DISTRIBUTED-MEMORY PARALLELIZATION OF THE AGGREGATED UNFITTED FINITE ELEMENT METHOD

Francesc Verdugo^{a,*}, Alberto F. Martín^{a,b}, and Santiago Badia^{a,b}

^aCIMNE – Centre Internacional de Mètodes Numèrics en Enginyeria
Esteve Terradas 5, 08860 Castelldefels, Spain.

^bDepartment of Civil and Environmental Engineering, Universitat Politècnica de Catalunya
Jordi Girona 1-3, Edifici C1, 08034 Barcelona, Spain.

Abstract

The aggregated unfitted finite element method (AgFEM) is a methodology recently introduced in order to address conditioning and stability problems associated with embedded, unfitted, or extended finite element methods. The method is based on removal of basis functions associated with badly cut cells by introducing carefully designed constraints, which results in well-posed systems of linear algebraic equations, while preserving the optimal approximation order of the underlying finite element spaces. The specific goal of this work is to present the implementation and performance of the method on distributed-memory platforms aiming at the efficient solution of large-scale problems. In particular, we show that, by considering AgFEM, the resulting systems of linear algebraic equations can be effectively solved using standard algebraic multigrid preconditioners. This is in contrast with previous works that consider highly customized preconditioners in order to allow one the usage of iterative solvers in combination with unfitted techniques. Another novelty with respect to the methods available in the literature is the problem sizes that can be handled with the proposed approach. While most of previous references discussing linear solvers for unfitted methods are based on serial non-scalable algorithms, we propose a parallel distributed-memory method able to efficiently solve problems at large scales. This is demonstrated by means of a weak scaling test defined on complex 3D domains up to 300M degrees of freedom and one billion cells on 16K CPU cores in the Marenostrum-IV platform. The parallel implementation of the AgFEM method is available in the large-scale finite element package FEMPAR.

Keywords: Unfitted finite element methods · Algebraic multigrid · High performance scientific computing

1 INTRODUCTION

Unfitted finite element (FE) methods (also known as embedded or immersed FE methods) are useful techniques in order to perform numerical computations associated with geometrically complex domains. They are specially useful in multi-phase and multi-physics applications with moving interfaces (e.g., fracture mechanics [1], fluid–structure interaction [2], free surface flows [3]), in applications with varying domain topologies (e.g., shape or topology optimization [4], 3D printing simulations [5]), in applications where the geometry is not described by CAD data (e.g., medical simulations based on CT-scan images [6]), or in large-scale parallel computations, where generating and partitioning large unstructured meshes is particularly difficult. The main benefit of unfitted FE methods is that they do not require the generation of *body-fitted* meshes (i.e., meshes whose faces conform to the domain boundary). Instead, they embed the domain of interest in a geometrically simple background grid (usually a uniform or an adaptive Cartesian grid), which can be generated much more efficiently. The popularity of this approach is illustrated by the large number of different methodologies that have been proposed following this rationale (see, e.g., the cutFEM method [7], the cutIGA method [8], the

*Corresponding author.

Emails: fverdugo@cimne.upc.edu (FV), amartin@cimne.upc.edu (AM), sbadia@cimne.upc.edu (SB)

August 8, 2019

Finite Cell Method [9], the AgFEM method [10], the cg-FEM method [11], the Immersed Boundary Method [12], the i-spline method [13], and some variants of the XFEM method [14]).

Unfortunately, unfitted FE methods have also well known drawbacks. One of the most notorious ones is the so-called *small cut cell problem*. The intersection of a background cell with the physical domain can be arbitrarily small and with arbitrarily high aspect ratios, which usually results in severely ill-conditioned systems of algebraic linear equations if no specific strategy is used to remedy it [15]. This flaw makes the solution of the underlying linear systems much more challenging than for standard FE methods based on body-fitted grids and it is still today one of the main limiting factors for the successful application of unfitted FE techniques in realistic large-scale applications. Sparse direct solvers [16] are usually considered when dealing with unfitted FE techniques (see, e.g., [9]) since they are robust enough to deal (up to a certain extent) with such ill-conditioned problems. However, its usage is prohibitive at large scales since their memory footprint and the algorithmic complexity scales supra-linearly with respect to the problem size. In FE analysis, the current way to effectively solve linear systems at large-scales is using iterative Krylov sub-space methods [17] in combination with parallel and scalable preconditioners. Unfortunately, the well known scalable preconditioners based on multigrid [18, 19] (either geometric [20] or algebraic [21, 22]), or multi-level domain decomposition [23, 24] are mainly designed for body-fitted meshes and cannot readily deal with the ill-conditioning associated with the small cut cell problem.

Specific preconditioners for unfitted methods have been proposed in the literature in different contexts. These techniques are generally based on the idea of splitting the system matrix into blocks associated with problematic degrees of freedom (DOFs), i.e., the ones affected by small cuts, and unproblematic ones, i.e., the ones not affected (see, e.g., [25, 15, 26, 27, 28]). The methods differ in the particular techniques used to handle the block associated with non-problematic DOFs, the block of problematic DOFs, and the coupling between these two blocks. Even though these methodologies have been able to allow the usage of iterative solvers in combination with unfitted FE methods in particular settings, they are mainly serial non-scalable algorithms. Considered problems are typically academic examples that go up to few hundreds of thousands of DOFs. To the authors' best knowledge, the only exceptions are the works in [29, 27]. Reference [27] is based on a single level Additive Schwarz preconditioner and reports results for problems up to almost 1 million DOFs solved with a hybrid OpenMP+MPI implementation with up to 16 MPI tasks and 64 OpenMP threads per MPI task. Even though this method is a step forward with respect to most previous works, it is well known that single-level Additive Schwarz preconditioners do not scale up algorithmically, and thus the method is not suitable for larger scales. On the other hand, the work in [29] considers a two-level balancing domain decomposition preconditioner [30] and reports results up to almost 6 million DOFs.

The goal of this work is to scale up to much larger problem sizes by considering an alternative approach. Instead of using a tailored preconditioner able to deal with linear systems affected by the small cut cell problem, we consider an enhanced FE formulation so that the resulting method leads to system matrices that are not affected by small cuts. This opens the door to consider well known scalable algorithms for conventional FE analysis such as algebraic multigrid (AMG) [21, 22], for which there are highly scalable parallel implementations in renowned scientific computing packages such as TRILINOS [31] or PETSc [32, 33]. This makes the method easy to use in parallel runs since the development of tailored linear solvers is not required. To this end, we consider the recently introduced aggregated unfitted finite element method (AgFEM) [10]. Other enhanced unfitted formulations could be also taken into account (e.g., the CutFEM method [7]), but this task is not in the scope of the current work. AgFEM is based on removal of basis functions associated with badly cut cells by introducing carefully designed constraints and its formulation shares the good properties of body-fitted FE methods such as stability, condition number bounds, optimal convergence, and continuity with respect to data. The detailed mathematical analysis of the method is included in [10] for elliptic problems and in [34] for the Stokes equation. Even though the method is potentially applicable to parallel large-scale computations, previous works [10, 34] focus mainly in the description of the

methodology in a serial context and its mathematical analysis. The main contribution of the current work is this two-fold. On the one hand, we extend AgFEM to large-scale computations by presenting an efficient distributed-memory implementation, which is based on standard functionality provided by distributed-memory FE packages. In this case, we have implemented the method in the large-scale FE software package FEMPAR [35]. The algorithmic strategies proposed for AgFEM in distributed-memory environments can also be used for the parallelization of other cell aggregation techniques for discontinuous Galerkin (DG) methods [36, 37, 38]. It will also be essential in CutFEM [7], when using substructuring domain decomposition techniques to end up with well-posed subdomain problems. On the other hand, we show that, by considering the parallel AgFEM, the resulting systems of linear algebraic equations can be effectively solved at large scales using standard AMG preconditioners. Here, we have considered the well-known PETSc's native AMG framework, referred to as GAMG [39]. We have customized as less as possible the parameters of this solver in order to demonstrate that users can take it as it is given "out of the box" by the library. These properties are demonstrated with weak scaling tests up to 16K cores and up to nearly 300M DOFs in the Marenostrum-IV platform at the Barcelona Supercomputing Center [40] using the Poisson equation on complex 3D domains as a model problem.

The outline of the paper is as follows. In Sect. 2, we briefly present the main components of the AgFEM method in a serial context as a basis for the subsequent extension to distributed-memory platforms. In particular, we introduce the geometrical setup (Sects. 2.1 and 2.2), the FE spaces of the AgFEM method (Sect. 2.3), the associated shape function basis (Sect. 2.4), and the FE assembly (Sect. 2.5). In Sect. 3, we discuss the distributed-memory implementation. We detail the domain-decomposition setup (Sect. 3.1), the distributed FE spaces (Sect. 3.2), several important algorithmic components (Sects. 3.3, 3.4, 3.5, and 3.6), and the parallel FE assembly (Sect. 3.7). In Sect. 4, we report the numerical examples. We detail the HPC environment (Sect. 4.1), the considered model problem (Sect. 4.2), the considered linear solver (Sect. 4.3), the particular setup for the experiments (Sect. 4.4), and the results of the weak scaling test (Sects. 4.5 and 4.6). Finally, we close the paper with the conclusions in Sect. 5.

