg, rhs, &rci_request);
11 // Finish initialization
12 preAlps_BlockJacobiApply(ecg.R, ecg.P);
13 preAlps_BlockOperator(ecg.P, ecg.AP);
14 // Main loop
15 while (stop != 1) {
16     ierr = preAlps_ECGIterate(&ecg, &rci_request);
17     if (rci_request == 0) {
18         // AP = A*P
19         preAlps_BlockOperator(ecg.P, ecg.AP);
20     }
21     else if (rci_request == 1) {
22         ierr = preAlps_ECGStoppingCriterion(&ecg, &stop);
23         if (stop == 1) break;

```

```

24     if (ecg.ortho_alg == ORTHOMIN)
25         // Z = M^-1*R
26         preAlps_BlockJacobiApply(ecg.R,ecg.Z);
27     else if (ecg.ortho_alg == ORTHODIR)
28         // Z = M^-1*AP
29         preAlps_BlockJacobiApply(ecg.AP,ecg.Z);
30     }
31 }
32 // Retrieve solution and free memory
33 preAlps_ECGFinalize(&ecg,sol);

```

### 4.3 Example program

The following listing is an abbreviated version of usage of Enlarged Conjugate Gradient routines for solving a parallel linear system. We provide in this example routines for calling block Jacobi as a preconditioner. However the user can combine ECG with any preconditioner. The routines for initializing and applying a preconditioner and associated data structures can evolve in our library. The user can also use his own routine for computing the matrix set of vectors product instead of using the `preAlps_BlockOperator` routine.

```

1  /* STD */
2  #include <stdlib.h>
3  #include <stdio.h>
4  #include <stddef.h>
5  /* MPI */
6  #include <mpi.h>
7  /* MKL */
8  #include <mkl.h>
9  /* CPaLAMeM */
10 #include <cpalamem_macro.h>
11 #include <cpalamem_instrumentation.h>
12 /* preAlps */
13 #include "operator.h"
14 #include "ecg.h"
15
16 /*****
17 /*                               CODE                               */
18 *****/
19 int main(int argc, char** argv) {
20
21     /*===== Initialize =====*/
22     MPI_Init(&argc, &argv);
23     int rank, size, ierr;
24     MPI_Comm_size(MPI_COMM_WORLD, &size);
25     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
26
27     /*===== Construct the operator using a CSR matrix =====*/
28     const char* matrixFilename = argv[1];
29     CPLM_Mat_CSR_t A = MatCSRNULL();
30     int M, m;
31     int* rowPos = NULL;

```

```

32  int* colPos = NULL;
33  int* dep = NULL;
34  int sizeRowPos, sizeColPos, sizeDep;
35  // Read and partition the matrix
36  preAlps_OperatorBuild(matrixFilename,MPI_COMM_WORLD);
37  // Get the CSR structure of A
38  preAlps_OperatorGetA(&A);
39  // Get the sizes of A
40  preAlps_OperatorGetSizes(&M,&m);
41  // Get row partitioning of A
42  preAlps_OperatorGetRowPosPtr(&rowPos,&sizeRowPos);
43  // Get col partitioning induced by this row partitioning
44  preAlps_OperatorGetColPosPtr(&colPos,&sizeColPos);
45
46  /*===== Construct the preconditioner =====*/
47  preAlps_BlockJacobiCreate(precond_type,
48                          &A,
49                          rowPos,
50                          sizeRowPos,
51                          colPos,
52                          sizeColPos);
53
54  /*===== Construct a random rhs =====*/
55  double* rhs = (double*) malloc(A.info.m*sizeof(double));
56  // Set the seed of the random generator
57  srand(0);
58  for (int i = 0; i < A.info.m; ++i)
59      rhs[i] = ((double) rand() / (double) RAND_MAX);
60
61  /*===== ECG solve =====*/
62  ECG_t ecg;
63  // Set parameters
64  ecg.comm = MPI_COMM_WORLD; /* MPI Communicator */
65  ecg.globPbSize = M;      /* Size of the global problem */
66  ecg.locPbSize = m;      /* Size of the local problem */
67  ecg.maxIter = 1000;    /* Maximum number of iterations */
68  ecg.enlFac = 2;        /* Enlarging factor */
69  ecg.tol = 1e-5;        /* Tolerance of the method */
70  ecg.ortho_alg = ORTHODIR; /* Orthogonalization algorithm */
71  // Get local and global sizes of operator A
72  int rci_request = 0;
73  int stop = 0;
74  double* sol = NULL;
75  sol = (double*) malloc(m*sizeof(double));
76  // Allocate memory and initialize variables
77  preAlps_ECGInitialize(&ecg,rhs,&rci_request);
78  // Finish initialization
79  preAlps_BlockJacobiApply(precond_type,ecg.R,ecg.P);
80  preAlp_BlockOperator(ecg.P,ecg.AP);
81  // Main loop
82  while (stop != 1) {
83      ierr = preAlps_ECGIterate(&ecg,&rci_request);
84      if (rci_request == 0) {
85          preAlps_BlockOperator(ecg.P,ecg.AP);
86      }

```

