



HAL
open science

On the Use of Block Low Rank Preconditioners for Primal Domain Decomposition Methods

Christophe Bovet, Théodore Gauthier, Pierre Gosselet

► **To cite this version:**

Christophe Bovet, Théodore Gauthier, Pierre Gosselet. On the Use of Block Low Rank Preconditioners for Primal Domain Decomposition Methods. *International Journal for Numerical Methods in Engineering*, 2024, pp.e7623. <10.1002/nme.7623>. <hal-04830430>

HAL Id: hal-04830430

<https://hal.science/hal-04830430v1>

Submitted on 11 Dec 2024

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.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

RESEARCH ARTICLE OPEN ACCESS

On the Use of Block Low Rank Preconditioners for Primal Domain Decomposition Methods

Christophe Bovet¹  | Théodore Gauthier^{1,2} | Pierre Gosselet²

¹Université Paris-Saclay, ONERA, Matériaux et Structures, Châtillon, France | ²LaMcube, Univ. Lille/CNRS/Centrale Lille, Lille, France

Correspondence: Christophe Bovet (christophe.bovet@onera.fr)

Received: 29 July 2024 | **Revised:** 24 October 2024 | **Accepted:** 29 October 2024

Funding: The authors received no specific funding for this work.

Keywords: adaptive multipreconditioning | block low rank factorizations | domain decomposition

ABSTRACT

This article investigates the use of the block low rank (BLR) factorization, recently proposed in the MUMPS solver, to define efficient and cheap preconditioners for primal domain decomposition methods, such as the Balancing Domain Decomposition method (BDD) and its adaptive multipreconditioned variant. To be scalable, these methods are equipped with an augmentation projector built from the local preconditioners nullspaces. The determination of these nullspaces is a complex task in the case of ill conditioned system, the use of block low rank compression makes this task even more complex as MUMPS' automatic detection no longer works properly. Two alternatives based on incomplete factorization with a well-chosen Schur complement are proposed. Also, the first massively parallel implementation of the adaptive multipreconditioned BDD solver (AMPBDD) is introduced. The performance of the methods is assessed with two weak scalability studies on problems up to 24,576 cores and about 790 millions of unknowns, on the Sator and Topaze supercomputers. BLR preconditioning proves to be an interesting strategy both in terms of memory usage and time to solution for reasonably conditioned problems.

1 | Introduction

In the last decade, non-overlapping domain decomposition methods have reached a high level of maturity, with sophisticated robustification techniques, and high performance implementations. Even though these questions are still the object of intense research, another question of interest is the ability to derive less numerically demanding variants of the methods which result in better performance in practice on sufficiently regular problems.

Considering a linear elasticity problem to fix the ideas, non-overlapping domain decomposition methods consist in solving independent problems on the subdomains, for a given boundary value imposed on the interface. Dual methods (FETI [1], AMPFETI [2–4]) search for the balanced traction

condition on the interface that nullify the displacement gap. They provide a zoology of preconditioners with variable quality and computational cost, allowing them to adapt to the conditioning of the system to be solved.

However, they are less suited for the simulation of fracture problems, such as damage and crack propagation, due to their high sensitivity to the computation of the nullspace of local stiffness operators. A wrong estimation of these kernels leads to the divergence of the Krylov solver and/or the FETI system not being equivalent to the initial one anymore. The FETI-DP method [5] only partially solves this issue. Indeed, this approach enforces the continuity between subdomains of certain (generalized) degrees of freedom (like corner nodes or averages on faces/edges), leading to all Neumann problems being well-posed, without local

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDeriv](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2024 The Author(s). *International Journal for Numerical Methods in Engineering* published by John Wiley & Sons Ltd.




FIGURE 1 | Following Liu [16], a case of a $L^T L$ factorization denoting $L = (\ell_{ij})$, the tree is explored in parallel from bottom to top.

- Sparsity of data can be further exploited by compressing the Front in low-rank format as explained below.

Before explaining block low-rank formats, low-rank matrices are defined. Let ϵ_p be a strictly positive threshold. Let n and k be non-negative integers such that $k \leq n$. Let \mathbf{M} be a $n \times n$ matrix. The numerical rank of \mathbf{M} at precision ϵ_p is k if and only if k is the lowest integer such that there exist \mathbf{X}_M and \mathbf{Y}_M , $n \times k$ rectangle matrices, such that:

$$\|\mathbf{M} - \mathbf{X}_M \mathbf{Y}_M^T\|_2 \leq \epsilon_p$$

with $\|\cdot\|_2$ the matrix norm induced by the Euclidean norm. The matrix \mathbf{M} is said to be low-rank for a given accuracy ϵ_p when storing \mathbf{X}_M and \mathbf{Y}_M requires less memory rather than \mathbf{M} . Thus, \mathbf{M} is low-rank whenever the following inequality holds:

$$2kn \leq n^2$$

If so, the low-rank approximation $\tilde{\mathbf{M}} = \mathbf{X}_M \mathbf{Y}_M^T$ of \mathbf{M} is stored as $(\mathbf{X}_M, \mathbf{Y}_M)$. In practice \mathbf{X}_M and \mathbf{Y}_M are obtained directly through Singular Values Decomposition or Rank Revealing QR (RRQR) factorization of \mathbf{M} in which case \mathbf{X} is orthogonal. As well as reducing memory requirements, the low-rank format reduces the complexity of algebraic operations. Indeed, operations on and between two low-rank matrices $\tilde{\mathbf{M}}$ and $\tilde{\mathbf{N}}$ can be achieved directly – exactly or approximately – on $\mathbf{X}_M, \mathbf{Y}_M, \mathbf{X}_N$ and \mathbf{Y}_N with reduced complexities. These operations include multiplication, addition, row or column swap, etc. Further explanations and computations of complexities regarding these operations can be readily found in the introduction of Bebendorf [17] or in Mary [8].

Obviously, global finite element matrices are not low-rank, but some extradiagonal subblocks representing long-range interactions may be. Thus, a *block* low-rank (BLR) factorization of \mathbf{A} is obtained by putting some of these off-diagonal blocks into low-rank format or “compressing” them whenever the expected gains are greater than the overhead of doing so. In practice, a graph based analysis is done as to decide which blocks will be compressed or not, rather than doing unnecessary RRQR factorizations to obtain the numerical rank. It permits assessing with a purely algebraic criterion the geometric distance between the nodes considered (see *admissibility condition* in Mary [8])

as one should expect the rank to decrease exponentially with reasonable hypothesis made and on elliptical partial differential equations [17].

Low-rank factorization consists in compressing the Frontal Matrices in block low-rank format to carry on the partial factorization with reduced complexity and memory footprint. For this work, we have chosen to use the MUMPS library. The user has control on several parameters:

- the precision ϵ_{BLR} which differs from ϵ_p previously mentioned only to a scaling;
- the ordering of some operations, leading to two variants: UCFS and UFSC.

Let us explain the last two denominations. During the factorization of a given Front F , a loop is done on its blocks. Data are accessed as late as possible within the Front F , and its blocks are Updated (that is, the contribution from pivots within the front are applied, step (U)) just before we start to treat them. Step (F) corresponds to the Factorization of the diagonal block before carrying out division by the unitary lower triangular matrix obtained on blocks below-the Solve step (S). To make numerical pivoting possible, it is needed to merge the Factor and Solve steps together, which requires Compression (C) to take place either before or after the Factor+Solve (FS) step. In either case, UCFS and UFSC make pivoting possible. The variant UCFS should reduce even more the complexity but at the cost of degraded numerical pivoting because it is done in low-rank, whereas UFSC should be more precise but will not benefit from early compression. It is stated in Mary [8] that the downside of UCFS is barely noticeable as a degraded solution could quickly be improved through cheap iterative refinement steps. We have used both variants in the current work to investigate if any differences could be highlighted either way. Regarding the iterative refinement process, our numerical tests have shown little interest in our case, it will not be used in the following.

Finally, one should note that there are other methods exploiting numerical ranks in the literature. One might be interested in the brief review of some of these methods in Mary [8], which states for instance that BLR factorization should be preferred for solving

systems repeatedly while Hierarchically Semi-Separable matrices (HSS) [17] should be used for aggressive preconditioning. Since our goal is to build a cheap preconditioner which approximates precisely enough the action of some generalized inverse repeatedly, this advocates for the use of BLR factorization.

3 | Primal Domain Decomposition Methods

This section briefly recalls the Balancing domain decomposition method (BDD [6]), its coupling with Adaptive Multipreconditioning (AMP [18]), and the use of inexact solvers.