2. THE AGGREGATED UNFITTED FINITE ELEMENT METHOD

Register for free at <https://www.scipedia.com> to download the version without the watermark

2.1 Immersed boundary setup

Let $\Omega \subset \mathbb{R}^d$ be an open bounded polygonal domain with $d \in \{2, 3\}$ the number of spatial dimensions. Like in any other embedded boundary method, let us introduce an *artificial* domain Ω^{art} such that it has a simple geometry (e.g., a cuboid) and it includes the *physical* domain $\Omega \subseteq \Omega^{\text{art}}$ (see Fig. 1a). We denote by $\mathcal{T}^{\text{art}}$ a partition of Ω^{art} into *cells* with characteristic cell size h . We build $\mathcal{T}^{\text{art}}$ as a Cartesian grid of d -dimensional parallelepipeds, i.e., hexahedra for $d = 3$ and quadrilaterals for $d = 2$ (see Fig. 1b). We consider Cartesian grids since they can be efficiently generated and partitioned in parallel with the `p4est` library [41]. Note, however, that the following rationale holds also for unstructured background meshes and other cell types such as triangles and tetrahedra.

We introduce the usual classification of cells in $\mathcal{T}^{\text{art}}$ as *internal*, *external* or *cut*. Specifically, $T \in \mathcal{T}^{\text{art}}$ is an *internal cell* if $T \subseteq \Omega$; T is an *external cell* if $T \cap \Omega = \emptyset$; otherwise, T is a *cut cell* (see Fig. 1b). For the sake of simplicity and without any loss of generality, we represent the domain boundary $\partial\Omega$ as the zero level-set of a given scalar function ψ^{ls} , namely $\partial\Omega \doteq \{x \in \mathbb{R}^d : \psi^{\text{ls}}(x) = 0\}$. We define the physical domain as the set of points where the levelset function is negative, namely $\Omega \doteq \{x \in \mathbb{R}^d : \psi^{\text{ls}}(x) < 0\}$. In this setup, classifying the cells into internal, external or cut is straightforward by sampling the levelset function at cell nodes. We note that the problem geometry could be described by other means, e.g., using 3D CAD data, by providing techniques to perform the cell classification and to perform numerical integration in cut cells (see, e.g., [42]). In any case, the way the geometry is handled does not affect the following exposition. We represent the set of interior (resp., external and cut) cells with $\mathcal{T}^{\text{in}}$ (resp., $\mathcal{T}^{\text{out}}$ and $\mathcal{T}^{\text{cut}}$). Furthermore, we define the

Figure 1: Embedded boundary setup.

set of *active cells* as the union of the sets of interior and cut cells, namely $\mathcal{T}^{\text{act}} \doteq \mathcal{T}^{\text{in}} \cup \mathcal{T}^{\text{cut}}$. The domain associated with active cells, namely $\Omega^{\text{act}} \doteq (\bigcup_{T \in \mathcal{T}^{\text{act}}} T) \cup (\bigcup_{T, T' \in \mathcal{T}^{\text{act}}} \bar{T} \cap \bar{T}')$ is referred to as the *active domain*, which satisfies $\Omega \subseteq \Omega^{\text{act}} \subseteq \Omega^{\text{art}}$. For simplicity, we refer to the set of active cells $\mathcal{T}^{\text{act}}$ simply as $\mathcal{T}$ (i.e., we drop the label "act" in the notation). In the following developments, it is convenient to assign a unique global identifier (referred to as *global id*) to each cell in $\mathcal{T}$. We denote by T^k the cell with the global id k , for $k = 1, \dots, |\mathcal{T}|$, where $|\cdot|$ is the number of elements in a set. With this notation, we can write $\mathcal{T} = \{T^k\}_{k \in \mathcal{K}}$, where $\mathcal{K} \doteq \{1, \dots, |\mathcal{T}|\}$ is the index set of all global cell ids. We also introduce $\mathcal{K}^{\text{in}}$ (resp. $\mathcal{K}^{\text{cut}}$) the subset of $\mathcal{K}$ that contains the ids of all interior (resp. cut) cells, namely $\mathcal{K}^{\text{in}} \doteq \{k \in \mathcal{K} : T^k \in \mathcal{T}^{\text{in}}\}$ and $\mathcal{K}^{\text{cut}} \doteq \{k \in \mathcal{K} : T^k \in \mathcal{T}^{\text{cut}}\}$.

2.2 Cell aggregation

We briefly present here the (serial) cell aggregation strategy proposed in [10], which is the cornerstone of the AgFEM method. The extension of this cell aggregation method to parallel distributed-memory frameworks is discussed later in Sect. 3.4. The cell aggregation process consists in building a surjective map that for each cell in $\mathcal{T}$ returns an interior cell in $\mathcal{T}^{\text{in}}$. For convenience, we represent this map as a transformation $R: \mathcal{K} \rightarrow \mathcal{K}^{\text{in}}$ such that $R(k) \in \mathcal{K}^{\text{in}}$ corresponding to an interior cell. The transformation R is called the *root cell map* since, for a given cell $T^k \in \mathcal{T}$, the associated interior cell $T^{R(k)} \in \mathcal{T}^{\text{in}}$ is referred to as the *root cell* of T^k . The concept of root cell plays an important role in the construction of the so-called aggregated finite element (agFE) spaces, which will be discussed later in Sect 2.3.

The root cell map R induces a partition of the cells in $\mathcal{T}$ into so-called *cell aggregates*. For each interior cell $T^k \in \mathcal{T}^{\text{in}}$, we define a cell aggregate $\mathcal{A}_k$ as the set of cells that share T^k as their root cell. Formally, the aggregate associated with cell T^k is defined as $\mathcal{A}_k \doteq \{T^{k'}\}_{k' \in R^{-1}(k)}$, where the inverse of R is defined as $R^{-1}(k) \doteq \{k' \in \mathcal{K} : k = R(k')\}$. Clearly, the set of aggregates $\{\mathcal{A}_k\}_{k \in \mathcal{K}^{\text{in}}}$ is a partition of $\mathcal{T}$ since each cell in $\mathcal{T}$ belongs to one (and only one) aggregate.

As stated in [10], in order to construct AgFEM spaces, the map R has to be such that the resulting cell aggregates $\mathcal{A}_k$ satisfy the following two properties for each $k \in \mathcal{K}^{\text{in}}$: (a) the set $\mathcal{A}_k$ contains one and only one interior cell, namely $|\mathcal{A}_k \cap \mathcal{T}^{\text{in}}| = 1$; and (b) the (unique) interior cell $\tilde{T}_0 \in \mathcal{A}_k \cap \mathcal{T}^{\text{in}}$ can be reached from any cut cell $\tilde{T}_n \in \mathcal{A}_k \cap \mathcal{T}^{\text{cut}}$ in the aggregate following a cell path $\tilde{T}_0, \tilde{T}_1, \dots, \tilde{T}_n$, where all cells are in the aggregate, i.e., $\tilde{T}_i \in \mathcal{A}_k, i = 0, \dots, n$, and two consecutive cells, $\tilde{T}_{i-1}$ and $\tilde{T}_i$, share a non-external face (for $d = 3$) or a non-external edge (for $d = 2$) for any $i = 1, \dots, n$. More precisely, we say that $\tilde{T}_{i-1}$ and $\tilde{T}_i$ share a non-external face ($d = 3$) or a non-external edge ($d = 2$) if $\tilde{T}_{i-1} \cap \tilde{T}_i$ is a manifold in $\mathbb{R}^d$ with co-dimension 1 and $\tilde{T}_{i-1} \cap \tilde{T}_i \cap \Omega \neq \emptyset$. Note that sharing only a corner (or an edge for $d = 3$) is not enough. For simplicity, we will use the words "edge" and "face" as completely equivalent synonyms in the following for the two-dimensional case ($d = 2$).

In serial runs, a map R fulfilling these properties is built with the strategy described in Alg. 1, and illustrated in Fig. 2 for a simple 2D case. The notation $\mathcal{K}^{\text{nei}}(k)$ in line 6 of the algorithm denotes, for

a given cell T^k with id $k \in \mathcal{K}$, the set containing the ids of the cells that share a face with cell T^k . In the following, we often denote the cells with ids in $\mathcal{K}^{\text{nei}}(k)$ as the *face neighbors* of cell T^k . Note that the algorithm consists in 3 basic steps (see Fig. 2). (a) We initialize the value of R for interior cells by setting each interior cell as its own root cell, and we mark all interior cells as "touched". (b) We perform a loop in cut cells. For a given cut cell T^k , we define its root cell taking the root cell of one of the face neighbors of T^k that have been already marked as "touched" (if such a neighbor exists). (c) After the loop, we mark all cells that have been assigned with a root cell id as "touched". Then, steps (b) and (c) are repeated as many times as needed until all the cells are assigned with a root cell id.

In line 11 of the algorithm, it is needed to choose one among the (possibly) several face neighbors that have already been "touched". Several criteria can be considered for this purpose. In our implementation, we chose the neighbor having the closest root cell (in terms of the euclidean distance between cell barycenters). In case of a tie between two or more neighbors, we chose one arbitrarily (e.g., the neighbor with smaller cell id). During the algorithm, we also populate a map $N : \mathcal{K} \rightarrow \mathcal{K}$ that stores the ids of the chosen face neighbor. More precisely, for a given cut cell T^k with id $k \in \mathcal{K}^{\text{cut}}$, $N(k)$ returns the id of face neighbor selected in line 11, whereas, for an interior cell T^k with $k \in \mathcal{K}^{\text{in}}$, $N(k)$ returns its own cell id, namely $N(k) = k$. The map N plays an important role in the distributed-memory implementation. Note that the aggregation scheme can be easily applied to arbitrary spatial dimensions. It only requires to find face neighbors and a classification of cells as internal, external, or cut. These procedures are widely available in embedded domain FE codes.