```

87     else if (rci_request == 1) {
88         ierr = preAlps_ECGStoppingCriterion(&ecg,&stop);
89         if (stop == 1) break;
90         if (ecg.ortho_alg == ORTHOMIN)
91             preAlps_BlockJacobiApply(ecg.R,ecg.Z);
92         else if (ecg.ortho_alg == ORTHODIR)
93             preAlps_BlockJacobiApply(ecg.AP,ecg.Z);
94     }
95 }
96 // Retrieve solution and free memory
97 preAlps_ECGFinalize(&ecg,sol);
98
99 if (rank == 0)
100     printf("=== ECG ===\n\titerations: %d\n\tnorm(res): %e\n",ecg.iter,ecg.res);
101
102 /*===== Finalize =====*/
103 // Free arrays
104 if (rhs != NULL) free(rhs);
105 if (sol != NULL) free(sol);
106 preAlps_OperatorFree();
107 MPI_Finalize();
108 return ierr;
109 }

```

## 5 LORASC

### 5.1 LORASC preconditioner

In this section, we briefly recall LORASC, a robust algebraic preconditioner of the form  $M = (L + \tilde{D})\tilde{D}^{-1}(\tilde{D} + L^T)$  as presented in the *Deliverable 4.2*. The graph of the input matrix is first partitioned by using  $k$ -way partitioning with vertex separators into  $N$  disjoint domains and a separator  $\Gamma$  formed by the interface vertices connecting the  $N$  domains. Such a partitioning can be obtained by using existing software as METIS [11]. The permuted matrix has a block arrow structure, as presented in equation (5.3), in which the first  $N$  diagonal blocks correspond to the disjoint domains, while the last diagonal block  $A_{\Gamma\Gamma}$  corresponds to the separator.

Consider a symmetric positive definite matrix  $A$  of size  $n \times n$ , which has a bordered block diagonal structure as in Equation (5.3). We refer in the following to LORASC preconditioner as  $M$ . The preconditioner  $M$  is defined by the following approximate factorization

$$\begin{aligned}
 M &= (L + \tilde{D})\tilde{D}^{-1}(\tilde{D} + L^T) \\
 &= \begin{pmatrix} A_{11} & & & \\ & \ddots & & \\ & & A_{NN} & \\ A_{\Gamma 1} & \cdots & A_{\Gamma N} & \tilde{S} \end{pmatrix} \begin{pmatrix} A_{11}^{-1} & & & \\ & \ddots & & \\ & & A_{NN}^{-1} & \\ & & & \tilde{S}^{-1} \end{pmatrix} \begin{pmatrix} A_{11} & & & A_{\Gamma^T} \\ & \ddots & & \vdots \\ & & A_{NN} & A_{\Gamma N}^T \\ & & & \tilde{S} \end{pmatrix},
 \end{aligned} \tag{5.1}$$



where  $\tilde{D} = \text{Block-Diag}(A_{11}, A_{22}, \dots, A_{NN}, \tilde{S})$ ,  $L$  is defined as

$$L = \begin{pmatrix} 0 & & & \\ & \ddots & & \\ & & 0 & \\ A_{\Gamma 1} & \cdots & A_{\Gamma N} & 0 \end{pmatrix}, \quad (5.2)$$

and  $\tilde{S}$  is an approximation of the Schur complement  $S = A_{\Gamma\Gamma} - \sum_{j=1}^N A_{\Gamma j} A_{jj}^{-1}$ .

**Definition 5.1** (LORASC preconditioner). *Let  $A$  be an  $n \times n$  symmetric positive definite matrix with a bordered block diagonal structure,*

$$A = \begin{pmatrix} A_{11} & & & A_{\Gamma 1}^T \\ & \ddots & & \vdots \\ & & A_{NN} & A_{\Gamma N}^T \\ A_{\Gamma 1} & \cdots & A_{\Gamma N} & A_{\Gamma\Gamma} \end{pmatrix}. \quad (5.3)$$

Let  $S = A_{\Gamma\Gamma} - \sum_{j=1}^N A_{\Gamma j} A_{jj}^{-1} A_{\Gamma j}^T$ . Given a tolerance  $\tau$ , a lower bound  $\varepsilon = \frac{1}{\tau}$  for the generalized eigenvalue problem  $Su = \lambda A_{\Gamma\Gamma} u$ , let  $\lambda_1, \lambda_2, \dots, \lambda_i$  be the generalized eigenvalues less than  $\varepsilon$ , i.e. for all  $k \in \{1, \dots, i\}$  then  $\lambda_k < \varepsilon$ , and let  $v_1, v_2, \dots, v_i$  be the corresponding  $A_{\Gamma\Gamma}$ -orthonormal generalized eigenvectors.

The LORASC preconditioner of  $A$  is defined as

$$M_{\text{LORASC}} := (L + \tilde{D})\tilde{D}^{-1}(\tilde{D} + L^T).$$

where  $\tilde{D} = \text{Block-Diag}(A_{11}, A_{22}, \dots, A_{NN}, \tilde{S})$  and  $L$  is given by (5.2). The matrix  $\tilde{S}$  is defined as

$$\tilde{S}^{-1} = A_{\Gamma\Gamma}^{-1} + E\Sigma E^T, \quad (5.4)$$

where  $E, \Sigma$  are defined as

$$E = (v_1 \ v_2 \ \dots \ v_i), \quad (5.5)$$

$$\Sigma = \text{Diag}(\sigma_1, \sigma_2, \dots, \sigma_i), \quad \text{with } \sigma_k = \frac{\varepsilon - \lambda_k}{\lambda_k}, \ k \in \{1, 2, \dots, i\}. \quad (5.6)$$

The parallel construction of LORASC using  $N$  processors is summarized in Algorithm 3.

## 5.2 Routines

LORASC preconditioner can be built separately and used in any sparse iterative solvers. The implemented routines are described as follows:

- `preAlps_LorascAlloc (preAlps_Lorasc_t **lorasc)` : creates an object of the type `preAlps_Lorasc_t`. The resulting object can be used by the end-user to replace the default parameters such as the `deflation_tolerance` ( $\varepsilon$  in definition 5.1).

---

**Algorithm 3** Parallel construction of LORASC using  $N$  processors

---

**Require:** SPD matrix  $A$ ,  $N$  processors, subset of  $N_\Gamma$  processors, tolerance  $\varepsilon$

- 1: Partition the graph of  $A$  into  $N$  disjoint domains using k-way partitioning with vertex separators, and permute the matrix into a block arrow structure as

$$A = \begin{pmatrix} A_{11} & & & A_{\Gamma 1}^T \\ & \ddots & & \vdots \\ & & A_{NN} & A_{\Gamma N}^T \\ A_{\Gamma 1} & \cdots & A_{\Gamma N} & A_{\Gamma\Gamma} \end{pmatrix}. \quad (5.7)$$

- 2: Distribute  $A$  on  $N$  processors such that processor  $j$  holds  $A_{\Gamma j}$  and  $A_{jj}$ , distribute  $A_{\Gamma\Gamma}$  on  $N_\Gamma$  processors by using a block row distribution
- 3: Processors  $j = 1, \dots, N$  compute locally the Cholesky factorization of  $A_{jj}$
- 4: A subset of  $N_\Gamma$  processors compute the Cholesky factorization of  $A_{\Gamma\Gamma}$  in parallel
- 5: All processors solve the generalized eigenvalue problem  $Su = \lambda A_{\Gamma\Gamma}u$  and compute the eigenvalues smaller than  $\varepsilon$  and associated eigenvectors using PARPACK, where  $S = A_{\Gamma\Gamma} - \sum_{j=1}^N A_{\Gamma j} A_{jj}^{-1} A_{\Gamma j}^T$  (the result is returned on processor 0, the shifted eigenvalues are stored in  $\Sigma$  and the eigenvectors in  $E$ )

**Ensure:**  $M = (L + \tilde{D})\tilde{D}^{-1}(\tilde{D} + L^T)$  where  $\tilde{D} = \text{Block-Diag}(A_{11}, A_{22}, \dots, A_{NN}, \tilde{S})$  and  $\tilde{S}^{-1} = A_{\Gamma\Gamma}^{-1} + E\Sigma E^T$

---

- `preAlps_LorascBuild` (`preAlps_Lorasc_t *lorasc`, `CPLM_Mat_CSR_t *A`, `CPLM_Mat_CSR_t *locAP`, `MPI_Comm comm`) : constructs LORASC preconditioner from an input matrix  $A$  and stores all the required internal workspace in the object `lorasc`. First, it partitions and permutes the matrix  $A$  into a block arrow structure, then distributes it to each processor. After this distribution, each processor stores in the output matrix `locAP` its block from a 1-D block row distribution of the permuted matrix  $A$ . Finally it constructs the preconditioner itself.
- `preAlps_LorascApply` (`preAlps_Lorasc_t *lorasc`, `double *x`, `double *y`) : applies LORASC preconditioner on a vector  $x$  and return the result in the vector  $y$ .
- `preAlps_LorascApplyMat` (`preAlps_Lorasc_t *lorasc`, `CPLM_Mat_Dense_t *A`, `CPLM_Mat_Dense_t *B`) : applies LORASC preconditioner on a dense matrix  $A$ , and returns the result in a dense matrix  $B$ . This routine does the same computation as `preAlps_LorascApply` routine with the difference that it applies the preconditioner on a dense matrix.
- `preAlps_LorascDestroy` (`preAlps_Lorasc_t **lorasc`) : frees the internal memory allocated by LORASC preconditioner and destroys `lorasc` object.

### 5.3 Example program

The following listing illustrates how LORASC can be used as a preconditioner in an example program for solving a sparse linear system. In this example, we use ECG as

iterative solver with an enlarging factor of 1, which is similar to using the classical CG algorithm.