3.1 | Balancing Domain Decomposition Method in a Nutshell

We consider a linear(ized) elasticity problem set on a domain Ω and discretized with the finite element method. This results in a large sparse linear system of equations of the form $\mathbf{K}\mathbf{u} = \mathbf{f}$ where \mathbf{u} is the vector of unknowns and \mathbf{f} the right-hand side. The operator \mathbf{K} (stiffness matrix) is assumed to be symmetric positive definite.

Let $(\Omega^s)_{1 \leq s \leq N_d}$ be a non overlapping partition of Ω such that: $\bar{\Omega} = \bigcup_{s=1}^{N_d} \bar{\Omega}^s$ and $\Omega^s \cap \Omega^p = \emptyset, \forall s \neq p$. The interface between the subdomains Ω^s and Ω^p is denoted by $\Upsilon^{sp} = \bar{\Omega}^s \cap \bar{\Omega}^p$, the union of all the interfaces of the subdomain Ω^s is denoted by Υ^s , and the union of the interfaces of all subdomains is denoted by Υ . In the substructured formulations, only local quantities (e.g. restricted to one subdomain) are assembled such as the matrices \mathbf{K}^s and the right-hand-side \mathbf{f}^s . The global system is equivalent to the substructured formulation:

$$\mathbf{K}^s \mathbf{u}^s = \mathbf{f}^s + \mathbf{T}^{s\top} \lambda_b^s \quad \forall 1 \leq s \leq N_d \quad (1)$$

$$\sum_{s=1}^{N_d} \mathbf{B}^s \mathbf{T}^s \mathbf{u}^s = \mathbf{0} \quad (2)$$

$$\sum_{s=1}^{N_d} \mathbf{A}^s \lambda_b^s = \mathbf{0} \quad (3)$$

where $\mathbf{T}^s : \Omega^s \rightarrow \Upsilon^s$ is the trace operator, \mathbf{A}^s and \mathbf{B}^s are primal and dual assembly operators respectively (see [19] for their definition). The Lagrange multiplier field λ_b^s enforces the continuity of the primal unknown across the subdomains interfaces. From a mechanical point of view, Equation (1) are the equilibrium of all subdomains, (2) corresponds to the continuity of the displacement across interfaces and (3) expresses the equilibrium of the interface (action-reaction principle).

All unknowns can be separated between internal unknowns (denoted with subscript i) and boundary ones (denoted with subscript b). Internal degrees of freedom can be eliminated in order to express (1)–(3) only in terms of boundary unknowns

$$\mathbf{S}^s \mathbf{u}_b^s = \hat{\mathbf{f}}_b^s + \lambda_b^s \quad \forall 1 \leq s \leq N_d \quad (4)$$

$$\sum_{s=1}^{N_d} \mathbf{B}^s \mathbf{u}_b^s = \mathbf{0} \quad (5)$$

$$\sum_{s=1}^{N_d} \mathbf{A}^s \lambda_b^s = \mathbf{0} \quad (6)$$

where \mathbf{S}^s and $\hat{\mathbf{f}}_b^s$ are primal Schur complements and condensed right-hand sides. The vector \mathbf{u}_b^s is the trace of the displacement at the boundary.

$$\mathbf{S}^s = \mathbf{K}_{bb}^s - \mathbf{K}_{bi}^s \mathbf{K}_{ii}^{s-1} \mathbf{K}_{ib}^s \quad (7)$$

$$\hat{\mathbf{f}}_b^s = \mathbf{f}_b^s - \mathbf{K}_{bi}^s \mathbf{K}_{ii}^{s-1} \mathbf{f}_i^s \quad (8)$$

Finally, we would like to point out that assembly operators are orthogonal in the following sense:

$$\sum_{s=1}^{N_d} \mathbf{B}^s \mathbf{A}^{s\top} = \mathbf{0} \quad (9)$$

which means that any local interface vector \mathbf{u}_b^s can be uniquely defined as a combination of a balanced vector and a continuous one $\mathbf{u}_b^s = \mathbf{B}^{s\top} \mathbf{u}_d + \mathbf{A}^{s\top} \mathbf{u}_p$.

The BDD method writes the interface problem in terms of one unique primal global unknown \mathbf{u}_p such that local interface vectors are given by $\mathbf{u}_b^s = \mathbf{A}^{s\top} \mathbf{u}_p$ and (5) is satisfied by construction thanks to the orthogonality property of assembly operators (9). Few algebraic manipulations lead to the primal formulation:

$$\underbrace{\sum_{s=1}^{N_d} \mathbf{A}^s \mathbf{S}^s \mathbf{A}^{s\top}}_{\mathbf{S}} \mathbf{u}_p - \underbrace{\sum_{s=1}^{N_d} \mathbf{A}^s \hat{\mathbf{f}}_b^s}_{\mathbf{f}_p} = \sum_{s=1}^{N_d} \mathbf{A}^s \lambda_b^s = \mathbf{0} \quad (10)$$

The global primal Schur complement $\mathbf{S} = \sum_{s=1}^{N_d} \mathbf{A}^s \mathbf{S}^s \mathbf{A}^{s\top}$ is never built explicitly. Since this system is solved using a Krylov iterative solver, only the result of a multiplication by \mathbf{S} is needed. This computation is well suited to parallel computers since \mathbf{S} is a sum of local contributions. Since Schur complements \mathbf{S}^s are dense matrices, they are not computed explicitly. The action of these Schur operators is evaluated implicitly with two sparse matrix vector products and one local solve with fixed boundary. Direct solvers without compression are used for these local solutions.

3.1.1 | First Level Preconditioner

The BDD preconditioner is the composition of the Neumann Neumann preconditioner \mathbf{M}_{NN}^{-1} with the coarse projector defined in the next subsection. The Neumann Neumann preconditioner mimics the additive structure of \mathbf{S} , it is chosen as a scaled sum of generalized inverse of primal Schur complements defined by

$$\mathbf{M}_{NN}^{-1} = \sum_{s=1}^{N_d} \tilde{\mathbf{A}}^s \mathbf{S}^{s\ddagger} \tilde{\mathbf{A}}^{s\top} \quad (11)$$

where $\tilde{\mathbf{A}}^s$ are scaled primal assembly operators such that $\sum_s \mathbf{A}^s \tilde{\mathbf{A}}^{s\top} = \mathbf{I}_\Upsilon$, and the superscript $\mathbf{S}^{s\ddagger}$ denotes for a generalized inverse of \mathbf{S}^s . Classical scaling operators are multiplicity scaling and stiffness scaling (often called **k**-scaling) [20].

The action of $\mathbf{S}^{s\ddagger}$ is obtained by solving a local problem with Neumann boundary conditions. Depending on the natural boundary

conditions of the problem, S^s may be singular. Corresponding subdomains are commonly qualified as “floating subdomains”.

3.1.2 | Second Level Preconditioner: Coarse Problem

The Neumann Neumann preconditioner is applied to the residual of the Krylov solver $\mathbf{z} = \mathbf{M}_{NN}^{-1} \mathbf{r}$. For floating subdomains, local right-hand-sides must lie inside the image of S^s which leads to the solvability conditions:

$$\mathbf{R}_b^{s\top} \tilde{\mathbf{A}}^{s\top} \mathbf{r} = \mathbf{0}, \quad \forall s \quad (12)$$

where \mathbf{R}_b^s is the nullspace of S^s . We rewrite this condition as $\mathbf{C}^\top \mathbf{r} = \mathbf{0}$ with

$$\mathbf{C} = \left(\tilde{\mathbf{A}}^1 \mathbf{R}_b^1 \mid \dots \mid \tilde{\mathbf{A}}^{N_d} \mathbf{R}_b^{N_d} \right) \quad (13)$$

These conditions provide the BDD coarse problem which is enforced using an augmented Krylov solver. An augmentation projector $\mathbf{\Pi}_C$ such that $\mathbf{C}^\top \mathbf{S} \mathbf{\Pi}_C = \mathbf{0}$ is defined, and the solution is sought as:

$$\mathbf{u}_p = \mathbf{u}_0 + \mathbf{\Pi}_C \tilde{\mathbf{u}} \quad (14)$$

$$\text{with } \mathbf{u}_0 = \mathbf{C}(\mathbf{C}^\top \mathbf{S} \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{f}_p \quad (15)$$

$$\mathbf{\Pi}_C = \mathbf{I} - \mathbf{C}(\mathbf{C}^\top \mathbf{S} \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{S} \quad (16)$$

The system to be solved by the Krylov solver is finally:

$$\mathbf{S} \mathbf{\Pi}_C \tilde{\mathbf{u}} = (\mathbf{f}_p - \mathbf{S} \mathbf{u}_0), \text{ preconditioned by } \mathbf{M}_{NN}^{-1} \quad (17)$$

This coarse problem provides a mechanism for rapidly propagating the mechanical load information to all the subdomains which is essential to build a scalable method.