Algorithm 1: Serial cell aggregation algorithm

```

1 Initialize  $R(k) \leftarrow -1$  and  $N(k) \leftarrow -1$  for all  $k \in \mathcal{K}$ 
2 Set  $R(k) \leftarrow k$  and  $N(k) \leftarrow k$  for all  $k \in \mathcal{K}^{\text{in}}$ 
3 Set as touched all cell ids  $k \in \mathcal{K}$  such that  $R(k) \neq -1$ 
4 for  $k \in \mathcal{K}^{\text{cut}}$  such that  $k$  has not been touched yet do
5    $\mathcal{K}^{\text{can}} \leftarrow \emptyset$ 
6   for  $k' \in \mathcal{K}^{\text{nei}}(k)$  such that cell id  $k'$  is already touched do
7     Let  $F$  be the face shared by cells  $T^k$  and  $T^{k'}$ 
8     if  $F \cap \mathcal{K}^{\text{in}} \neq \emptyset$  then
9       Add  $k'$  to the set of candidates  $\mathcal{K}^{\text{can}}$ 
10  if  $\mathcal{K}^{\text{can}} \neq \emptyset$  then
11    Choose an arbitrary cell id  $k' \in \mathcal{K}^{\text{can}}$  (e.g., the one with the closest root cell)
12     $R(k) \leftarrow R(k')$ 
13     $N(k) \leftarrow k'$ 
14 if  $R(k) = -1$  for some cell id  $k \in \mathcal{K}$  then
15   Go to line 3

```

2.3 Standard and aggregated finite element spaces

Using the root cell map R , the authors in [10] define agFE spaces, which are the core of the AgFEM method. The aim of this section is to present such agFE spaces in a serial context, for its later extension to the distributed-memory case. To this end, we first define the usual conforming FE space associated with the active part of the mesh, namely

$$\mathcal{V}^{\text{std}} \doteq \{v \in C^0(\Omega^{\text{act}}) : v|_T \in \mathcal{V}(T) \text{ for any } T \in \mathcal{T}\}.$$

We denote by $\mathcal{V}(T)$ a vector space of functions defined on a generic cell $T \in \mathcal{T}$. In the numerical examples, we consider hexahedral cells and tri-linear shape functions. That is, we define $\mathcal{V}(T) \doteq Q_q(T)$ as the space of polynomials that are of degree less or equal to q ($q = 1$ in the examples) with respect to each variable in $(x_1, \dots, x_d) \in T$. Even though d -dimensional simplices are not

Figure 2: Illustration of the cell aggregation scheme defined in Algorithm 1. A blue circle indicates cells for which the map R is already defined with a valid value. Blue arrows indicate the face neighbor that is selected in line 11 of the algorithm.

considered in the numerical experiments, they can be also used in the presented method. In this case, we would define $\mathcal{V}(T) \doteq \mathcal{P}_q(T)$ as the space of polynomials of order less or equal to q in the variables $(x_1, \dots, x_d) \in T$. In order to simplify notation, we define the elemental functional spaces $\mathcal{V}(T)$ in the physical cell $T \subset \Omega^{\text{act}}$ (even though in the code we use reference parametric spaces as usual). The space $\mathcal{V}^{\text{std}}$ is the *standard* functional space used in unfitted FE methods (see, e.g., [29, 15, 9]). It is well known that it leads to arbitrary ill conditioned systems of linear algebraic equations (if no extra technique is used to remedy it). This is the main motivation in order to introduce the agFE spaces (see, e.g., [10, 34]), since they lead to linear systems whose condition number is not affected by small cuts.

In order to introduce the agFE spaces we need some extra notation. For the sake of simplicity in the exposition and without loss of generality, we restrict ourselves to scalar-valued Lagrangian FEs.

The extension to vector-valued or tensor-valued Lagrangian elements is straightforward by applying the same approach component by component. The extension to other types of FEs (e.g., Nedgelec elements [43]) is also possible. It only requires to work with the specific DOFs of the element instead of the Lagrangian nodal values.

Let $\{x^{a,k}\}_{a \in \Lambda(T^k)}$ be the set containing the (physical) coordinates of the Lagrangian nodes associated with cell $T^k \in \mathcal{T}$. In particular, $x^{a,k} \in \overline{\Omega^{\text{act}}}$ denotes the coordinate vector of node with id a (in the local scope of cell T^k), and $\Lambda(T^k) \doteq \{1, \dots, |\mathcal{V}(T^k)|\}$ is the set containing all local node ids in T^k . Since (for scalar-valued Lagrangian spaces) each local node is associated with one DOF of the space $\mathcal{V}(T^k)$, we also refer to the local indices $a \in \Lambda(T^k)$ as the local DOF ids of cell T^k and $\Lambda(T^k)$ as the set of local DOF ids in T^k . We also need to introduce the set $\{\phi^{a,k}\}_{a \in \Lambda(T^k)}$ of Lagrangian shape functions in the (physical) cell T^k . Following an analogous notation as for the cell nodal coordinates, the symbol $\phi^{a,k}$ refers to the shape function associated with the local DOF id $a \in \Lambda(T^k)$ in cell T^k . Clearly, $\phi^{a,k}(x^{b,k}) = \delta_{ab}$ for any $a, b \in \Lambda(T^k)$, and $v^k = \sum_{a \in \Lambda(T^k)} v^{a,k} \phi^{a,k}$, with $v^{a,k} \doteq v^k(x^{a,k})$ being the cell nodal / DOF values for a given function $v^k \in \mathcal{V}(T^k)$.

On the other hand, we denote by $\mathcal{X} \doteq \{x^i\}_{i \in \mathcal{I}}$ and $\{\phi^i\}_{i \in \mathcal{I}}$ the set of *global* nodal coordinates and the set of *global* shape functions of space $\mathcal{V}^{\text{std}}$, resp., where x^i and ϕ^i are the coordinate vector and shape function of the mesh node with global id i and $\mathcal{I} \doteq \{1, \dots, |\mathcal{V}^{\text{std}}|\}$ is the set of global node ids / DOF ids associated with the standard FE space $\mathcal{V}^{\text{std}}$. Clearly $\phi^i(x^j) = \delta_{ij}$, for any $i, j \in \mathcal{I}$, and $v = \sum_{i \in \mathcal{I}} v^i \phi^i$, with $v^i \doteq v(x^i)$ being the global nodal / DOF values for given function $v \in \mathcal{V}^{\text{std}}$.

We also need to introduce, for each cell $T^k \in \mathcal{T}$, the usual transformation between the cell local DOF ids and the corresponding global DOF ids. To this end, we denote by $g^{a,k} \in \mathcal{I}$ the global DOF id associated with the local DOF id a in cell T^k . Using the local-to-global map described by the values

$g^{a,k}$, we define the set of global DOF ids touched by a cell T^k , i.e., $\mathcal{G}(k) \doteq \{g^{a,k}\}_{a \in \Lambda(T^k)}$. In addition, we introduce the notation $[\phi^i]_k$ representing the local shape function in cell T^k associated with the global one ϕ^i , namely $[\phi^i]_k \doteq \phi^{a,k}$ with $i = g^{a,k}$. Note that $[\phi^i]_k$ and ϕ^i coincide inside the cell T^k , but they are different elsewhere.

Finally, we introduce the following classification of the global DOF ids into *interior* and *exterior* DOF ids. The set of interior DOF ids is defined as $\mathcal{I}^{\text{in}} \doteq \{i \in \mathcal{I} : i \in \mathcal{G}(k) \text{ for some } k \in \mathcal{K}^{\text{in}}\}$, i.e., all global DOF ids touched by the interior cells. On the other hand, the set of exterior DOF ids is defined as $\mathcal{I}^{\text{out}} = \mathcal{I} \setminus \mathcal{I}^{\text{in}}$. Clearly, the ids in $\mathcal{I}^{\text{out}}$ are only touched by cut cells. In Fig. 3b, interior and exterior DOF ids are represented with blue dots and red crosses, resp. Let us also define a map $K^{\text{own}} : \mathcal{I}^{\text{out}} \rightarrow \mathcal{K}$ that takes a DOF id $i \in \mathcal{I}^{\text{out}}$ and returns the cell id of one of the (possibly) several cells that contain the DOF id i (see Fig. 3c). In our implementation, we define $K^{\text{own}}(i)$ as the smaller cell id $k \in \mathcal{K}$ such that $i \in \mathcal{G}(k)$. By composing the root cell map R with K^{own} , we obtain another map $O : \mathcal{I}^{\text{out}} \rightarrow \mathcal{K}^{\text{in}}$ that, for each exterior DOF id $i \in \mathcal{I}^{\text{out}}$, returns an interior cell id. This latter map is defined as $O(i) \doteq R(K^{\text{own}}(i))$. We denote cell $T^{O(i)}$ as the *owner cell* of the DOF id i , or the *root cell* associated with DOF id i , see Fig. 3c.

Figure 3: Illustration of the spaces $\mathcal{V}^{\text{std}}$ and $\mathcal{V}^{\text{agg}}$ for first order Lagrangian elements.

With these notations, the agFE space reads

$$\mathcal{V}^{\text{agg}} \doteq \{v \in \mathcal{V}^{\text{std}} : v^i = \sum_{j \in \mathcal{M}(i)} C_{ij} v^j \text{ for any } i \in \mathcal{I}^{\text{out}}\}, \quad (1)$$

where $C_{ij} \doteq [\phi^j]_{O(i)}(x^i)$ and $\mathcal{M}(i) \doteq \mathcal{G}(O(i))$. Note that the space $\mathcal{V}^{\text{agg}}$ is defined taking as starting point $\mathcal{V}^{\text{std}}$, and adding some judiciously defined constraints to all exterior DOFs (see Fig. 3a and 3b). Only interior DOFs are free in the space $\mathcal{V}^{\text{agg}}$. Thus, the dimension of $\mathcal{V}^{\text{agg}}$ is equal to $|\mathcal{I}^{\text{in}}|$. Moreover, $\mathcal{V}^{\text{agg}} \subseteq \mathcal{V}^{\text{std}}$ by construction. Each exterior DOF $i \in \mathcal{I}^{\text{out}}$ is constrained as a linear combination of the DOF ids in the set $\mathcal{M}(i)$. Consequently, $\mathcal{M}(i)$ denotes the *master DOF ids* of a given exterior DOF id $i \in \mathcal{I}^{\text{out}}$. Note that $\mathcal{M}(i)$ contains all global DOF ids in the owner cell of i . Thus, constraints in (1) state that the value of v at an outer node with id $i \in \mathcal{I}^{\text{out}}$ is computed as the extrapolation of the values of v at the nodes of the owner cell of i (e.g., in Fig. 3c, the value at the red cross is computed as an extrapolation of the values at the blue dots).