```

1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <string.h>
4 #include <mpi.h>
5 #include <CPLM_mat_load_mm.h>
6 #include <CPLM_mat_csr.h>
7 #include "preAlps_utils.h"
8 #include "preAlps_preconditioner.h"
9 #include "preAlps_lorasc.h"
10 #include "preAlps_ecg.h"
11
12 int main(int argc, char** argv){
13
14     int i, bsize, ierr, nbprocs, my_rank;
15     char matrix_filename[150]="", rhs_filename[150]="";
16     CPLM_Mat_CSR_t A = CPLM_MatCSRNULL(), locAP = CPLM_MatCSRNULL();
17     double *x = NULL, *b = NULL;
18
19     /* Generic preconditioner object */
20     preAlps_preconditioner_t *precond = NULL;
21
22     /* Lorasc preconditioner */
23     preAlps_Lorasc_t *lorascA = NULL;
24
25     /* Start MPI*/
26     MPI_Init(&argc, &argv);
27     MPI_Comm_size(MPI_COMM_WORLD, &nbprocs);
28     MPI_Comm_rank(MPI_COMM_WORLD, &my_rank);
29
30     /* Load the matrix on processor 0 */
31     if(my_rank==0){
32         /* Get user parameters */
33         for(i=1;i<argc-1;i+=2){
34             if (strcmp(argv[i],"-mat") == 0) strcpy(matrix_filename,argv[i+1]);
35             if (strcmp(argv[i],"-rhs") == 0) strcpy(rhs_filename,argv[i+1]);
36         }
37
38         /* Read the matrix in CSR format */
39         CPLM_LoadMatrixMarket(matrix_filename, &A);
40
41         /* Read the rhs */
42         CPLM_DVectorLoad(rhs_filename, &b, &bsize, 0);
43     }
44
45     /* Memory allocation for LORASC preconditioner */
46     ierr = preAlps_LorascAlloc(&lorascA); preAlps_checkError(ierr);
47
48     /* Set LORASC parameters */
49     lorascA->deflation_tolerance = 1e-2; //eigenvalues deflation tolerance
50
51     /* Build the preconditioner and distribute the matrix according to the Block arrow
       structure */

```