3.2 | Block Low Rank BDD Preconditioner

In order to build a cheap preconditioner with a small memory footprint for the BDD method, the basic idea is to replace the full-rank resolution of local problems with a compressed resolution. The BLR BDD Preconditioner \mathbf{M}_{BLR}^{-1} can be written as:

$$\mathbf{M}_{BLR}^{-1} = \sum_{s=1}^{N_d} \tilde{\mathbf{A}}^s \mathbf{S}_{BLR}^{s\top} \tilde{\mathbf{A}}^{s\top} \quad (18)$$

where $\mathbf{S}_{BLR}^{s\top}$ stands for a resolution with a BLR factorization. Available choices for the scaling remain unchanged. The compression threshold ϵ_{BLR} and the variants (UCFS, UFSC) are parameters of the preconditioner. The former offers real flexibility in terms of the cost and quality of the preconditioner. For simplicity, all subdomains use the same threshold and variant. Mumps proposes an iterative refinement process when using compressed factorization. This iterative refinement is not used in the following since it represents a significant additional cost and there is no guarantee that the solution will be improved.

The main disadvantage of the compressed preconditioner concerns the floating subdomains and the coarse problem. Indeed,

depending on the compression level, the nullspace of local operator of floating subdomains may be lost or Mumps may be not able to compute it correctly, thus leading to a degraded scalability. To overcome this problem, other ways of constructing the coarse problem are examined in Section 4. Another possibility to recover a coarse grid mechanism is to rely on multipreconditioning and not on local operators nullspace.

3.3 | Adaptive Multipreconditioning

Multipreconditioning was proposed for iterative solvers [21] and adapted to domain decomposition methods [2]. It is a strategy which exploits the additive structure of the preconditioner in order to generate as many search directions as subdomains. In the preconditioning step of conjugate gradient, instead of computing the preconditioned residual as $\mathbf{z}_i = \sum_{s=1}^{N_d} \tilde{\mathbf{A}}^s \mathbf{S}^{s\top} \tilde{\mathbf{A}}^{s\top} \mathbf{r}_i$, the following block of vectors is generated: $\mathbf{Z}_i = \left(\dots \tilde{\mathbf{A}}^s \mathbf{S}^{s\top} \tilde{\mathbf{A}}^{s\top} \mathbf{r}_i \dots \right)$. Of course the classical direction writes $\mathbf{z}_i = \mathbf{Z}_i \mathbf{1}$, where $\mathbf{1}$ is the vector filled with ones. The idea of multipreconditioning is to let the algorithm find the optimal combination of directions under the form $\mathbf{z}_i = \mathbf{Z}_i \boldsymbol{\alpha}_i$, where $\boldsymbol{\alpha}_i$ is the unknown vector of subdomains' magnitude of contribution.


Multipreconditioning must be used in conjunction with full reorthogonalization, and it is a numerically expensive option. Nevertheless, it proved to be an efficient cure to FETI and BDD's bad conditioning situations. Indeed, it was proved [22] that multipreconditioning is a technique to approximate on the fly the bad modes that would be detected and eliminated by GENE0 coarse spaces [23].

In order to limit the numerical costs of multipreconditioning, a clever adaptation strategy was proposed by Spillane [18] where the effectiveness of each direction is predicted based on a costless criterion, making it possible to accumulate directions which contribute weakly to the decrease of the error and limit the memory and CPU footprint. A large-scale assessment of this approach has been carried out for the FETI method [3]. The method was further improved with more sophisticated aggregation of search directions depending on the subdomains' connectivity [4].


As mentioned in previous subsection, BLR-preconditioning may cause the disappearance of nullspace modes and BDD-coarse space may not be a numerical necessity to preserve the well-posedness of Neumann problems. Even though there is a mechanical urge to preserve rigid body motions coarse space in order to comply with Saint-Venant's principle, we wish to evaluate the ability of multipreconditioning to naturally bring out this information and ensure scalability.

4 | Coarse Problem Computation: Nullspace and Generalized Inverses

The coarse problem of the BDD method relies on the computation local preconditioner nullspace and generalized inverses. The exact computation of these kernels is not mandatory since it plays at the preconditioning level. However, it still impacts the rate of convergence. In this section we propose three different methods to evaluate the defect k^s , the nullspace of \mathbf{K}^s and its generalized



(a) Checkerboard cube. Red and blue areas correspond to the two different materials.



(b) Automatic decomposition. Each color represents a different subdomain.

FIGURE 2 | Heterogeneous cube (configuration with $N_d = 216$, $n_c = 6$). (a) Checkerboard cube. Red and blue areas correspond to the two different materials. (b) Automatic decomposition. Each color represents a different subdomain.

TABLE 1 | Checkerboard cube, weak parallel scalability: configurations.

n_c	N_d	#DOFs total	#cores
4	64	12.52M	384
6	216	41.99M	1,296
8	512	99.22M	3,072
10	1000	193.44M	6,000
16	4096	790.12M	24,576

inverse. For readability, we drop the exponent s of the local operator \mathbf{K}^s .

4.1 | Mumps Automatic Nullspace Detection (M)

The first method is simply to use Mumps' capabilities to evaluate the operator nullspace and null pivots. There are two user defined parameters for the detection of the kernel dimension in Mumps, CNTL(1) and CNTL(3). The control parameter CNTL(1) is a relative threshold for numerical pivoting. The default value CNTL(1) = 10^{-2} is used in this work. The second control parameter CNTL(3) is a threshold to detect null pivots. According to the documentation, a pivot is considered to be null if the infinite norm of its row/column is smaller than a threshold $thres$. The default value of CNTL(3) = 0 provides an automatic process to determine this threshold, $thres = \epsilon \times 10^{-5} \times \|A_{pre}\|$ where A_{pre} is the preprocessed matrix to be factorized and ϵ is machine precision. A positive value of CNTL(3) leads to the user defined threshold $thres = CNTL(3) \times \|A_{pre}\|$.

As shown in a previous work [24], the automatic kernel detection can be put on severe test when dealing with ill-conditioned systems. Often, the automatic threshold does not detect the right kernel size. It is however possible to recover the right kernel with a user defined threshold, but the admissible range for CNTL(3) becomes narrow. If Mumps allows both BLR compression and

nullspace calculation to be enabled, we expect the estimation of the correct nullspace to be even more complex in those cases.

4.2 | Incomplete Factorization and Fixing-Nodes Framework

The other two methods reuse the graph based approach proposed in a previous work [24]. This framework is briefly recalled here, we refer to the original paper and the references therein for more details. The overall methodology relies on the partial factorization of the operator and on the analysis of a well-chosen Schur complement. Let c be a nonempty subset of $\{1, \dots, n\}$, called fixing variables, the incomplete LL^T factorization is:

$$\mathbf{K} = \begin{bmatrix} \mathbf{K}_{\bar{c}\bar{c}} & \mathbf{K}_{\bar{c}c} \\ \mathbf{K}_{c\bar{c}} & \mathbf{K}_{cc} \end{bmatrix} = \begin{bmatrix} \mathbf{L}_{\bar{c}\bar{c}} & \mathbf{0} \\ \mathbf{L}_{c\bar{c}} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_{\bar{c}\bar{c}}^T & \mathbf{L}_{c\bar{c}}^T \\ \mathbf{0} & \mathbf{S}_{cc} \end{bmatrix} \quad (19)$$

The construction of c in the paper [24], based on graph centrality measures, not only ensures that $\mathbf{K}_{\bar{c}\bar{c}}$ remain full-rank, but it minimizes its condition number, which is an important feature when dealing with large ill conditioned systems. A generalized inverse \mathbf{K}^+ of \mathbf{K} is given by

$$\mathbf{K}^+ = \begin{bmatrix} \mathbf{L}_{\bar{c}\bar{c}}^{-T} & -\mathbf{L}_{\bar{c}\bar{c}}^{-T} \mathbf{L}_{c\bar{c}}^T \mathbf{S}_{cc}^\dagger \\ \mathbf{0} & \mathbf{S}_{cc}^\dagger \end{bmatrix} \begin{bmatrix} \mathbf{L}_{\bar{c}\bar{c}}^{-1} & \mathbf{0} \\ -\mathbf{L}_{c\bar{c}} \mathbf{L}_{\bar{c}\bar{c}}^{-1} & \mathbf{I} \end{bmatrix} \quad (20)$$

Since the Schur complement \mathbf{S}_{cc} is a small dense matrix, the use of the Moore-Penrose generalized inverse, obtained by SVD, is reliable and affordable here. From a practical point of view, once the fixing variables have been selected, the partial factorization is performed with the Mumps library. The fact that Mumps can activate both BLR compression and partial factorization is a very interesting opportunity here.