2.4 Shape functions

Before starting the discussion on the distributed-memory implementation, we detail how to assemble the system of linear algebraic equations associated with the agFE space $\mathcal{V}^{\text{agg}}$ in the serial case. To this end, we first need to introduce a shape function basis of $\mathcal{V}^{\text{agg}}$, which is the purpose of this section. Let

us define a restriction operator that, given a function $v \in \mathcal{V}^{\text{std}}$, returns a function $\tilde{v} \in \mathcal{V}^{\text{agg}}$ defined as:

$$\tilde{v} \doteq \sum_{j \in \mathcal{I}^{\text{in}}} v(x^j) \phi^j + \sum_{i \in \mathcal{I}^{\text{out}}} \sum_{j \in \mathcal{M}(i)} C_{ij} v(x^j) \phi^i. \quad (2)$$

Note that function $\tilde{v}$ is computed in terms of the interior DOF ids only and it fulfills the constraints in (1) by construction. Moreover, v and $\tilde{v}$ are different only in cut cells. The restriction does not affect the value of v in interior cells. A shape function basis $\{\tilde{\phi}^g\}_{g \in \mathcal{I}^{\text{in}}}$ of $\mathcal{V}^{\text{agg}}$ is obtained by applying restriction (2) to the shape functions of $\mathcal{V}^{\text{std}}$ associated with interior DOFs. More precisely, the shape function of $\mathcal{V}^{\text{agg}}$ associated with the DOF id $g \in \mathcal{I}^{\text{in}}$ is defined as

$$\tilde{\phi}^g \doteq \sum_{j \in \mathcal{I}^{\text{in}}} \phi^g(x^j) \phi^j + \sum_{i \in \mathcal{I}^{\text{out}}} \sum_{j \in \mathcal{M}(i)} C_{ij} \phi^g(x^j) \phi^i.$$

Since the global shape functions of $\mathcal{V}^{\text{std}}$ are such that $\phi^i(x^j) = \delta_{ij}$ for any $i, j \in \mathcal{I}$, previous formula simplifies to

$$\tilde{\phi}^g \doteq \phi^g + \sum_{i \in \mathcal{Z}(g)} C_{ig} \phi^i, \quad (3)$$

where we have introduced the set $\mathcal{Z}(g) \doteq \{i \in \mathcal{I}^{\text{out}} : g \in \mathcal{M}(i)\}$ referred to as the *slave DOF ids* of an interior DOF id $g \in \mathcal{I}^{\text{in}}$. Note that the functions $\tilde{\phi}^i$ satisfy $\tilde{\phi}^i(x^j) = \delta_{ij}$ for any $i, j \in \mathcal{I}^{\text{in}}$. Thus, $\{\tilde{\phi}^i\}_{i \in \mathcal{I}^{\text{in}}}$ is the shape function basis of $\mathcal{V}^{\text{agg}}$ associated with the set of nodes $\{x^j\}_{j \in \mathcal{I}^{\text{in}}}$. Note that the shape functions of the agFE are associated with the interior nodes of the mesh.

2.5 Finite element assembly

We present now the assembly operations needed to build the matrix $\mathbf{A}$ and vector $\mathbf{b}$ representing the discrete counterparts of a given bi-linear form $a(\cdot, \cdot)$ and a linear form $b(\cdot)$ with respect to the agFE space $\mathcal{V}^{\text{agg}}$. The strategy we have followed is related with the way linear constraints are imposed in the assembly process in reference [44].

Using the previously defined shape functions of $\mathcal{V}^{\text{agg}}$, the entries of the matrix $\mathbf{A}$ and the vector $\mathbf{b}$ are defined as $[\mathbf{A}]_{ij} \doteq a(\tilde{\phi}^i, \tilde{\phi}^j)$ and $[\mathbf{b}]_i \doteq b(\tilde{\phi}^i)$ for each $i, j \in \mathcal{I}^{\text{in}}$. As usual, the global matrix and vector are computed as the assembly of elemental matrices and vectors $[\mathbf{A}_k]_{ab} \doteq a_k(\phi^{a,k}, \phi^{b,k})$ and $[\mathbf{b}_k]_a \doteq b_k(\phi^{a,k})$, $a, b \in \Lambda(T^k)$, associated with a generic cell $T^k \in \mathcal{T}$. Here, $a_k(\cdot, \cdot)$ (resp. $b_k(\cdot)$) is the restriction of $a(\cdot, \cdot)$ (resp. $b(\cdot)$) to cell $T^k \in \mathcal{T}$.

For local DOF ids $a \in \Lambda(T^k)$ such that its corresponding global DOF id $g^{a,k} \in \mathcal{I}^{\text{in}}$ is interior, the assembly is standard, i.e., the value $[\mathbf{b}_k]_a$ is assembled into the global entry $[\mathbf{b}]_i$ such that $i = g^{a,k}$. However, for local ids such that $g^{a,k} \notin \mathcal{I}^{\text{in}}$, the assembly is slightly different. To clarify this, we use the definition of $\tilde{\phi}^g$ (see Eq. (3)) in order to get a more detailed expression for $[\mathbf{A}]_{ij}$ and $[\mathbf{b}]_i$, namely

$$[\mathbf{b}]_i = b(\phi^i) + \sum_{n \in \mathcal{Z}(i)} C_{ni} b(\phi^n), \quad \text{for } i \in \mathcal{I}^{\text{in}}, \text{ and} \quad (4)$$

$$[\mathbf{A}]_{ij} = a(\phi^i, \phi^j) + \sum_{n \in \mathcal{Z}(i)} C_{ni} a(\phi^n, \phi^j) + \sum_{m \in \mathcal{Z}(j)} C_{mj} a(\phi^i, \phi^m) + \sum_{n \in \mathcal{Z}(i)} \sum_{m \in \mathcal{Z}(j)} C_{ni} C_{mj} a(\phi^n, \phi^m),$$

for $i, j \in \mathcal{I}^{\text{in}}$. In addition, we rewrite formula (4) in terms of cell-local quantities:

$$[\mathbf{b}]_i = \sum_{k \in \mathcal{B}(i)} [\mathbf{b}_k]_a + \sum_{n \in \mathcal{Z}(i)} \sum_{k \in \mathcal{B}(n)} C_{ni} [\mathbf{b}_k]_b \quad \text{for } i \in \mathcal{I}^{\text{in}}, \quad (5)$$

with $i = g^{a,k}$, $n = g^{b,k}$, and $\mathcal{B}(i) \doteq \{k \in \mathcal{K} : i = g^{c,k} \text{ for some } c \in \Lambda(T^k)\}$ being the set containing the ids of cells sharing DOF id i . It is clear by observing formula (5) that the assembly of the entry

$[\mathbf{b}_k]_a$ associated with a local DOF id $a \in \Lambda(T^k)$ such that $g^{a,k} \in \mathcal{I}^{\text{in}}$ is done as usual, i.e., $[\mathbf{b}_k]_a$ is added to the corresponding global entry $[\mathbf{b}]_i$ with $i = g^{a,k}$. However, the assembly is different for a local DOF id $a \in \Lambda(T^k)$ such that $g^{a,k} \notin \mathcal{I}^{\text{in}}$. By observing the last term in Eq. (5), the local entry $[\mathbf{b}_k]_b$ contributes to the global entry $[\mathbf{b}]_i$ for all global DOF ids $i \in \mathcal{I}^{\text{in}}$ that satisfy $g^{b,k} \in \mathcal{Z}(i)$, which is equivalent to say that $i \in \mathcal{M}(g^{b,k})$. Thus, the value $[\mathbf{b}_k]_b$ scaled by a factor C_{ni} , with $n = g^{b,k}$, is added to all global entries $[\mathbf{b}]_i$ that satisfy $i \in \mathcal{M}(g^{b,k})$. That is, the local entries $[\mathbf{b}_k]_a$ associated with constrained DOF ids $g^{a,k} \notin \mathcal{I}^{\text{in}}$ contribute to the global entries corresponding to the master DOF ids in $\mathcal{M}(g^{a,k})$. This is a convenient way of imposing constraints, since it does not lead to saddle point problems as it would be the case if we would use the method of Lagrange multipliers.

The same rationale applies for the entries of the system matrix. In this case, the assembly operation to be performed depends on the two local DOF ids associated with $[\mathbf{A}_k]_{ab}$, leading to four different cases. A summary of all different assembly operations is given in Table 1 for all possible cases. The symbol $+=$ stands for an in-place sum, i.e., the value of the right is added into the value on the left. Note that the procedures needed to assemble in $\mathcal{V}^{\text{agg}}$ are closely related to the ones required in octree-based adaptive mesh refinement codes with non-conforming meshes. The imposition of continuity constraints at so called *hanging nodes* (see, e.g., [45]) leads to assembly operations equivalent to the ones presented in Table 1. Thus, agFE spaces can be easily incorporated in existing octree-based adaptive codes.