```

52  ierr = preAlps_LorascBuild(lorascA, &A, &locAP, MPI_COMM_WORLD);
      preAlps_checkError(ierr);
53
54  /* Create a generic preconditioner object compatible with EcgSolver*/
55  preAlps_PreconditionerCreate(&precond, PREALPS_LORASC, (void *) lorascA);
56
57  /* Solve the system using ECGSolve*/
58  ECG_t ecg;          /* Enlarge CG solver */
59  /* Set the solver parameters */
60  ecg.comm = MPI_COMM_WORLD; /* MPI Communicator */
61  ecg.globPbSize = A.info.m; /* Size of the global problem */
62  ecg.locPbSize = locAP.info.m; /* Size of the local problem */
63  ecg.maxIter = 10000;      /* Maximum number of iterations */
64  ecg.enlFac = 1;          /* Enlarging factor */
65  ecg.tol = 1e-8;         /* Tolerance of the method */
66  ecg.ortho_alg = ORTHOMIN; /* Orthogonalization algorithm */
67
68
69  /* Call Ecg Solve */
70  /* ... */
71  /* See listing in ECG section */
72  /* ... */
73
74  if (my_rank == 0)
75      printf("=== ECG ===\n\titerations: %d\n\tnorm(res): %e\n", ecg.iter, ecg.res);
76
77  /* Destroy Lorasc preconditioner */
78  ierr = preAlps_LorascDestroy(&lorascA); preAlps_checkError(ierr);
79
80  /* Destroy the generic preconditioner object*/
81  preAlps_PreconditionerDestroy(&precond);
82
83  /* Free memory*/
84  if(b!=NULL) free(b);
85  CPLM_MatCSRFree(&locAP);
86
87  if(my_rank==0){
88      CPLM_MatCSRFree(&A);
89  }
90
91  MPI_Finalize();
92  return EXIT_SUCCESS;
93 }

```

## 6 Experiments

### 6.1 Enlarged Conjugate Gradient

#### 6.1.1 Environment

Concerning ECG, the experiments were performed on a machine located at Umeå University as part of High Performance Computing Center North (HPC2N), called Keb-

nekaise. It is an heterogeneous machine formed by a mix of Intel Xeon E5-2690v4 with 2x14 cores (and E7-8860v4 for large memory computations), Nvidia K80 GPU and Intel Xeon Phi 7250 (Knight’s Landing) with 68 cores. In our experiments, we use the so-called compute nodes, which means formed by Intel Xeon E5-2690v4. For a detailed description of the machine, we refer to the online documentation (<https://www.hpc2n.umu.se/resources/hardware/kebneka>).

### 6.1.2 Description of test matrices

We compare the performance of ECG and PETSc PCG on a set of matrices that are also used in [6, 1, 15] where they are described in more details. In this document, we present the results for one matrix in this set which is detailed in Table 2 where we present its size and its number of nonzeros.

The Ela matrices arise from the linear elasticity problem with Dirichlet and Neumann boundary conditions defined as follows

$$\operatorname{div}(\sigma(u)) + f = 0 \quad \text{on } \Omega \quad (6.1)$$

$$u = 0 \quad \text{on } \partial\Omega_D \quad (6.2)$$

$$\sigma(u) \cdot n = 0 \quad \text{on } \partial\Omega_N \quad (6.3)$$

where  $\Omega$  is a unit square cube. The matrices Ela $N$  correspond to this equation discretized using a triangular mesh with  $N \times 10 \times 10$  points on the corresponding vertices and P1 finite elements scheme.  $\partial\Omega_D$  is the Dirichlet boundary,  $\partial\Omega_N$  is the Neumann boundary,  $f$  is some body force,  $u$  is the unknown displacement field.  $\sigma(\cdot)$  is the Cauchy stress tensor given by Hooke’s law: it can be expressed in terms of Young’s Modulus  $E$  and Poisson’s ratio  $\nu$ . For a more detailed description of the problem see [14] and [6]. We consider discontinuous  $E$  and  $\nu$ ,  $(E_1, \nu_1) = (2 \times 10^{11}, 0.25)$  and  $(E_2, \nu_2) = (10^7, 0.45)$ .

Table 2: Test matrix we use in our tests, its size, the number of nonzeros and the type of problem it is coming from.

	Size	Nonzeros	Problem
Ela400	145 563	4 907 997	3D Linear Elasticity P1 FE

In the experiments we use a block Jacobi preconditioner with as many blocks as processors. Before calling ECG, each processor factorizes the diagonal block of  $A$  corresponding to the local row panel that it owns. At each iteration of ECG, each processor performs a backward and forward solve locally in order to apply the preconditioner. Hence, the application does not need any communication.

### 6.1.3 Performance of ECG

We compile `preAlps` (and its dependencies) using Intel toolchain installed on the machine: `mpiicc` (based on `icc` version 17.0.1 20161005) and `MKL` version 2017.1.132.

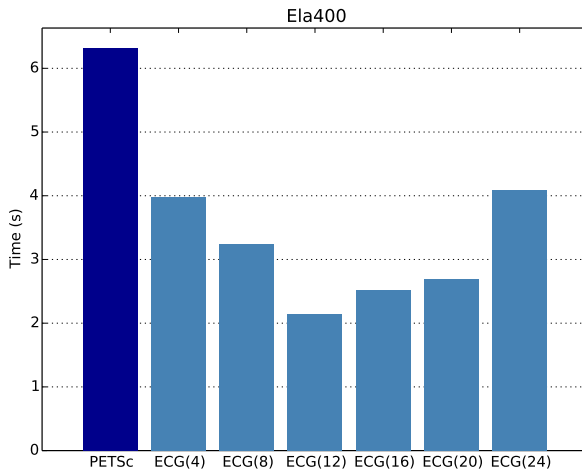


Figure 1: Runtimes for Ela400 with a tolerance of  $10^{-5}$  on 48 processors.

We are using PETSc in order to compare our ECG implementation to PETSc PCG implementation. PETSc 3.7.6 is compiled using the following configuration,

---