What remains to be done is to choose a criterion to determine Moore-Penrose generalized inverse of \mathbf{S}_{cc} , and a way to compute a basis of the nullspace.

TABLE 2 | Small checkerboard cube, summary of the results with $E_r/E_b = 10^0$ (homogeneous case).

Heterogeneity $E_r/E_b = 10^0$				Solver CG			Solver AMPCG			
BLR	ϵ_{BLR}	Kernel	#C	#iter	t(s)	S^{s+} (GB)	#iter	t(s)	S^{s+} (GB)	#s.dir.
		M	192	65	72	2.5	59	71	2.5	90
UCFS	10^{-1}	M	0	169	109	1.4	168	112	1.4	199
UCFS	10^{-3}	M	0	120	90	1.8	115	93	1.7	146
UCFS	10^{-5}	M	0	257	165	2.1	58	77	2.1	358
UFSC	10^{-1}	M	0	168	108	1.4	168	112	1.4	199
UFSC	10^{-3}	M	0	115	89	1.8	120	94	1.7	151
UFSC	10^{-5}	M	0	257	167	2.1	62	79	2.1	347
		G	192	65	79	2.5	59	77	2.5	90
UCFS	10^{-1}	G	192	98	89	1.4	97	90	1.4	128
UCFS	10^{-3}	G	192	61	72	1.8	58	72	1.8	89
UCFS	10^{-5}	G	192	65	77	2.1	59	76	2.1	90
UFSC	10^{-1}	G	192	98	88	1.4	97	91	1.4	128
UFSC	10^{-3}	G	192	61	72	1.8	59	73	1.7	90
UFSC	10^{-5}	G	192	65	77	2.2	59	75	2.1	90
		E	192	65	81	2.4	59	78	2.4	90
UCFS	10^{-1}	E	192	165	125	1.3	151	122	1.3	182
UCFS	10^{-3}	E	192	66	75	1.6	63	76	1.6	94
UCFS	10^{-5}	E	192	65	77	2.0	59	76	2.0	90
UFSC	10^{-1}	E	192	151	119	1.3	151	121	1.3	182
UFSC	10^{-3}	E	192	66	75	1.7	63	76	1.6	94
UFSC	10^{-5}	E	192	65	78	2.0	59	76	2.0	90

4.2.1 | Low Energy Modes (E)

Let $(\sigma_j)_{1 \leq j \leq c}$ be the singular values of S_{cc} such that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_c \geq 0$. With this method, a relative criterion $\sigma_j \leq \epsilon \sigma_1$ is used to estimate the “null” singular values. The singular value decomposition of S_{cc} also provides the nullspace of the Schur complement R_c and the nullspace of the full matrix is deduced from R_c :

$$R = \begin{bmatrix} -K_{cc}^{-1} K_{cc} R_c \\ R_c \end{bmatrix} \quad (21)$$

where K_{cc}^{-1} makes use of the BLR compression. Here both the coarse space C and the coarse projector Π_C take into account the BLR compression.

4.2.2 | Hybrid Geometric–Algebraic Detection (G)

In our experiments, it appeared that BLR compression may have a strong impact on the estimation of the defect (size of the nullspace) and on the basis input in the coarse problem, while Saint-Venant’s principle urges us to preserve actual rigid body motions for the coarse problem. Thus, we propose another strategy inspired by the hybrid geometric–algebraic approach of Farhat and Gérardin [25].

The method requires knowing the nullspace in the case of a totally floating subdomain. Let R_u be a basis of this totally unrestrained nullspace. In 3D elastostatics on connected domains, R_u

is made of the six rigid body modes (3 translations and 3 rotations). The method of Farhat and Gérardin [25] permits to calculate the combinations of rigid body motions which are not precluded by the Dirichlet conditions. These combination form the actual nullspace R of the subdomain.

Once the dimension k of the nullspace is known, the generalized inverse is computed using Equation (20) where the Moore–Penrose generalized inverse S_{cc}^\dagger considers that the k -smallest singular values are zero. This treatment differs from the original paper of Farhat and Gérardin [25] where exactly k fixing nodes were deduced from the knowledge of the nullspace.

The hybrid geometric–algebraic method leads to coarse space C and projector Π_C being the same as those constructed without compression. Only the generalized inverse is impacted by the BLR compression.

5 | Numerical Experiments

5.1 | Remarks on the Implementation and Dependencies

The proposed methods have been implemented in the finite element suite Z-Set 9.1¹. In all configurations, the local direct solves are performed with the MUMPS library (version 5.5.1) [26]. MUMPS is linked with the BLAS library provided by Intel MKL.

TABLE 3 | Small checkerboard cube, summary of the results with $E_r/E_b = 10^2$ (moderate heterogeneity).

Heterogeneity $E_r/E_b = 10^2$				Solver CG			Solver AMPCG			
BLR	ϵ_{BLR}	Kernel	#C	#iter	t(s)	S^{s+} (GB)	#iter	t(s)	S^{s+} (GB)	#s.dir.
		M	192	142	115	2.5	102	102	2.5	288
UCFS	10^{-1}	M	0	> 500		1.4	> 500		1.4	
UCFS	10^{-3}	M	0	345	204	1.8	232	168	1.8	515
UCFS	10^{-5}	M	0	335	209	2.2	124	120	2.1	572
UFSC	10^{-1}	M	0	> 500		1.4	> 500		1.4	
UFSC	10^{-3}	M	0	353	206	1.7	240	173	1.8	537
UFSC	10^{-5}	M	0	335	210	2.1	158	137	2.2	548
		G	192	142	125	2.5	102	110	2.5	288
UCFS	10^{-1}	G	192	320	209	1.4	308	211	1.4	370
UCFS	10^{-3}	G	192	208	153	1.7	138	125	1.8	324
UCFS	10^{-5}	G	192	142	120	2.1	106	108	2.1	261
UFSC	10^{-1}	G	192	323	211	1.4	319	215	1.4	350
UFSC	10^{-3}	G	192	221	160	1.7	147	129	1.8	333
UFSC	10^{-5}	G	192	142	120	2.1	107	110	2.1	262
		E	192	142	126	2.4	102	112	2.4	288
UCFS	10^{-1}	E	192	> 500		1.3	> 500		1.3	
UCFS	10^{-3}	E	192	297	206	1.6	215	174	1.6	399
UCFS	10^{-5}	E	192	142	122	2.0	106	109	2.0	261
UFSC	10^{-1}	E	192	> 500		1.3	> 500		1.3	
UFSC	10^{-3}	E	192	259	186	1.7	216	175	1.7	433
UFSC	10^{-5}	E	192	142	123	2.0	106	110	2.0	261

The coarse problem is solved with the Pardiso direct solver. The Eigen library² is used for dense linear algebra. Communication are handled by the MPI protocol. The MPI library depends on the supercomputer used.

5.2 | Description of the Weak Scaling Test Case

For $n_c \in \{4, \dots, 16\}$, we consider a set of three-dimensional heterogeneous cubes made of n_c^3 identical sub-cubes (see Figure 2). Each sub-cube is discretized with the same ruled mesh made of 64,000 eight-node brick elements (c3d8), leading to a total number of $3 \times (40 \times n_c + 1)^3$ degrees of freedom. With this setup, the H/h ratio equals 40 where h is the diameter of the finite elements and H that of the subdomains.

The cube is clamped on one face and subjected to a prescribed unitary displacement in the three space directions on the opposite face, all other faces being traction-free. The material behavior is isotropic linear elastic, with a Poisson's coefficient of 0.3 and two values of Young's modulus assigned following a checkerboard pattern in order to obtain a coefficient jump E_r/E_b between two adjacent sub-cubes. Three ratios of Young's modulus are used: 10^0 , 10^2 and 10^4 . Finally, an unstructured decomposition in $N_d = n_c^3$ subdomains is obtained with a graph partitioning software which leads to interfaces not aligned with the heterogeneity. For a given number of subdomains, the partitioning is computed once and reused for all solvers configurations and for

both coefficient jumps. The choice $N_d = n_c^3$, combined with the use of an automatic graph partitioning software leads to a lot of traversing heterogeneities that are known to strongly deteriorate the convergence of domain decomposition methods. Such a configuration is represented in Figure 2 for $n_c = 6$.