Case	Assembly operation
Right hand side vector $g^{a,k} \in \mathcal{I}^{\text{in}}$	$[\mathbf{b}]_i += [\mathbf{b}_k]_a, i = g^{a,k}$ (standard case)
$g^{a,k} \notin \mathcal{I}^{\text{in}}$	$[\mathbf{b}]_i += C_{ni} [\mathbf{b}_k]_a, n = g^{a,k}, \forall i \in \mathcal{M}(n)$
System matrix $g^{a,k} \in \mathcal{I}^{\text{in}}$ and $g^{b,k} \in \mathcal{I}^{\text{in}}$	$[\mathbf{A}]_{ij} += [\mathbf{A}_k]_{ab}, i = g^{a,k}, j = g^{b,k}$ (standard case)
$g^{a,k} \notin \mathcal{I}^{\text{in}}$ and $g^{b,k} \in \mathcal{I}^{\text{in}}$	$[\mathbf{A}]_{ij} += C_{ni} [\mathbf{A}_k]_{ab}, n = g^{a,k}, j = g^{b,k}, \forall i \in \mathcal{M}(n)$
$g^{a,k} \in \mathcal{I}^{\text{in}}$ and $g^{b,k} \notin \mathcal{I}^{\text{in}}$	$[\mathbf{A}]_{ij} += C_{mj} [\mathbf{A}_k]_{ab}, i = g^{a,k}, m = g^{b,k}, \forall j \in \mathcal{M}(m)$
$g^{a,k} \notin \mathcal{I}^{\text{in}}$ and $g^{b,k} \notin \mathcal{I}^{\text{in}}$	$[\mathbf{A}]_{ij} += C_{ni} C_{mj} [\mathbf{A}_k]_{ab}, n = g^{a,k}, m = g^{b,k}, \forall i \in \mathcal{M}(n), \forall j \in \mathcal{M}(m)$

Table 1: Summary of the assembly operations related with the agFE space $\mathcal{V}^{\text{agg}}$. These assembly operations are related with the methodology introduced in [44].

3 DISTRIBUTED-MEMORY IMPLEMENTATION

3.1 Domain decomposition setup

After the presentation of the AgFEM in a serial context, we discuss its implementation in a distributed-memory framework. We start by introducing some notation associated with the partition of the mesh into sub-domains, which is the basis of the distributed-memory implementation. Let $\mathcal{D}$ be a partition of Ω^{art} into sub-domains obtained by the union of cells in the background mesh $\mathcal{T}^{\text{art}}$, i.e., for each cell $T \in \mathcal{T}^{\text{art}}$, there is a sub-domain $D \in \mathcal{D}$ such that $T \subset D$ (see Fig. 4a). The generation of the sub-domain partition $\mathcal{D}$ is specially simple in the context of embedded boundary methods since Cartesian grids are allowed. Here, we consider the load-balancing mechanism available in p4est, which is based on space-filling curves [46] and has shown excellent scalability up to hundreds of thousands of cores. Note that our approach does not require graph-based partitioning techniques such as metis or parmetis [47].

For convenience, we consider a numbering of the sub-domains in $\mathcal{D}$, where D_s is the sub-domain with id s and $\mathcal{S} \doteq \{1, \dots, |\mathcal{D}|\}$ is the index set containing all sub-domain ids. With this notation the

Figure 4: Domain decomposition setup illustrated for two sub-domains. (A) Partition into sub-domains. (B–C) Local and ghost cells for each sub-domain. Local cells are represented with solid lines, whereas ghost cells are represented with dashed ones.

set of sub-domains $\mathcal{D}$ is expressed as $\mathcal{D} = \{D_s\}_{s \in \mathcal{S}}$. For a given $s \in \mathcal{S}$, we define the set of *local cells* $\mathcal{T}_s^L$ as the cells included in D_s , i.e., $\mathcal{T}_s^L \doteq \{T \in \mathcal{T}^{\text{art}} : T \subset D_s\}$. On the other hand, we define the set of *ghost cells* $\mathcal{T}_s^G$ of D_s as the off-subdomain cells that share at least one vertex, edge or face with D_s , namely $\mathcal{T}_s^G \doteq \{T \in (\mathcal{T}^{\text{art}} \setminus \mathcal{T}_s^L) : \bar{T} \cap \bar{D}_s \neq \emptyset\}$ (see Fig. 4b, and 4c). The union of local and ghost cells, $\mathcal{T}_s \doteq \mathcal{T}_s^L \cup \mathcal{T}_s^G$, is called the set of *locally relevant cells* in sub-domain D_s .

For the following developments, we need a *local* numbering for the locally relevant cells $\mathcal{T}_s$. That is, we write $\mathcal{T}_s = \{T_s^\ell\}_{\ell \in \mathcal{L}_s}$, where T_s^ℓ denotes the cell with local id ℓ in the scope of sub-domain D_s and $\mathcal{L}_s \doteq \{1, \dots, |\mathcal{T}_s|\}$ is the set of all local cell ids associated with D_s . In addition, we define $\mathcal{L}_s^L \doteq \{\ell \in \mathcal{L}_s : T_s^\ell \in \mathcal{T}_s^L\}$ and $\mathcal{L}_s^G \doteq \{\ell \in \mathcal{L}_s : T_s^\ell \in \mathcal{T}_s^G\}$, which contain the local cell ids of local and ghost cells resp. We also need to identify which is the sub-domain of each cell in terms of local cell ids. To this end, we denote as $S_s(\ell)$ the id of the sub-domain, where the cell with id $\ell \in \mathcal{L}_s$ is a local cell. Clearly, $S_s(\ell) = s$ for all local cell ids $\ell \in \mathcal{L}_s^L$ and $S_s(\ell) \neq s$ for ghost cell ids $\ell \in \mathcal{L}_s^G$.

Register for free at <https://www.scipedia.com> to download the version without the watermark

For simplicity and without any loss of generality, we will assume (as we have done in the serial case) that exterior cells have been removed since they do not play any role. In particular, it implies that $\mathcal{T}_s^L$, $\mathcal{T}_s^G$, and $\mathcal{T}_s$ contain only interior or cut cells. As already stated in Sect. 2.1, each cell in the active mesh $\mathcal{T}$ receives a global id $k \in \mathcal{K}$, which is unique across all cells in $\mathcal{T}$. Since $\mathcal{T}_s$ is a subset of $\mathcal{T}$, all cells in $\mathcal{T}_s$ have a global cell id. We denote by $K_s(\ell)$ the global cell id associated with a local cell id $\ell \in \mathcal{L}_s$. Conversely, we denote as $L_s(k)$ the local cell id associated with the global cell id $k \in \mathcal{K}_s$. Here, $\mathcal{K}_s \doteq \{k \in \mathcal{K} : T^k \in \mathcal{T}_s\}$ is the set of all global ids assigned to cells in $\mathcal{T}_s$. With this notation, we can write $\mathcal{T}_s = \{T^k\}_{k \in \mathcal{K}_s}$. Finally, we need to introduce $\mathcal{S}_s^{\text{nei}} \doteq \{s' \in \mathcal{S} \setminus \{s\} : \bar{D}_{s'} \cap \bar{D}_s \neq \emptyset\}$, i.e., the set containing the sub-domain ids of the *nearest neighbors* of D_s .

3.2 Distributed finite element spaces

The goal of this section is to express the agFE space $\mathcal{V}^{\text{agg}}$ in terms of sub-domain local quantities in the framework of the domain decomposition setup. By doing so, most of the notation needed in the subsequent text will be introduced in a natural way. For the sake of simplicity, we assume from now on (if not otherwise indicated) that we are located in the local scope of a generic sub-domain D_s with id $s \in \mathcal{S}$. Also for simplicity (and without any loss of generality), we assume that we deal with scalar-valued Lagrangian FE interpolations.

First, we define $\mathcal{V}_s^{\text{std}} \doteq \mathcal{V}^{\text{std}}|_{D_s}$ the restriction of the standard FE space $\mathcal{V}^{\text{std}}$ to the sub-domain D_s . Given a FE function $v \in \mathcal{V}^{\text{std}}$, we also introduce its sub-domain restriction: $v_s \doteq v|_{D_s}$. Note that v_s is characterized by a subset of the DOFs of v , which are the ones associated with the nodes of the mesh $\mathcal{T}_s^L$, i.e., the nodes in the set $\mathcal{X}_s \doteq \{x \in \mathcal{X} : x \in \bar{T} \text{ for some } T \in \mathcal{T}_s^L\}$. As we have done previously for cells, we introduce a sub-domain local numbering for nodes in $\mathcal{X}_s$, i.e., $\mathcal{X}_s = \{x_s^j\}_{j \in \mathcal{J}_s}$, where x_s^j

denotes the nodal coordinate of the mesh node with local id j in sub-domain D_s and $\mathcal{J}_s \doteq \{1, \dots, |X_s|\}$ is the set of all sub-domain local node ids (also referred as set of local DOF ids since there is a one-to-one correspondence between nodes and DOFs). Thus, each node in X_s has two different ids, a local one in the local mesh $\mathcal{T}_s$ and a global one in the entire mesh $\mathcal{T}$. Let $\mathcal{I}_s \doteq \{i \in \mathcal{I} : x^i \in X_s\}$ be the set of the global node ids (or global DOF ids) in sub-domain D_s and let $I_s : \mathcal{J}_s \rightarrow \mathcal{I}_s$ be the map that, for a given local node id j , returns the corresponding global id $i = I_s(j)$ such that $x^i = x_s^j$. With these notations, we can write a generic function $v_s \in \mathcal{V}_s^{\text{std}}$ as a linear combination of sub-domain local quantities, i.e., $v_s(x) = \sum_{j \in \mathcal{J}_s} v_s^j \phi_s^j(x)$, where $v_s^j \doteq v_s(x_s^j) = v^i = v(x^i)$ and $\phi_s^j \doteq \phi^i|_{D_s}$ with $i = I_s(j)$ for all $j \in \mathcal{J}_s$. Clearly, v_s^j and ϕ_s^j are the DOF and the shape function associated with the sub-domain local id j .