```
1 ./configure --with-cc=mpicc --with-cxx=0 --with-fc=0 COPTFLAGS="-O3 -march=native
  -mtune=native" --with-debugging=0 --with-blas-lapack-dir=${MKLRROOT}
  --with-mkl_pardiso-dir=${MKLRROOT} --download-metis
```

---

In particular, PETSc is using MKL-PARDISO as exact solver for sparse matrices in the block Jacobi preconditioner. CPaLAMeM has been linked with PETSc, and the METIS library downloaded and installed by PETSc.

When increasing the enlarging factor  $t$  the number of iterations decreases and the number of flop per iteration is increasing. In other words, if  $t$  is small the communication cost should dominate and if  $t$  is large the computational cost should dominate. There is no simple way to choose a priori the optimal enlarging factor  $t$ . In Figure 1 we plot the runtime as a function of  $t$  when solving Ela400 on 48 processors. Overall, ECG performs better than PETSc PCG, both with the same preconditioner. Indeed, ECG runtime is 2 to 3 times faster than PCG runtime. On the other hand, experimental results confirm that there exists a good choice of  $t$ . Indeed, the runtime is decreasing until  $t = 12$  but then it increases. It is important not to choose an enlarging factor too small, for example, ECG(4) is two times slower than ECG(12). On the other side, increasing  $t$  is not useful up to a certain number of right-hand sides, for example, ECG(12) is nearly two times faster than ECG(24).

In Figure 2 we summarize the results obtained when performing a strong scalability study on Ela400 with PETSc PCG and ECG(12). We start with 24 processors and double the number of processors until 192. First, ECG(12) scaling is sub-linear and PETSc PCG scales a little more than linearly. However, this test case is relatively small and thus the runtime with 192 processors is very small (around 1 second). When using 8 times more processors (from 24 to 192) PETSc PCG becomes around 8.9 times faster and ECG(12) around 6.5 times faster. However, ECG(12) remains around 2 times faster than PETSc PCG. With 192 processors ECG(12) is 86% faster and with 48 processors it is 3 times faster.

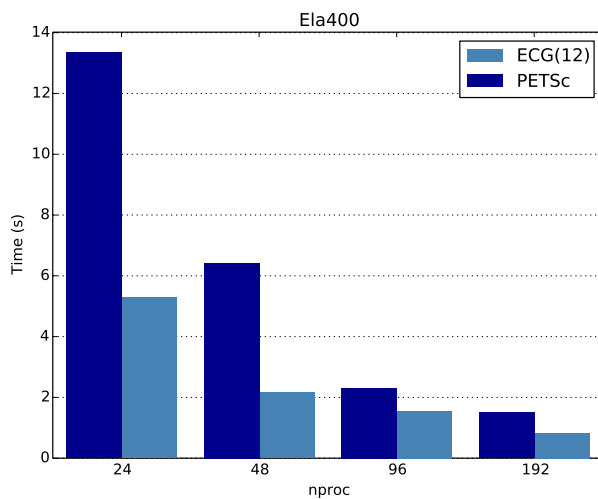


Figure 2: Strong scaling study for Ela400 with a tolerance of  $10^{-5}$  and an enlarging factor of 12.

Table 3: Runtimes comparison between ECG(12) and PETSc PCG for Ela400 with a tolerance of  $10^{-5}$  when increasing the number of processors.

	nprocs	runtimes (s)
ECG(12)	24	5.28
	48	2.16
	96	1.55
	192	0.83
PETSc PCG	24	13.35
	48	6.43
	96	2.31
	192	1.51

## 6.2 LORASC

### 6.2.1 Environment

In this section we discuss the performance of LORASC on a parallel machine formed by 28 compute nodes. This follows the description provided in [6]. Each node is equipped with a 24 cores socket based on Intel Xeon E5-2670 (Sandy Bridge), each core has a frequency of 2.6 GHz. We assign one MPI task per core. As described in [6], the construction of LORASC is performed in parallel and requires computing the Cholesky factorization of the diagonal blocks of  $A$  and solving a generalized eigenvalue problem. When  $N$  processors are used, first the input matrix  $A$  is permuted by using METIS into a matrix with a block arrow structure as in equation (5.3). The Cholesky factorization of each domain is computed by one of the  $N$  processors using PARDISO from MKL's scientific library. The Cholesky decomposition of  $A_{\Gamma\Gamma}$  is computed in parallel by a subset of the  $N$  processors using MUMPS 5.0.2 [2]. In our experiments, the number of processors used for  $A_{\Gamma\Gamma}$  is determined by the minimum of  $N$ , and  $size(A_{\Gamma\Gamma})/2000$ . Hence each processor will have at least 2000 rows of  $A_{\Gamma\Gamma}$ . We observed that increasing the number of processors used for the factorization of  $A_{\Gamma\Gamma}$  over this threshold leads to more communication between processors, and no performance gain. Once the factorization of the diagonal blocks is computed, the generalized eigenvalue problem is solved by using the parallel version of ARPACK, PARPACK 2.1 [13]. We provide to PARPACK a parallel routine that applies  $S$  to a vector and  $A_{\Gamma\Gamma}^{-1}$  to a vector.

Once LORASC is built, the linear system is solved using ECG solver with an enlarging factor of 1 which corresponds to the classical CG.

### 6.2.2 Performance of LORASC

We compare the performance of LORASC with block Jacobi from PETSc. On  $N$  processors, block Jacobi preconditioner uses the Cholesky factorization of  $N$  diagonal blocks of the matrix  $A$  permuted as following. The vertices corresponding to the separator  $\Gamma$  are partitioned into  $N$  sets. The rows and columns of the input matrix  $A$  are permuted such that the  $j$ -th set of vertices of the separator  $\Gamma$  are ordered after the vertices of the  $j$ -th domain. Processor  $j$  owns the rows corresponding to these vertices. The residual tolerance of the solver is fixed to  $10^{-6}$  for all the preconditioners. In the following tables, *No prec* refers to the version of CG without preconditioning.