All preconditioners make use of the stiffness scaling, they differ by the local operator S^{s+} (with or without BLR compression) and the way to construct the coarse space C . To make it easier to identify the method used to build the coarse problem, each method is assigned a letter (see the column Kernel in Table 2 for instance):

- M refers to the Mumps automatic nullspace detection (Section 4.1),
- G stands for the geometric-algebraic detection (Section 4.2.2),
- E corresponds to the low energy modes (Section 4.2.1).

The convergence is triggered when $\|r_i\|/\|r_0\| \leq \epsilon = 10^{-6}$. When AMPCG is used, the number of aggregates is 32 and the τ -test threshold is set to 10^{-2} .

Six cores are allocated to each subdomain, a shared memory parallelism is used at several steps including (but not limited to) local operators and coarse problem factorization. The study starts

TABLE 4 | Small checkerboard cube, summary of the results with $E_r/E_b = 10^4$ (high heterogeneity).

Heterogeneity $E_r/E_b = 10^4$				Solver CG			Solver AMPCG			
BLR	ϵ_{BLR}	Kernel	#C	#iter	t(s)	S^{s+} (GB)	#iter	t(s)	S^{s+} (GB)	#s.dir.
		M	149	> 500		2.5	> 500		2.5	
UCFS	10^{-1}	M	0	> 500		1.4	> 500		1.4	
UCFS	10^{-3}	M	0	> 500		1.7	313	267	1.7	1189
UCFS	10^{-5}	M	0	> 500		2.1	183	186	2.1	1122
UFSC	10^{-1}	M	0	> 500		1.4	> 500		1.4	
UFSC	10^{-3}	M	0	> 500		1.7	286	254	1.8	1198
UFSC	10^{-5}	M	0	> 500		2.1	> 500		2.1	
		G	192	393	291	2.5	108	141	2.5	851
UCFS	10^{-1}	G	192	> 500		1.4	393	308	1.4	877
UCFS	10^{-3}	G	192	> 500		1.8	275	256	1.8	1046
UCFS	10^{-5}	G	192	393	276	2.1	108	140	2.1	895
UFSC	10^{-1}	G	192	> 500		1.4	335	276	1.4	904
UFSC	10^{-3}	G	192	> 500		1.8	229	222	1.7	1034
UFSC	10^{-5}	G	192	393	273	2.1	108	140	2.1	922
		E	192	393	288	2.4	108	143	2.4	851
UCFS	10^{-1}	E	192	> 500		1.3	> 500		1.3	
UCFS	10^{-3}	E	192	> 500		1.6	327	301	1.6	1097
UCFS	10^{-5}	E	192	> 500		2.0	223	232	2.0	1036
UFSC	10^{-1}	E	192	> 500		1.3	> 500		1.3	
UFSC	10^{-3}	E	192	> 500		1.6	377	351	1.6	1213
UFSC	10^{-5}	E	192	> 500		2.0	208	221	2.0	1056

from 64 subdomains and goes up to 4096 subdomains which corresponds to a total number of 24,576 cores and 790.12 millions unknowns. Table 1 summarizes the different configurations.

5.3 | Weak Scalability Study on the Sator Supercomputer

5.3.1 | Presentation of the Hardware

Sator is Onera's in-house supercomputer. It is a parallel scalar cluster with 43,600 cores supplied by NEC. Thanks to three groups of computing nodes (Broadwell, Skylake and Cascade Lake), the Linpack performance of Sator is 1.8 PFlop/s. In this work, only the Cascade Lake partition has been used. It is made of 400 compute nodes with Intel Xeon "Cascade Lake 6240R" bi-processors (19,200 cores). Each node has 2×24 cores at 2.4 GHz and 192 GB of RAM (4GB RAM per core). The interconnection network is based on an Intel Omnipath 100Gbps fabric, in a Fat-tree topology. Communications are handled with Intel MPI 22.2.0. Since the largest queue in Sator is limited, the weak scalability only goes up to 3,072 cores in this section.

5.3.2 | Focus on a Small Test Case ($N_d = 64$)


In order to reduce the number of calculations and select only the most promising configurations, the focus is made on the smallest

test case with 64 subdomains and 384 cores. Several counters and timers are provided to compare the results:


- The size of the coarse problem is shown in column #C.
- The column t(s) represents the total time of the simulation, including the construction and the factorization of the local operators, the computation of the coarse problem and the time spent in the iterations.
- The column S^{s+} (GB) shows the memory footprint of the local preconditioner.
- For AMPCG, the number of search directions is given in column #s.dir. (for CG it equals the number of iterations since we use full reorthogonalization).

5.3.2.1 | Homogeneous Problem. The results of the homogeneous test case are summarized in Table 2. This test case being well conditioned, all variants converge in less than 500 iterations. As expected, the convergence is strongly degraded when Mumps loses the nullspace due to the BLR compression. AMPCG is able to compensate for this loss for $\epsilon_{BLR} = 10^{-5}$, at the cost of a much larger search space. However, the multipreconditioning does not improve the convergence for moderate and large compression.

The geometric-algebraic (G) and the low energy mode (E) provide similar and much better convergence rates. They differ only for the highest level of compression where the



(a) Total wall time.



(b) Number of iterations.

FIGURE 3 | Checkerboard cube, weak parallel scalability (homogeneous case $E_r/E_b = 10^0$): total time and number of iterations (the minimization space size is equal to the number of iterations for CG). Sator supercomputer. (a) Total wall time. (b) Number of iterations.

geometric–algebraic method performs better. The low energy modes probably drift away from the original operator’s nullspace for the highest compression. A degraded convergence is expected in this situation as shown by Dohrmann [14].


The BLR compression significantly reduces the memory footprint of the local preconditioner. The gain is about 40% for $\epsilon_{BLR} = 10^{-1}$, 27% for a moderate compression ($\epsilon_{BLR} = 10^{-3}$) and 20% for a small one ($\epsilon_{BLR} = 10^{-5}$). Interestingly, moderate and low compression improve both resolution time and memory footprint here (for both CG and AMPCG). Also, the geometric–algebraic method with high compression leads to the same total time than the uncompressed results while reducing the memory footprint of the preconditioner of 40%. Finally, the two BLR variants UCFS and UFSC lead to very similar results.

5.3.2.2 | Moderate Heterogeneity. The results obtained with $E_r/E_b = 10^2$ are summarized in Table 3. Again, the convergence is strongly degraded when Mumps does not detect the correct nullspace. The (G) method performs better, especially for moderate and high compression. It is the only one that reaches


convergence with CG for a high BLR compression. The multi-preconditioning clearly improves the convergence and time to solution. However, both (M) and (E) do not reach convergence with a high compression. Regarding the difference between the two BLR variants, no clear trend can be identified. Finally, the memory gain provided by the compression seems not affected by the material heterogeneity.

5.3.2.3 | High Heterogeneity. The results obtained with $E_r/E_b = 10^4$ are summarized in Table 4. The system is ill-conditioned due to the high heterogeneity. Mumps does not compute the correct coarse space even without BLR compression and very few configurations with the CG converge in less than 500 iterations. In this case, only the (G) handling of rigid body motion and classical CG supports a small compression factor, resulting in a maximum time gain of 6%.


Multipreconditioning makes it possible to use larger compression ratio, but the performance is poor in terms of iterations and time, only memory consumption is improved but the gain for the storage of the preconditioner is reduced by the large number



(a) Total wall time.




(b) Number of iterations.




(c) Minimization space size.


FIGURE 4 | Checkerboard cube, weak parallel scalability (moderate heterogeneity $E_r/E_b = 10^2$): total time, number of iterations and minimization space size. Sator supercomputer. (a) Total wall time. (b) Number of iterations. (c) Minimization space size.



(a) Total wall time.




(b) Number of iterations.




(c) Minimization space size.

FIGURE 5 | Checkerboard cube, weak parallel scalability (high heterogeneity $E_r/E_b = 10^4$): total time, number of iterations and minimization space size. Sator supercomputer. (a) Total wall time. (b) Number of iterations. (c) Minimization space size.



(a) Total wall time.



(b) Number of iterations.

FIGURE 6 | Checkerboard cube, weak parallel scalability (homogeneous case $E_r/E_b = 10^0$): wall time, number of iterations and minimization space size. Topaze supercomputer. (a) Total wall time. (b) Number of iterations.

of search directions to be kept. The best results with MPCG are obtained with the (G) handling of rigid body motion and a small compression where both the memory footprint and the total time are improved.