On the other hand, we introduce $\mathcal{V}_s^{\text{agg}} \doteq \mathcal{V}^{\text{agg}}|_{D_s}$ the restriction of the agFE space $\mathcal{V}^{\text{agg}}$ to sub-domain D_s . In order to characterize $\mathcal{V}_s^{\text{agg}}$ in terms of sub-domain local quantities, we need to introduce some further notation. First, let us define $\phi_{s,\ell}^a$ the shape function associated with a generic DOF id $a \in \Lambda(T_s^\ell)$ in cell T_s^ℓ , namely $\phi_{s,\ell}^a \doteq \phi^{a,k}$ with $k = K_s(\ell)$. Analogously, we introduce $x_{s,\ell}^a \doteq x^{a,k}$ and $g_{s,\ell}^a \doteq g^{a,k}$, with $k = K_s(\ell)$, which represent the nodal coordinates and the global DOF id associated with a generic DOF id $a \in \Lambda(T_s^\ell)$ in cell T_s^ℓ . Note that the symbols $g_{s,\ell}^a$ and $g^{a,k}$ express the same quantities, but the former expresses them in terms of local cell ids $\ell \in \mathcal{L}_s$, and the latter in terms of global cell ids $k \in \mathcal{K}_s$. This is true also for $x_{s,\ell}^a$ and $x^{a,k}$, and $\phi_{s,\ell}^a$ and $\phi^{a,k}$, resp. We also define $j_{s,\ell}^a$ as the sub-domain local DOF id in $\mathcal{J}_s$ associated with the cell local DOF id $a \in \Lambda(T_s^\ell)$ in cell T_s^ℓ . Formally, $j_{s,\ell}^a$ is expressed as $j_{s,\ell}^a \doteq I_s^{-1}(g_{s,\ell}^a)$ in terms of the inverse application I_s^{-1} (which is not needed to be generated in practice). The process that generates $g_{s,\ell}^a$ and $j_{s,\ell}^a$ for each a, ℓ , and s is standard in FE packages. The sub-domain local cell-wise DOF ids $j_{s,\ell}^a$ are simply generated by applying independently in each sub-domain the standard procedure for generating DOF ids in the serial case (see, e.g., [35] for details). On the other hand, the generation of the global DOF ids $g_{s,\ell}^a$ is also done by means of standard tools (see, e.g., [48, 45]) as it is detailed later at the end of Sect. 3.6.

We also need to introduce a classification of the locally relevant cells into interior, exterior, and cut. To this end, we define the index sets $\mathcal{L}_s^{\text{in}} \doteq \{\ell \in \mathcal{L}_s : T_s^\ell \in \mathcal{T}^{\text{in}}\}$ (resp. for $\mathcal{L}_s^{\text{out}}$ and $\mathcal{L}_s^{\text{cut}}$) containing the subset of $\mathcal{L}_s$ corresponding to interior cells (resp. exterior and cut). In our case, the set $\mathcal{L}_s^{\text{out}}$ is empty since we have assumed that exterior cells have been removed from the mesh. The classification of cells as interior or cut induces the following partition of sub-domain local DOF ids in $\mathcal{J}_s$. Let $\mathcal{J}_s^{\text{in}} \doteq \{j \in \mathcal{J}_s : j = j_{s,\ell}^a \text{ for some } a \in \Lambda(T_s^\ell) \text{ and } \ell \in \mathcal{L}_s^{\text{in}}\}$ be the set of *interior* sub-domain local DOF ids, and let $\mathcal{J}_s^{\text{out}} \doteq \mathcal{J}_s \setminus \mathcal{J}_s^{\text{in}}$ be the set of *exterior* sub-domain local DOF ids. Note that the DOF ids in $\mathcal{J}_s^{\text{in}}$ are the ones touched by interior cells with ids in $\mathcal{L}_s^{\text{in}}$, whereas the ids in $\mathcal{J}_s^{\text{out}}$ are only touched by cut cells.

The last ingredient to characterize $\mathcal{V}_s^{\text{agg}}$ in terms of sub-domain quantities is a distributed version of the root cell map R and the map K^{own} . On the one hand, the distributed root cell map $R_s : \mathcal{L}_s \rightarrow \mathcal{K}^{\text{in}}$ is a transformation that takes a local cell id $\ell \in \mathcal{L}_s$ and returns the global cell id $R_s(\ell) \in \mathcal{K}^{\text{in}}$ of the root cell of T_s^ℓ , namely $R_s(\ell) \doteq R(K_s(\ell))$. We encode the input value of $R_s(\cdot)$ with a local cell id in $\mathcal{L}_s$ since we only need to define R_s for locally relevant cells in $\mathcal{T}_s$. However, the returned value of R_s has to be encoded with a global cell id in $\mathcal{K}^{\text{in}}$, and not with a local cell id in $\mathcal{L}_s$, since the root cell of T_s^ℓ can be outside the set of locally relevant cells $\mathcal{T}_s$ (which are the only ones that have a local id in sub-domain D_s). Note also that, even though we have formally defined R_s using R , the map R is never built in practice. Constructing and storing R so that each processor can retrieve the value $R(k)$ for any global cell id $k \in \mathcal{K}$ in the entire mesh would kill parallel (memory) scalability. Instead, we only build the distributed map R_s as shown later in Sect. 3.4. In addition to R_s , we also construct another map $S_s^{\text{root}} : \mathcal{L}_s \rightarrow \mathcal{S}$ that for each local cell id $\ell \in \mathcal{L}_s$ returns the id $S_s^{\text{root}}(\ell) \in \mathcal{S}$ of the sub-domain in which the corresponding root cell is a local cell, namely $s' = S_s^{\text{root}}(\ell)$ is such that $T_{s'}^{\ell'} \in \mathcal{T}_{s'}^{\text{L}}$ with $\ell' = L_{s'}(R_s(\ell))$. Finally, we need to build a transformation $L_s^{\text{own}} : \mathcal{J}_s^{\text{out}} \rightarrow \mathcal{L}_s$, which is a distributed

version of K^{own} . For a local DOF id $j \in \mathcal{J}_s^{\text{out}}$, the value $L_s^{\text{own}}(j)$ is the local cell id $\ell \in \mathcal{L}_s$ of one of the (possibly several) cells T_s^ℓ such that $j = j_{s,\ell}^a$ for some $a \in \Lambda(T_s^\ell)$. That is, $L_s^{\text{own}}(j)$ is the local id of one of the (possibly) several cells that contain the sub-domain DOF id j . When more than one cell can be selected, we choose one arbitrarily, e.g., the cell that has smaller global id $k \in \mathcal{K}_s$. Note that one can choose one cell arbitrarily, but the chosen criterion has to be consistent across all sub-domains. E.g., choosing the cell with smaller global cell id is consistent, but choosing the cell with smaller local id is not.

With these notations, we can finally write $\mathcal{V}_s^{\text{agg}}$ in terms of sub-domain quantities:

$$\mathcal{V}_s^{\text{agg}} = \left\{ v_s \in \mathcal{V}_s^{\text{std}} : v_s^j = \sum_{a \in \Lambda(T_{s'}^{\ell'})} C_s^{ja} v_{s',\ell'}^a \text{ for all } j \in \mathcal{J}_s^{\text{out}} \right\}, \quad (6)$$

with $C_s^{ja} \doteq \phi_{s',\ell'}^a(x_s^j)$, $\ell' = L_{s'}(R_s(L_s^{\text{own}}(j)))$ and $s' = S_s^{\text{root}}(L_s^{\text{own}}(j))$. By observing Eq. (6), the main difficulty of the implementation of the method becomes apparent. Note that in the scope of the current sub-domain D_s , not all quantities are available in order to define the constraints in Eq. (6). In order to evaluate the coefficient C_s^{ja} , one needs access to the shape function in physical space $\phi_{s',\ell'}^a$, which is one of the local shape basis functions of the cell with local id $\ell' = L_{s'}(R_s(L_s^{\text{own}}(j)))$ in sub-domain $s' = S_s^{\text{root}}(L_s^{\text{own}}(j))$, which can be interpreted as the root cell associated with the DOF id j . Building up $\phi_{s',\ell'}^a$ requires the physical coordinates of the nodes of the root cell, namely $x_{s',\ell'}^b$ for $b \in \Lambda(T_{s'}^{\ell'})$. Since, in general, the sub-domain where the root cell is located is different from the current sub-domain, i.e., $s' \neq s$, the quantities $x_{s',\ell'}^b$ are not available in the scope of D_s (see Fig. 5). Moreover, one cannot directly use the communication pattern used to exchange cell-based data between nearest neighbor sub-domains because id s' is not even necessarily the id of a nearest neighbor sub-domain, namely $s' \notin S_s^{\text{nei}}$ in general. Thus, an ad-hoc communication strategy has to be built in order to import quantities like $x_{s',\ell'}^b$ into the scope of the current processor.

Figure 5: Illustration of a case where an exterior DOF is constrained by a root cell in a remote sub-domain. The value at x_s^j is constrained as the extrapolation of the nodal values of cell $T_{s'}^{\ell'}$. In order to set up such constraint it is necessary to send data from sub-domain $D_{s'}$ to sub-domain D_s , which are not nearest neighbor sub-domains.