Figure 3 presents the performance of LORASC on matrices arising from the discretization of the 3D linear elasticity problem. The dimension of the matrices increases from  $n = 4719$  to  $n = 72963$ . In the results presented in Figure 3, the number of processors, and hence the number of domains, is fixed to 16. For all the results, we include both the time to construct the preconditioner and the time to solve the system. The runtime of LORASC increases from 0.2 to 3 when the matrix size increases from 4719 to 72963. LORASC is significantly faster than Block Jacobi, up to a factor of 6 for the matrix of dimension  $n = 72963$ . In these results, all the eigenvalues smaller than  $\varepsilon = 10^{-2}$  are deflated in LORASC. For example, for  $n = 9438$ , 8 eigenvalues are deflated, and for  $n = 72963$ , 23 eigenvalues are deflated.

Figure 4 presents the performance of LORASC and Block Jacobi for a medium size matrix with  $n = 145563$  and 4907997 nonzero elements. The number of processors increases



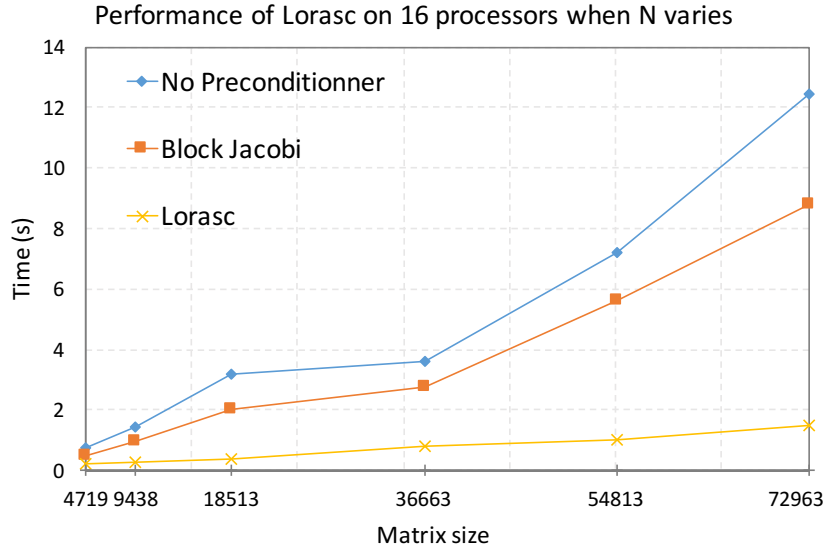


Figure 3: Performance of LORASC on 16 processors on 3D linear elasticity problems when  $n$  varies from 4719 to 72963.

from 16 to 256. The number of domains used in LORASC and the number of diagonal blocks used in Block Jacobi are equal to the number of processors. In addition, table 4 displays for both preconditioners the number of iterations (**Iter**), the time to construct the preconditioner (**Prec**), the time to solve (**Solve**), and the total time of the iterative process (**Total = Prec + Solve**). Similarly in Figure 4 the runtime is broken down into the time to construct the preconditioner (**Preconditioner**), and the solve time (**Solve**). One of the parameters that we need to pass to PARPACK is the number of eigenvalues and eigenvectors that need to be computed. In this experiment the number of eigenvalues computed by PARPACK is fixed to 150, and the eigenvalues smaller than the threshold  $\varepsilon = 10^{-2}$  are deflated. In the future, we will work on estimating the number of eigenvalues smaller than the threshold  $\varepsilon$ , and thus have an estimate of the number of eigenvalues that need to be computed by PARPACK and deflated by LORASC.

The results in Figure 4 show that LORASC is faster than Block Jacobi when up to 128 processors are used. It is 7 times faster than Block Jacobi on 16 processors. However, the time required to compute LORASC increases when the number of domains increases. More than 80% of the time is spent in solving the generalized eigenvalue problem with PARPACK. Since the size of the interface  $\Gamma$  becomes larger when the number of domains increases, the construction of LORASC requires solving a larger generalized eigenvalue problem, and this becomes more expensive, even if we use a larger number of processors to solve it. The time to build Block Jacobi is very small, but since the number of iterations to converge grows when the number of processors increases, the solve time is important. On 256 processors, even if Block Jacobi requires a factor of 64 more iterations to converge than LORASC, the time to apply Block Jacobi is almost equal with the time to construct LORASC.

These results show that for medium size problems and for a small to medium number

$N$	Block Jacobi				LORASC			
	Iter	Prec	Solve	Total	Iter	Prec	Solve	Total
16	4058	3.5E-4	132.1	132.1	102	15.4	2.7	18.1
32	6020	3.5E-4	91.8	91.8	131	12.6	1.4	14.1
64	9096	3.7E-4	64.3	64.3	184	13.3	1.0	14.3
128	13321	3.7E-4	41.3	41.3	227	22.4	1.2	23.6
256	15573	3.7E-4	37.3	37.3	243	36.3	1.8	38.1

Table 4: Parallel runtime of Block Jacobi and LORASC on  $N$  processors for a matrix arising from the discretization of a 3D linear elasticity problem. The matrix has dimension  $n = 145563$  and 4907997 nonzeros. **Iter** denotes the number of iterations until convergence of preconditioned CG, **Prec** denotes the time to construct the preconditioner and **Solve** denotes the solve time.

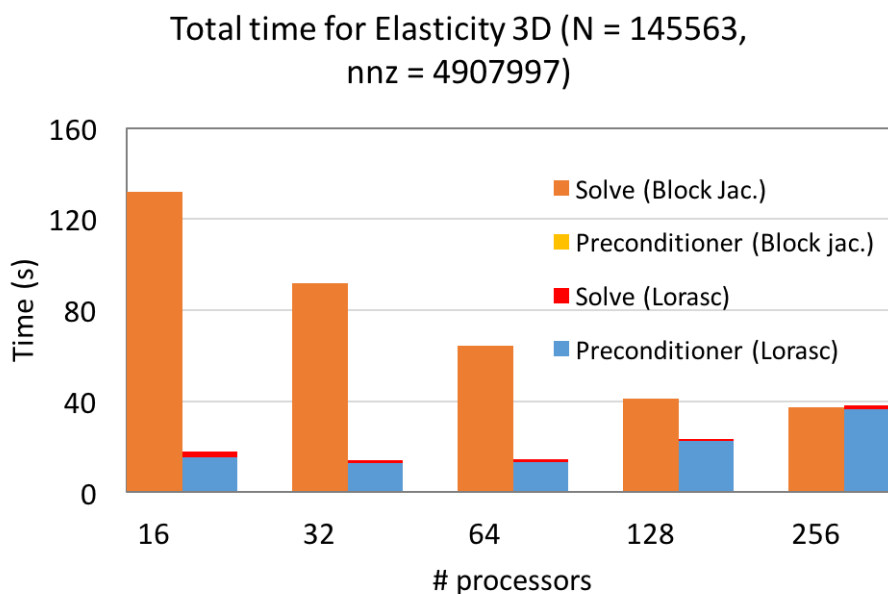


Figure 4: Strong scaling performance of LORASC on elasticity 3D problem of size  $n = 145563$  with 4907997 nonzeros.

of processors, LORASC is an efficient preconditioner. When the number of processors increases and the interface  $\Gamma$  becomes too large, the resolution of the generalized eigenvalue problem becomes expensive. However, if multiple linear systems with the same matrix  $A$  need to be solved, and the right hand sides are not available simultaneously, then LORASC remains an interesting option since the solve phase is very fast.

## 7 Routines

/\*

\* Read a sparse matrix from file (mtx or petsc binary) then partition it