5.3.3 | Weak Scalability Results

After analysing the previous results and in order to reduce the number of data, only the best configurations are shown in the following. Multipreconditioning is only considered for the moderate and high heterogeneity. The level of BLR compression is adapted to the heterogeneity of the material, the higher the heterogeneity, the lower the level of compression. Also, since the two variants UCFS and UFCS lead to very similar results, only UCFS is used in the following.

5.3.3.1 | Homogeneous Problem. The parallel performance of the homogeneous test case are shown in Figure 3 and Table A1. The trends observed in Section 5.3.2 are confirmed. Moderate and low BLR compression do not significantly penalize the rate of convergence. As in section 5.3.2, (G) with a moderate

compression ratio ($\epsilon_{BLR} = 10^{-3}$) converges faster, both in terms of total time and number of iterations. It sounds surprising, but somehow the compressed preconditioner works better than the classic one. With a high compression ratio $\epsilon_{BLR} = 10^{-1}$, the convergence rate of (G) is slowed down significantly. The purpose of a such a configuration is mainly to reduce the memory footprint of the preconditioner. Variants (E) slightly increase the solution time due to a higher number of iterations and/or due to the overhead caused by the partial factorization (as observed in a previous work [24]).

5.3.3.2 | Moderate Heterogeneity. For the moderate heterogeneity test case, five curves are considered: the CG solver without and with a low compression ratio, and the MPCG solver with a low compression ratio. Both (G) and (E) are considered when using compression. The parallel performance with $E_r/E_b = 10^2$ are shown in Figure 4 and Table A2. As before, a constant number of iterations is not expected due to the automatic domain decomposition. Also, the larger the problem, the larger is the condition number due to material heterogeneity. As expected, multipreconditioned solvers tend to be faster to






FIGURE 7 | Checkerboard cube, weak parallel scalability (moderate heterogeneity $E_r/E_b = 10^2$): wall time, number of iterations and minimization space size. Topaze supercomputer. (a) Total wall time. (b) Number of iterations. (c) Minimization space size.

converge thanks to an enlarged search space. The convergence rate of the CG solver is quite satisfactory and remains competitive in terms of time to solution. Both (G) and (E) give similar results in terms of number of iterations and search space


size. However, the time to solution is much shorter for the (G) method, a closer look at the internal timers suggests that the time spent in backward and forward substitutions is faster with this method.



(a) Total wall time.



(b) Number of iterations.



(c) Minimization space size.

FIGURE 8 | Checkerboard cube, weak parallel scalability (high heterogeneity $E_r/E_b = 10^4$): wall time, number of iterations and minimization space size. Topaze supercomputer. (a) Total wall time. (b) Number of iterations. (c) Minimization space size.

5.3.3.3 | High Heterogeneity. Only two curves are shown for the highly heterogeneous test case: MPCG solver with geometric–algebraic nullspace detection, without and with low compression ratio. The MPCG solver with (E) without compression leads to the same convergence as with (G). The MPCG with (E) and BLR compression does not converge in less than 500 iterations for the test case with 3,072 cores. The parallel performance with $E_r/E_b = 10^4$ are shown in Figure 5 and Table A3. For this ill-conditioned test case, BLR compression slightly degrades the convergence rate but the time to solution and the search space size remain similar.

5.4 | Weak Scalability Study on the Topaze Supercomputer

This section presents the scalability study carried out on the Topaze supercomputer. The main interest here is that the available memory per core is only 2 GB, which initially motivated the use of BLR compression. Also, the compute nodes use AMD processors and it is the first time that our implementation is benchmarked on such an architecture.

5.4.1 | Presentation of the Topaze Supercomputer

The Topaze supercomputer is managed by the French Computing Center for Research and Technology (CCRT, <http://www-ccrt.cea.fr>). It is made of 864 nodes, 2.45 GHz AMD Milan bi-socket with 64 cores per socket. With 864 compute nodes (111,592 cores) and a theoretical Peak performance of 4.34 PFlop/s, Topaze is ranked 238 in the TOP500 (list from Nov. 2023). One specificity of Topaze is that the RAM per core is only 2 GB which motivates the use of compression techniques. Compute nodes are connected through a EDR InfiniBand network in a pruned Fat-tree topology. The communication are handled with OpenMPI 4.1.4.

5.4.2 | Weak Scaling Results

5.4.2.1 | Homogeneous Problem. For the homogeneous test case and in order to reduce the number of simulations, only the CG solver is used with or without BLR compression. The weak scaling results are shown in Figure 6. Full results are summarized in Table B1. Whatever the solver is, the number of iterations slightly increases with the size of the problem due to the automatic subdomain decomposition. Again, the configuration with a moderate compression provides the best performance, both in terms of iterations and time to solution. For the largest test case with 24,576 cores and 790.12M dofs, the time to solution is about 300s which represents a gain of about 40%. Also, the configuration with high compression provides the same time to solution than the uncompressed one, despite a greater number of iterations.

5.4.2.2 | Moderate Heterogeneity. For the moderate heterogeneity test case, only three curves are considered: the CG solver with (G) nullspace without and with a low compression ratio, and the MPCG solver with a low compression ratio. The results are shown in Figure 7 and Table B2. Once again, the CG solver performs well with low BLR compression, the convergence

rate is the same as without compression and the time to solution is reduced. Due to the larger search space, the MPCG solver with low BLR compression gives the best convergence rate. However, the cost of orthogonalising this search space tends to dominate the computation time for large problems ($\geq 6,000$ cores).

5.4.2.3 | High Heterogeneity. The weak scaling results obtained with $E_r/E_b = 10^4$ are shown in Figure 8 and Table B3. Here only MPCG without compression is able to converge in less than 500 iterations for large problems. The test case with 24,756 cores ran out of memory. Multipreconditioning provides robustness at the cost of a large search space: the number of iterations is only doubled between 384 and 6,000 cores. For this type of problem, a restart of the MPCG solver should be implemented, in the same spirit as, for example, the GMRES-DR algorithm [27]. This is however out of the scope of the present study.

6 | Conclusion and Perspectives

In order to adapt to modern supercomputer designs where the available memory per core is constantly decreasing, this paper proposes to use block low-rank factorization methods to equip primal domain decomposition methods with low memory footprint preconditioner. The BLR compression makes it difficult for the Mumps solver to detect the correct kernel to use. The nullspace is often not detected and the BDD method falls back to the Neumann-Neumann method: scalability is lost. Two alternative strategies have been tested: the hybrid geometric–algebraic approach and the low energy modes. The former makes the coarse problem independent of BLR compression, but requires the knowledge of the nullspace in the case of a completely floating subdomain. The latter is fully algebraic and takes compression into account, but numerical results suggest that the hybrid geometric–algebraic approach is preferable whenever available. Indeed, low energy modes seem to drift away from the original operator’s nullspace for a high level of compression, which significantly degrades the convergence rate [14]. The BLR preconditioner has also been combined with adaptive multipreconditioning in order to increase the robustness of the solver with respect to material heterogeneity.

Weak scalability studies were presented using two supercomputers (Sator and Topaze) and three heterogeneity ratios. Numerical results show that BLR compression can improve both memory and solution time. It is especially interesting for reasonably well conditioned problems. For the largest homogeneous test case with 24,576 cores and 790.12M dofs, the time to solution is about 300s, which represents a 40% gain over the uncompressed preconditioner, while the memory footprint of the preconditioner is reduced by 20%.

The results also show that AMPBDD is robust with respect to material heterogeneity but generates a large search space. Unfortunately, multipreconditioning is unable to compensate for the loss of the correct coarse space in most situations. The largest ill-conditioned test case has approximately 200 million of unknowns and runs on 6,000 cores. Block low rank factorization is not sufficient here, and a GMRES-DR-style restart procedure will need to be investigated in the near future. However, this is the first large-scale evaluation of this solver. AMPBDD is particularly

useful for simulating crack propagation problems because the nullspace computation only plays at the preconditioner level. One prospect of this work is the extension of AMPBDD phase field fracture [28] to larger scale problems solved on low memory supercomputers.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Data will be made available on request for the benchmarks presented in Section 5. The data that support the findings of this study are available from the corresponding author upon reasonable request.