3.3 Implementation overview

In the following sections (Sects. 3.4–3.7), we discuss the implementation of the main algorithmic phases of the distributed-memory AgFEM method. We start by presenting the distributed version of the cell aggregation method presented in Alg. 1 (see Sect. 3.4), where we show the actual computation of the distributed root cell map R_s and the auxiliary maps K_s^{own} and S_s^{root} . Then, in Sect. 3.5, we detail how to build the communication strategy in order to import data from root cells that are located in remote sub-domains. Finally, we detail the data distribution model used for representing the linear system of equations in the distributed-memory implementation (Sect. 3.6) and the parallel FE assembly associated with the agFE space (Sect. 3.7).

Algorithm 2: Parallel cell aggregation algorithm

```

1 Initialize  $R_s(\ell) \leftarrow -1$ ,  $S_s^{\text{root}}(\ell) \leftarrow -1$  and  $N_s(\ell) \leftarrow -1$  for all local cells  $\ell \in \mathcal{L}_s^L$ 
2 Set  $R_s(\ell) \leftarrow K_s(\ell)$ ,  $S_s^{\text{root}}(\ell) \leftarrow s$  and  $N_s(\ell) \leftarrow K_s(\ell)$  for all interior local cells  $\ell \in \mathcal{L}_s^L \cap \mathcal{L}_s^{\text{in}}$ 
3 Recover the values of  $R_s(\ell)$  and  $S_s^{\text{root}}(\ell)$  in ghost cells  $\ell \in \mathcal{L}_s^G$  with a nearest neighbor exchange
4 Set as touched all cell ids  $\ell \in \mathcal{L}_s$  such that  $R_s(\ell) \neq -1$ 
5 for  $\ell \in \mathcal{L}_s^L \cap \mathcal{L}_s^{\text{cut}}$  such that  $\ell$  has not been touched yet do
6    $\mathcal{L}_s^{\text{can}} \leftarrow \emptyset$ 
7   for  $\ell' \in \mathcal{L}_s^{\text{nei}}(\ell)$  such that cell id  $\ell'$  is already touched do
8     Let  $F$  be the face shared by cells  $T_s^\ell$  and  $T_s^{\ell'}$ 
9     if  $F \cap \Omega \neq \emptyset$  then
10       Add  $\ell'$  to the set of candidates  $\mathcal{L}_s^{\text{can}}$ 
11   if  $\mathcal{L}_s^{\text{can}} \neq \emptyset$  then
12     Choose an arbitrary cell id  $\ell' \in \mathcal{L}_s^{\text{can}}$  (e.g., the one with the closest root cell)
13      $R_s(\ell) \leftarrow R_s(\ell')$ 
14      $S_s^{\text{root}}(\ell) \leftarrow S_s^{\text{root}}(\ell')$ 
15      $N_s(\ell) \leftarrow K_s(\ell')$ 
16 if  $R_s(\ell) \neq -1$  for all cell ids  $\ell \in \mathcal{L}_s$  then
17    $b_s \leftarrow \text{true}$ 
18 else
19    $b_s \leftarrow \text{false}$ 
20 if  $b_{s'}$  is still false in some sub-domain  $s' \in \mathcal{S}$  then
21   Go to line 3 in all processors

```

3.4 Parallel cell aggregation

The goal of this section is to parallelize the cell aggregation method presented in Alg. 1. To this end, the most naive approach would be to run Alg. 1 on the local cells $\mathcal{T}_s^L$ in each processor independently, trying to greedily associate each cut cell with a root cell in the same sub-domain. However, this approach does not work when a sub-domain has cut cells that cannot be aggregated to any local cell (i.e., the cut cell does not share a face with any local internal cell), and therefore, it is not general enough. Moreover, even when the method is algorithmically applicable (e.g., for domain decomposition setups where each sub-domain has at least one interior cell), it does not necessarily lead to the same aggregates as the serial method in Alg. 1. The mathematical analysis of the AgFEM method in [10] states that the approximation properties of the agFE spaces are negatively affected as the aggregate size increases. Thus, this naive parallel strategy that can potentially lead to large aggregates has to be avoided.

Instead, we consider a more sophisticated parallel algorithm, which leads to the same aggregates as the ones obtained with the serial aggregation method. This enhanced distributed-memory cell aggregation method is detailed in Alg. 2 and illustrated in Fig. 6 for a simple two dimensional case. The notation $\mathcal{L}_s^{\text{nei}}(\ell)$, in line 7 of the algorithm represents the set of local cell ids of the face neighbors of cell T_s^ℓ . That is, the notations $\mathcal{K}^{\text{neig}}(k)$ (see Sect. 2.2) and $\mathcal{L}_s^{\text{nei}}(\ell)$ both represent the face neighbors of a given cell, but the former in terms of global cell ids and the latter in terms of local cell ids. The parallel cell aggregation method in Alg. 2 follows essentially the same steps as the serial counterpart plus some communications between neighbor sub-domains, which allow one to aggregate a cell into a root belonging to a different sub-domain. The method consists in three main steps (see Fig. 6). (a) We initialize the value of R_s for all interior local cells and perform a nearest neighbor communication in order to determine the value of R_s at ghost cells. At this point, we mark all interior (local and ghost) cells as "touched". (b) As in the serial case, we perform a loop in (local) cut cells. For a given cut cell, we define its root cell taking the root cell of one of its face neighbors that have been already

Figure 6: Illustration of the cell aggregation scheme defined in Algorithm 2 for the two sub-domains presented in Fig. 4. A blue circle indicates cells for which the map R_s is already defined with a valid value. Blue arrows indicate the face neighbor that is selected in line 12 of the algorithm.

marked as "touched" (if such a neighbor exists). (c) After the loop, we perform a nearest neighbor communication in order to update the value of R_s at ghost cells. At this point, we mark as "touched" all (local and ghost) cells for which R_s is already set with a valid root cell id. Finally, we repeat steps (b) and (c) in all processors as many times as needed until all the cells are assigned with a root cell id. In the numerical experiments, 3 repetitions of steps (b) and (c) were needed to aggregate all cut cells for all cases.

Note that, during the execution of Alg. 2, we also build the map $S_s^{\text{root}} : \mathcal{L}_s \rightarrow \mathcal{S}$, which was formally presented in Sect. 3.2. In addition, we build an auxiliary map $N_s : \mathcal{L}_s \rightarrow \mathcal{K}$, which is the distributed version of the map $N : \mathcal{K} \rightarrow \mathcal{K}$ presented in Sect. 2.2. Precisely, $N_s(\ell)$ returns the global id of the cell that was selected in line 12 of the algorithm, i.e., the next cell in the path towards the root cell. This information will be required later in the next section (Sect. 3.5) in order to build the communication strategy for importing data from remote root cells.

3.5 Importing data from root cells

As we have discussed in Sect. 3.2, we need to import data associated with remote root cells in order to build the AgFEM constraints in Eq. (6). Formally, the set containing the global cell ids of those root cells is defined as $\mathcal{K}_s^{\text{R}} \doteq \{k \in \mathcal{K} \setminus \mathcal{K}_s : k = R_s(\ell) \text{ for some } \ell \in \mathcal{L}_s\}$. For assembling the system matrix and right hand side vector, we need to import the cell-wise nodal coordinates and global DOF ids for all cells with global ids in $\mathcal{K}_s^{\text{R}}$ as it will be justified in Sect. 3.7. This information is fetched from the corresponding remote sub-domains (as we discuss later) and stored in the local scope of the processor in a separate buffer allocation. In order to address data in this buffer, each cell with global id k in the set $\mathcal{K}_s^{\text{R}}$ is labeled with a new sub-domain local id z , defined as the position of k in the list that represents the index set $\mathcal{K}_s^{\text{R}}$. This transformation is represented here by a map $Z_s : \mathcal{K}_s^{\text{R}} \rightarrow \mathcal{L}_s^{\text{R}} \doteq \{1, \dots, |\mathcal{K}_s^{\text{R}}|\}$ such that, given a cell global id $k \in \mathcal{K}_s^{\text{R}}$, it returns the corresponding z id. We denote by $\tilde{T}_s^z$ the cell

Algorithm 3: Parallel direct path reconstruction

```

1  $Q_s^{\text{rcv}} \leftarrow \emptyset$ 
2  $\mathcal{R}_{s,s'}^{\text{rcv}} \leftarrow \emptyset$  for all  $s' \in \mathcal{S}$ 
3 for  $\ell \in \mathcal{L}_s^{\text{cut}}$  do
4    $s' \leftarrow S_s^{\text{root}}(\ell)$ 
5   if  $s' \neq s$  then
6     Add  $s'$  to  $Q_s^{\text{rcv}}$ 
7     Add  $R_s(\ell)$  to  $\mathcal{R}_{s,s'}^{\text{rcv}}$ 

```

associated with label $z \in \mathcal{L}_s^R$, namely $\tilde{T}_s^z \doteq T^k$ with $z = Z_s(k)$. In addition, we denote as $\tilde{x}_{s,z}^a \doteq x^{a,k}$, $z = Z_s(k)$, the coordinates of the a -th node of cell $\tilde{T}_s^z$ for each $a \in \Lambda(\tilde{T}_s^z)$. We use an analogous notation for $\tilde{g}_{s,z}^a \doteq g^{a,k}$, $\tilde{\phi}_{s,z}^a \doteq \phi^{a,k}$, and $\tilde{v}_{s,z}^a \doteq v^{a,k}$ (i.e., cell-wise global DOF ids, cell shape functions, and cell DOF values of a given FE function $v \in \mathcal{V}^{\text{agg}}$).