```
* using METIS K-Way algorithm into the number of processor in the communicator
* comm and distribute it among those processors.
* input: matrixFilename: MatrixMarket or PETSc binary file
*       comm           : MPI communicator
*/
int preAlps_OperatorBuild(const char* matrixFilename, MPI_Comm comm);

/*
 * Release memory allocated during preAlp_OperatorBuild.
 */
void preAlps_OperatorFree();

/*
 * Print informations about the operator.
 */
void preAlps_OperatorPrint(int rank);

/*
 * Apply the operator to a group of vector X and put the result into AX.
 * input : X : CPLM_Mat_Dense_t
 * output: AX: CPLM_Mat_Dense_t
 */
int preAlps_OperatorGetSizes(int* M, int* m);

/*
 * Return the local part of the operator as a row panel CSR matrix.
 */
int preAlps_BlockOperator(CPLM_Mat_Dense_t* X, CPLM_Mat_Dense_t* AX);

/*
 * Return the global size and local size of the operator A.
 */
int preAlps_OperatorGetA(CPLM_Mat_CSR_t* A);

/*
 * Return a vector containing the global row index corresponding to the beginning
 * of each row panel.
 */
int preAlps_OperatorGetRowPosPtr(int** rowPos, int* sizeRowPos);

/*
 * Return a vector containing the global col index corresponding to the beginning
 * of each local column block (column block i corresponds to values at the interface
 * between processor i and my_rank)
 */
int preAlps_OperatorGetColPosPtr(int** colPos, int* sizeColPos);