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Endnotes

¹ <http://www.zset-software.com/>

² <http://eigen.tuxfamily.org/>

References

1. C. Farhat and F. X. Roux, "The Dual Schur Complement Method With Well-Posed Local Neumann Problems," *Contemporary Mathematics* 157 (1994): 193, <https://doi.org/10.1137/0914047>.
2. P. Gosselet, D. Rixen, F. X. Roux, and N. Spillane, "Simultaneous FETI and Block FETI: Robust Domain Decomposition With Multiple Search Directions," *International Journal for Numerical Methods in Engineering* 104, no. 10 (2015): 905–927. nme.4946, <https://doi.org/10.1002/nme.4946>.
3. C. Bovet, A. Parret-Fréaud, N. Spillane, and P. Gosselet, "Adaptive Multipreconditioned FETI: Scalability Results and Robustness Assessment," *Computers and Structures* 193 (2017): 1–20, <https://doi.org/10.1016/j.compstruc.2017.07.010>.
4. C. Bovet, A. Parret-Fréaud, and P. Gosselet, "Two-Level Adaptation for Adaptive Multipreconditioned FETI," *Advances in Engineering Software* 152 (2021): 102952, <https://doi.org/10.1016/j.advengsoft.2020.102952>.
5. C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen, "FETI-DP: a Dual-Primal Unified FETI Method - Part I: a Faster Alternative to the Two-Level FETI Method," *International Journal for Numerical Methods in Engineering* 50, no. 7 (2001): 1523–1544, <https://doi.org/10.1002/nme.76>.
6. J. Mandel, "Balancing Domain Decomposition," *Communications in Numerical Methods in Engineering* 9, no. 3 (1993): 233, <https://doi.org/10.1002/cnm.1640090307>.
7. C. R. Dohrmann, "A Preconditioner for Substructuring Based on Constrained Energy Minimization," *SIAM Journal for Scientific Computing* 25 (2003): 246, <https://doi.org/10.1137/s1064827502412887>.
8. T. Mary, "Block Low-Rank Multifrontal Solvers: Complexity, Performance, and Scalability" (PhD diss, Université Paul Sabatier-Toulouse III, 2017).
9. N. J. Higham and T. Mary, "A New Preconditioner that Exploits Low-Rank Approximations to Factorization Error," *SIAM Journal on Scientific Computing* 41, no. 1 (2019): A59–A82, <https://doi.org/10.1137/18M1182802>.
10. J. H. Bramble, J. E. Pasciak, and A. T. Vassilev, "Analysis of Non-Overlapping Domain Decomposition Algorithms With Inexact Solves," *Mathematics of Computation* 67, no. 221 (1998): 1–19, <https://doi.org/10.1090/S0025-5718-98-00879-5>.
11. C. Börgers, "The Neumann-Dirichlet Domain Decomposition Method With Inexact Solvers on the Subdomains," *Numerische Mathematik* 55, no. 2 (1989): 123–136, <https://doi.org/10.1007/BF01406510>.
12. G. Haase, U. Langer, and A. Meyer, "The Approximate Dirichlet Domain Decomposition Method," *Part I: An Algebraic Approach. Computing* 47, no. 2 (1991): 137–151, <https://doi.org/10.1007/BF02253431>.
13. G. Haase, U. Langer, and A. Meyer, "The Approximate Dirichlet Domain Decomposition Method. Part II: Applications to 2nd-Order Elliptic B.V.P.s," *Computing* 47, no. 2 (1991): 153–167, <https://doi.org/10.1007/BF02253432>.
14. C. R. Dohrmann, "An Approximate BDDC Preconditioner," *Numerical Linear Algebra With Applications* 14, no. 2 (2007): 149–168, <https://doi.org/10.1002/nla.514>.
15. J. Li and O. B. Widlund, "On the use of Inexact Subdomain Solvers for BDDC Algorithms," *Computer Methods in Applied Mechanics and Engineering* 196, no. 8 (2007): 1415–1428, <https://doi.org/10.1016/j.cma.2006.03.011>.
16. J. W. Liu, "The Multifrontal Method for Sparse Matrix Solution: Theory and Practice," *SIAM Review* 34, no. 1 (1992): 82–109.
17. M. Bebendorf, *Hierarchical Matrices* (Berlin, Heidelberg: Springer, 2008).
18. N. Spillane, "An Adaptive Multipreconditioned Conjugate Gradient Algorithm," *SIAM Journal on Scientific Computing* 38, no. 3 (2016): A1896–A1918, <https://doi.org/10.1137/15M1028534>.
19. P. Gosselet and C. Rey, "Non-Overlapping Domain Decomposition Methods in Structural Mechanics," *Archives of Computational Methods in Engineering* 13, no. 4 (2006): 515–572.
20. D. J. Rixen and C. Farhat, "A Simple and Efficient Extension of a Class of Substructure Based Preconditioners to Heterogeneous Structural Mechanics Problems," *International Journal for Numerical Methods in Engineering* 44, no. 4 (1999): 489–516, [https://doi.org/10.1002/\(SICI\)1097-0207\(19990210\)44:4<489::AID-NME514>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0207(19990210)44:4<489::AID-NME514>3.0.CO;2-Z).
21. R. Bridson and C. Greif, "A Multipreconditioned Conjugate Gradient Algorithm," *SIAM Journal on Matrix Analysis and Applications* 27, no. 4 (2006): 1056–1068 (electronic), <https://doi.org/10.1137/040620047>.
22. M. C. Leistner, P. Gosselet, and D. J. Rixen, "Recycling of Solution Spaces in Multi-Preconditioned FETI Methods Applied to Structural Dynamics," *International Journal for Numerical Methods in Engineering* 116, no. 2 (2018): 141–160, <https://doi.org/10.1002/nme.5918>.
23. N. Spillane and D. J. Rixen, "Automatic Spectral Coarse spaces for Robust FETI and BDD Algorithms," *International Journal for Numerical Methods in Engineering* 95, no. 11 (2013): 953–990, <https://doi.org/10.1002/nme.4534>.
24. C. Bovet, "On the use of Graph Centralities to Compute Generalized Inverse of Singular Finite Element Operators: Applications to the Analysis of Floating Substructures," *International Journal for Numerical Methods in Engineering* 124, no. 9 (2022): 1933–1964, <https://doi.org/10.1002/nme.7193>.
25. C. Farhat and M. Géradin, "On the General Solution by a Direct Method of a Large Scale Singular System of Linear Equations: Application to the Analysis of Floating Structures," *International Journal for Numerical Methods in Engineering* 41, no. 4 (1998): 675–696, [https://doi.org/10.1002/\(SICI\)1097-0207\(19980228\)41:4<675::AID-NME305>3.0.CO;2-8](https://doi.org/10.1002/(SICI)1097-0207(19980228)41:4<675::AID-NME305>3.0.CO;2-8).
26. P. R. Amestoy, I. S. Duff, J. Koster, and J. Y. L'Excellent, "A Fully Asynchronous Multifrontal Solver Using Distributed Dynamic Scheduling," *SIAM Journal on Matrix Analysis and Applications* 23, no. 1 (2001): 15–41, <https://doi.org/10.1137/S0895479899358194>.
27. R. B. Morgan, "GMRES with Deflated Restarting," *SIAM Journal on Scientific Computing* 24, no. 1 (2002): 20–37, <https://doi.org/10.1137/S1064827599364659>.

Appendix A

Weak Scalability Results on the Sator Supercomputer

TABLE A1 | Checkerboard cube, weak parallel scalability (homogeneous case $E_r/E_b = 10^0$): total time and number of iterations. Sator supercomputer.

Heterogeneity $E_r/E_b = 10^0$						
Solver	BLR	ϵ_{BLR}	Kernel		#iter	
CG	UCFS	10^{-1}	G	98	126	138
CG	UFSC	10^{-1}	G	98	126	138
CG	UCFS	10^{-3}	G	61	72	73
CG	UFSC	10^{-3}	G	61	73	73
CG			M	65	96	97
CG	UCFS	10^{-3}	E	66	92	114
CG	UFSC	10^{-3}	E	66	97	112
CG	UCFS	10^{-5}	E	65	96	97
CG	UFSC	10^{-5}	E	65	96	97
Solver	BLR	ϵ_{BLR}	Kernel		Total time [s]	
CG	UCFS	10^{-1}	G	88.86	113.3	132.1
CG	UFSC	10^{-1}	G	88.23	113.8	133.2
CG	UCFS	10^{-3}	G	72.19	85.75	94.12
CG	UFSC	10^{-3}	G	72.2	86.06	95.24
CG			M	72.33	99.24	110.4
CG	UCFS	10^{-3}	E	75.38	98.37	121.7
CG	UFSC	10^{-3}	E	75.33	101.0	122.4
CG	UCFS	10^{-5}	E	77.32	103.6	114.3
CG	UFSC	10^{-5}	E	77.65	104.4	115.5
Solver	BLR	ϵ_{BLR}	Kernel		Search space size	
CG	UCFS	10^{-1}	G	98	126	138
CG	UFSC	10^{-1}	G	98	126	138
CG	UCFS	10^{-3}	G	61	72	73
CG	UFSC	10^{-3}	G	61	73	73
CG			M	65	96	97
CG	UCFS	10^{-3}	E	66	92	114
CG	UFSC	10^{-3}	E	66	97	112
CG	UCFS	10^{-5}	E	65	96	97
CG	UFSC	10^{-5}	E	65	96	97
Number of subdomains				64	216	512
Number of cores				384	1296	3072

TABLE A2 | Checkerboard cube, weak parallel scalability (moderate heterogeneity $E_r/E_b = 10^2$): total time, number of iterations and minimization space size. Sator supercomputer.