/*
```

```

    * Return a vector containing the indexes of the neighbor processors.
    */
int preAlps_OperatorGetDepPtr(int** dep, int* sizeDep);

/*
 * Factorize the diagonal block associated to the local row panel owned locally
 * input: A          : a CSR matrix CPLM_Mat_CSR_t
 *        rowPos     : the vector returned by preAlps_OperatorGetRowPosPtr
 *        sizeRowPos : the size of rowPos
 *        colPos     : the vector returned by preAlps_OperatorGetColPosPtr
 *        sizeColPos : the size of colPos
 */
int preAlps_BlockJacobiCreate(CPLM_Mat_CSR_t* A,
                             int* rowPos,
                             int sizeRowPos,
                             int* colPos,
                             int sizeColPos);

/*
 * Solve  $Mx = rhs$  and put the result into rhs.
 * Internal usage only.
 */
int preAlps_BlockJacobiInitialize(CPLM_DVector_t* rhs);

/*
 * Solve  $M B_{out} = A_{in}$  with M a block Jacobi preconditioner.
 */
int preAlps_BlockJacobiApply(CPLM_Mat_Dense_t* A_in, CPLM_Mat_Dense_t* B_out);

/*
 * Free the memory allocated during the construction of the preconditioner.
 */
void preAlps_BlockJacobiFree();

```

## 8 Conclusion

In this document we described a first version of `preAlps` library devoted to solving large sparse linear systems of equations by using preconditioned iterative methods. In its current version, the library implements ECG iterative solver and LORASC preconditioner.

In the context of enlarged Krylov methods, our future work will focus on implementing the dynamic reduction of the number of vectors added during each iteration of ECG such that the computational cost per iteration is reduced while the same convergence rate is maintained. This technique is described in [7]. In the context of LORASC, our performance results show that its application during the iterative process is fast and scalable, but on larger number of processors, its construction becomes expensive since it requires solving a generalized eigenvalue problem of a larger size. Our future work will focus on approaches that allow to speed up the resolution of the generalized eigenvalue problem, which for the moment increases when the number of processors

increases.

## 9 Acknowledgments

This project is funded from the European Union’s Horizon 2020 research and innovation programme under the NLAFET grant agreement No 671633.

## References

- [1] Y. Achdou and F. Nataf. Low frequency tangential filtering decomposition. *Numerical Linear Algebra with Applications*, 14:129–147, 2007.
- [2] P. R. Amestoy, I. S. Duff, J. Koster, and J.-Y. L’Excellent. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.*, 23(1):15–41, 2001.
- [3] Edward Anderson, Zhaojun Bai, Christian Bischof, Susan Blackford, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users’ guide*, volume 9. SIAM, 1999.
- [4] Satish Balay, Shrirang Abhyankar, M Adams, Peter Brune, Kris Buschelman, L Dalcin, W Gropp, Barry Smith, D Karpeyev, Dinesh Kaushik, et al. *Petsc users manual revision 3.7*. Technical report, Argonne National Lab.(ANL), Argonne, IL (United States), 2016.
- [5] L. Susan Blackford, Jaeyoung Choi, Andy Cleary, Eduardo D’Azevedo, James Demmel, Inderjit Dhillon, Jack Dongarra, Sven Hammarling, Greg Henry, Antoine Petit, et al. *ScaLAPACK users’ guide*, volume 4. SIAM, 1997.
- [6] L. Grigori, F. Nataf, and S. Youssef. Robust algebraic schur complement based on low rank correction. Technical report, ALPINES-INRIA, Paris-Rocquencourt, 6 2014.
- [7] L. Grigori and O. Tissot. Reducing the communication and computational costs of enlarged krylov subspaces conjugate gradient. Research Report RR-9023, Feb 2017.
- [8] Laura Grigori, Sophie Moufawad, and Frederic Nataf. Enlarged Krylov subspace conjugate gradient methods for reducing communication. *SIAM J. Matrix Anal. Appl.*, 2016.
- [9] Martin H. Gutknecht. Block Krylov space methods for linear systems with multiple right-hand sides: an introduction. 2006.
- [10] A.Kalhan J. Dongarra, V.Eijkhout. Reverse communication interface for linear algebra templates for iterative methods. Technical report, May 1995.
- [11] George Karypis and Vipin Kumar. METIS –unstructured graph partitioning and sparse matrix ordering system, version 2.0. 1995.
- [12] George Karypis, Kirk Schloegel, and Vipin Kumar. Parmetis: Parallel graph partitioning and sparse matrix ordering library. *Version 1.0, Dept. of Computer Science, University of Minnesota*, 1997.
- [13] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *Arpack User’s Guide: Solution of Large Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, 1997.

- [14] F. Nataf, F. Hecht, P. Jolivet, and C. Prud'Homme. Scalable domain decomposition preconditioners for heterogeneous elliptic problems. *SC13, (Denver, Colorado, United States)*, 2013.
- [15] Q. Niu, L. Grigori, P. Kumar, and F. Nataf. Modified tangential frequency filtering decomposition and its fourier analysis. *Numerische Mathematik*, 116:123–148, 2010.
- [16] D. P. O'Leary. The block conjugate gradient algorithm and related methods. *Linear Algebra and Its Applications*, 29:293–322, 1980.
- [17] Endong Wang, Qing Zhang, Bo Shen, Guangyong Zhang, Xiaowei Lu, Qing Wu, and Yajuan Wang. Intel math kernel library. In *High-Performance Computing on the Intel® Xeon Phi*, pages 167–188. Springer, 2014.