Heterogeneity $E_r/E_b = 10^2$						
Solver	BLR	ϵ_{BLR}	Kernel		#iter	
CG			G	142	177	197
CG	UCFS	10^{-5}	G	142	177	197
CG	UFSC	10^{-5}	G	142	177	197
CG			E	142	177	197
CG	UCFS	10^{-5}	E	142	181	204
CG	UFSC	10^{-5}	E	142	179	206
MPCG	UCFS	10^{-3}	G	138	283	468
MPCG	UFSC	10^{-3}	G	147	280	462
MPCG	UCFS	10^{-5}	G	106	116	131
MPCG	UFSC	10^{-5}	G	107	116	133
MPCG			E	102	120	130
MPCG	UFSC	10^{-3}	E	216	334	469
MPCG	UCFS	10^{-5}	E	106	120	128
MPCG	UFSC	10^{-5}	E	106	120	133
Solver	BLR	ϵ_{BLR}	Kernel		Total time [s]	
CG			G	124.8	160.3	189.5
CG	UCFS	10^{-5}	G	119.7	153.3	180.7
CG	UFSC	10^{-5}	G	120.5	153.2	181.7
CG			E	126.5	160.4	189.3
CG	UCFS	10^{-5}	E	122.4	158.3	252.2
CG	UFSC	10^{-5}	E	122.7	156.0	189.3
MPCG	UCFS	10^{-3}	G	124.8	253.7	488.2
MPCG	UFSC	10^{-3}	G	129.0	265.0	482.7
MPCG	UCFS	10^{-5}	G	107.9	143.5	169.2
MPCG	UFSC	10^{-5}	G	109.9	142.8	167.5
MPCG			E	111.6	147.3	269.5
MPCG	UFSC	10^{-3}	E	175.3	308.8	636.6
MPCG	UCFS	10^{-5}	E	109.3	145.7	231.7
MPCG	UFSC	10^{-5}	E	109.7	144.6	229.9
Solver	BLR	ϵ_{BLR}	Kernel		Search space size	
CG			G	142	177	197
CG	UCFS	10^{-5}	G	142	177	197
CG	UFSC	10^{-5}	G	142	177	197
CG			E	142	177	197
CG	UCFS	10^{-5}	E	142	181	204
CG	UFSC	10^{-5}	E	142	179	206
MPCG	UCFS	10^{-3}	G	324	562	1057
MPCG	UFSC	10^{-3}	G	333	652	1051
MPCG	UCFS	10^{-5}	G	261	426	596
MPCG	UFSC	10^{-5}	G	262	426	567
MPCG			E	288	399	595
MPCG	UFSC	10^{-3}	E	433	706	1151
MPCG	UCFS	10^{-5}	E	261	399	624
MPCG	UFSC	10^{-5}	E	261	399	567
Number of subdomains				64	216	512
Number of cores				384	1296	3072

TABLE A3 | Checkerboard cube, weak parallel scalability (high heterogeneity $E_r/E_b = 10^4$): total time, number of iterations and minimization space size. Sator supercomputer.

Heterogeneity $E_r/E_b = 10^4$						
Solver	BLR	ϵ_{BLR}	Kernel		#iter	
MPCG			G	108	139	194
MPCG	UCFS	10^{-5}	G	108	141	206
MPCG	UFSC	10^{-5}	G	108	153	202
MPCG			E	108	139	194
Solver	BLR	ϵ_{BLR}	Kernel		Total time [s]	
MPCG			G	140.6	260.4	450.4
MPCG	UCFS	10^{-5}	G	139.5	262.8	462.4
MPCG	UFSC	10^{-5}	G	139.7	272.1	449.1
MPCG			E	142.8	261.2	448.8
Solver	BLR	ϵ_{BLR}	Kernel		Search space size	
MPCG			G	851	1564	3759
MPCG	UCFS	10^{-5}	G	895	1620	3771
MPCG	UFSC	10^{-5}	G	922	1610	3612
MPCG			E	851	1564	3759
Number of subdomains				64	216	512
Number of cores				384	1296	3072

Appendix B

Weak Scalability Results on the Topaze Supercomputer

TABLE B1 | Checkerboard cube, weak parallel scalability (homogeneous case $E_r/E_b = 10^0$): total time and number of iterations. Topaze supercomputer.

Heterogeneity $E_r/E_b = 10^0$						
Solver	BLR	ϵ_{BLR}	Kernel		#iter	
CG			G	65	96	97
CG	UCFS	10^{-1}	G	98	126	138
CG	UFSC	10^{-1}	G	98	126	138
CG	UCFS	10^{-3}	G	62	72	74
CG	UFSC	10^{-3}	G	61	73	74
CG			M	65	96	97
Solver	BLR	ϵ_{BLR}	Kernel		Total time [s]	
CG			G	87.5	168.8	261.9
CG	UCFS	10^{-1}	G	89.34	156.3	254.4
CG	UFSC	10^{-1}	G	90.37	153.1	260.8
CG	UCFS	10^{-3}	G	73.5	128.5	184.9
CG	UFSC	10^{-3}	G	74.19	127.9	177.7
CG			M	84.58	173.1	255.4
Solver	BLR	ϵ_{BLR}	Kernel		Search space size	
CG			G	65	96	97
CG	UCFS	10^{-1}	G	98	126	138
CG	UFSC	10^{-1}	G	98	126	138
CG	UCFS	10^{-3}	G	62	72	74
CG	UFSC	10^{-3}	G	61	73	74
CG			M	65	96	97
Number of subdomains				64	216	512
Number of cores				384	1296	3072

TABLE B2 | Checkerboard cube, weak parallel scalability (moderate heterogeneity $E_r/E_b = 10^2$): total time, number of iterations and minimization space size. Topaze supercomputer.

Heterogeneity $E_r/E_b = 10^2$								
Solver	BLR	ϵ_{BLR}	Kernel	#iter				
CG			G	141	182	210	238	304
CG	UCFS	10^{-5}	G	141	182	211	239	306
CG	UFSC	10^{-5}	G	141	182	211	238	306
MPCG	UCFS	10^{-5}	G	95	117	130	135	154
MPCG	UFSC	10^{-5}	G	95	114	130	139	154
Solver	BLR	ϵ_{BLR}	Kernel	Total time [s]				
CG			G	204.8	251.3	437.2	505.1	867.6
CG	UCFS	10^{-5}	G	169.2	225.2	385.8	430.8	740.8
CG	UFSC	10^{-5}	G	162.7	217.1	392.2	458.8	761.1
MPCG	UCFS	10^{-5}	G	161.3	221.9	353.2	402.1	945.7
MPCG	UFSC	10^{-5}	G	164.4	220.5	355.5	420.9	887.4
Solver	BLR	ϵ_{BLR}	Kernel	Search space size				
CG			G	141	182	210	238	304
CG	UCFS	10^{-5}	G	141	182	211	239	306
CG	UFSC	10^{-5}	G	141	182	211	238	306
MPCG	UCFS	10^{-5}	G	311	458	657	941	2572
MPCG	UFSC	10^{-5}	G	311	455	657	943	2541
Number of subdomains				64	216	512	1000	4096
Number of cores				384	1296	3072	6000	24576

TABLE B3 | Checkerboard cube, weak parallel scalability (high heterogeneity $E_r/E_b = 10^4$): total time, number of iterations and minimization space size. Topaze supercomputer.

Heterogeneity $E_r/E_b = 10^4$								
Solver	BLR	ϵ_{BLR}	Kernel	#iter				
MPCG			G	108	139	194	225	
MPCG			E	108	139	194	226	
Solver	BLR	ϵ_{BLR}	Kernel	Total time [s]				
MPCG			G	225.3	497.9	900.2	1299.0	
MPCG			E	236.1	501.9	850.7	1233.0	
Solver	BLR	ϵ_{BLR}	Kernel	Search space size				
MPCG			G	851	1564	3759	4906	
MPCG			E	851	1564	3759	4907	
Number of subdomains				64	216	512	1000	4096
Number of cores				384	1296	3072	6000	24576