Note that the cells with global id in $\mathcal{K}_s^R$ are local cells in the sub-domains with id in the set $Q_s^{\text{rcv}} \doteq \{s' \in \mathcal{S} : s' = S_s^{\text{root}}(\ell) \text{ for some } \ell \in \mathcal{L}_s^{\text{cut}}\}$. That is, we import data from all sub-domains that contain the root cells of the cut cells in the current sub-domain. We note that the root cells of *ghost* cut cells also need to be imported as $L_s^{\text{own}}(j)$ may be in $\mathcal{L}_s^G$ for some $j \in \mathcal{J}_s^{\text{out}}$. We also note that we have constructed the map S_s^{root} during the parallel cell aggregation algorithm in order to be able to construct the set Q_s^{rcv} . It is also easy to see that the set $\mathcal{R}_{s,s'}^{\text{rcv}} \doteq \{k \in \mathcal{K}_{s'} : k = R_s(\ell) \text{ for some } \ell \in \mathcal{L}_s^{\text{cut}}\}$ contains, for a given $s' \in Q_s^{\text{rcv}}$, the global ids of root cells in sub-domain $D_{s'}$ from which we (potentially) need to recover data.

In order to grasp the sets Q_s^{rcv} and $\mathcal{R}_{s,s'}^{\text{rcv}}$, we can consider the cell path that goes from each cut cell to its corresponding root cell, jumping along the face neighbors that are selected during the cell aggregation algorithm (see the blue arrows in Fig 6d). The set Q_s^{rcv} contains the ids of the sub-domains, where the cell paths starting in the current sub-domain D_s end. On the other hand, the set $\mathcal{R}_{s,s'}^{\text{rcv}}$ contains the ids of the last cell in each path that starts at D_s and ends at $D_{s'}$. The practical construction of the sets Q_s^{rcv} and $\mathcal{R}_{s,s'}^{\text{rcv}}$ from the maps R_s and S_s^{root} is indicated in Alg. 3. Note that this procedure is done at each sub-domain independently and does not require communication at all.

Once the sets Q_s^{rcv} and $\mathcal{R}_{s,s'}^{\text{rcv}}$ are built, we know from which sub-domains, and from which cells in those sub-domains, we need to recover data. However, we still need to know to which sub-domains we need to send data from the current sub-domain, and which is the data to be sent to each of those sub-domains. The "send" counterparts of the sets Q_s^{rcv} and $\mathcal{R}_{s,s'}^{\text{rcv}}$ are formally defined as $Q_s^{\text{snd}} \doteq \{s' \in \mathcal{S} : s = S_{s'}^{\text{root}}(\ell') \text{ for some } \ell' \in \mathcal{L}_{s'}^{\text{cut}}\}$ and, for a given $s' \in Q_s^{\text{snd}}$, $\mathcal{R}_{s,s'}^{\text{snd}} \doteq \{k \in \mathcal{K}_s : k = R_{s'}(\ell') \text{ for some } \ell' \in \mathcal{L}_{s'}^{\text{cut}}\}$. We can interpret the set Q_s^{snd} as the ids of the sub-domains, where all the cell paths ending at D_s start. On the other hand, the set $\mathcal{R}_{s,s'}^{\text{snd}}$ contains the ids of the last cell in each path that ends at D_s and starts at $D_{s'}$. In contrast to the sets Q_s^{rcv} and $\mathcal{R}_{s,s'}^{\text{rcv}}$, it is clear that we cannot construct the sets Q_s^{snd} and $\mathcal{R}_{s,s'}^{\text{snd}}$ without communications since their definitions are in terms of the maps $R_{s'}$ and $S_{s'}^{\text{root}}$ seen from remote sub-domains $D_{s'}$ and thus not available at the current one D_s . However, thanks to the construction of the auxiliary map $N_s : \mathcal{L}_s \rightarrow \mathcal{K}$ during the parallel cell aggregation algorithm, we can build the sets Q_s^{snd} and $\mathcal{R}_{s,s'}^{\text{snd}}$ by using only standard communications between nearest neighbor sub-domains.

The detailed construction of Q_s^{snd} and $\mathcal{R}_{s,s'}^{\text{snd}}$ is given in Alg. 4. In the algorithm, we create a tuple for each cell path that starts at the current sub-domain (i.e., one tuple for each local cut cell). These tuples have three items each (see line 3): the global id of the first cell in the path, the global id of the next cell in the path, and the current sub-domain id. Conceptually, we gradually move forward these tuples along their cell paths until they arrive to the final cell in the path, which can be possibly in another sub-domain. At the end of the process each sub-domain collects all the tuples corresponding to the cell paths that end in this sub-domain. With this information each sub-domain is able to reconstruct the sets

Q_s^{snd} and $\mathcal{R}_{s,s'}^{\text{snd}}$. Note that Alg. 4 only involves communications with nearest neighbor sub-domains. Finally, we detail in Alg. 5 how we import the required data from remote sub-domains by using the sets Q_s^{snd} and $\mathcal{R}_{s,s'}^{\text{snd}}$ and Q_s^{rcv} and $\mathcal{R}_{s,s'}^{\text{rcv}}$.

Algorithm 4: Parallel inverse path reconstruction

```

1   $\mathcal{A}_s \leftarrow \emptyset$ 
2  for  $\ell \in \mathcal{L}_s^L \cap \mathcal{L}_s^{\text{cut}}$  do
3    Add the tuple  $(K_s(\ell), N_s(\ell), s)$  to the set  $\mathcal{A}_s$ 
4  while there is a tuple  $(k, n, z) \in \mathcal{A}_s$  such that  $n \in \mathcal{K}_s^L$  and  $L_s(n) \in \mathcal{L}_s^{\text{cut}}$  do
5    for  $(k, n, z) \in \mathcal{A}_s$  do
6      if  $n \in \mathcal{K}_s^L$  and  $L_s(n) \in \mathcal{L}_s^{\text{cut}}$  then
7         $n' \leftarrow N_s(L_s(n))$ 
8        Replace the tuple  $(k, n, z)$  with  $(k, n', z)$  in the set  $\mathcal{A}_s$ 
9   $Q_s^{\text{snd}} \leftarrow \emptyset$ 
10  $\mathcal{R}_{s,s'}^{\text{snd}} \leftarrow \emptyset$  for all  $s' \in Q_s^{\text{snd}}$ 
11 for  $(k, n, z) \in \mathcal{A}_s$  do
12   if  $n \in \mathcal{K}_s^L$  and  $L_s(n) \in \mathcal{L}_s^{\text{in}}$  then
13     Add  $z$  to the set  $Q_s^{\text{snd}}$ 
14     Add  $k$  to the set  $\mathcal{R}_{s,z}^{\text{snd}}$ 
15     Delete the tuple  $(k, n, z)$  from  $\mathcal{A}_s$ 
16 for  $s' \in S_s^{\text{nei}}$  do
17    $\mathcal{A}_{s,s'}^{\text{snd}} \leftarrow \emptyset$ 
18   for  $(k, n, z) \in \mathcal{A}_s$  do
19     if  $s' = S_s(L_s(n))$  then
20       Add the tuple  $(k, n, z)$  to  $\mathcal{A}_{s,s'}^{\text{snd}}$ 
21   Send the set  $\mathcal{A}_{s,s'}^{\text{snd}}$  to processor  $s'$ 
22 Delete the set  $\mathcal{A}_s$ 
23  $\mathcal{A}_s \leftarrow \emptyset$ 
24 for  $s' \in S_s^{\text{nei}}$  do
25   Receive the set  $\mathcal{A}_{s,s'}^{\text{rcv}}$  from processor  $s'$  as a counterpart of the send operation in line 21
26   Add all tuples in  $\mathcal{A}_{s,s'}^{\text{rcv}}$  to the set  $\mathcal{A}_s$ 
27 if  $\mathcal{A}_{s'} \neq \emptyset$  still in some sub-domain  $s' \in \mathcal{S}$  then
28   Go to line 4 in all processors

```

3.6 Linear algebra data distribution layout

Before detailing the FE assembly, we need to specify the data model used for representing the matrices and vectors in a distributed-memory context. Here, we consider row-wise partitioned globally addressable matrices and vectors (as we detail in the next paragraph). This choice is motivated by the fact that the vast majority of parallel linear algebra packages (e.g., PETSc [32, 33], or TRILINOS [31]) use this data distribution model. Other data layouts are also possible, such as sub-assembled distributed matrices for non-overlapping domain decomposition solvers [49, 23], but they are out of the scope of the current work.

In a row-wise partitioning model, the global matrix and the global vector are partitioned by rows, i.e., each processor owns a consecutive, non-overlapping, range of global rows of $\mathbf{A}$ and $\mathbf{b}$, respectively. Since each row of $\mathbf{A}$ and $\mathbf{b}$ is associated with a global DOF id in $\mathcal{I}^{\text{in}}$, the consecutive row partition implies that each sub-domain D_s owns a consecutive subset of the global DOF ids in $\mathcal{I}^{\text{in}}$, denoted as

Algorithm 5: Import required data from remote root cells

```

1 for  $s' \in Q_s^{\text{snd}}$  do
2    $\mathcal{B}_{s,s'} \leftarrow \emptyset$ 
3   for  $k \in \mathcal{R}_{s,s'}^{\text{snd}}$  do
4      $\ell \leftarrow L_s(k)$ 
5     Add to  $\mathcal{B}_{s,s'}$  the pair  $(x_{s,\ell}^a, g_{s,\ell}^a)$  for all  $a \in \Lambda(T_s^\ell)$ 
6   Send the buffer  $\mathcal{B}_{s,s'}^{\text{snd}}$  to sub-domain  $s'$ 
7 for  $s' \in Q_s^{\text{rcv}}$  do
8   Receive the buffer  $\mathcal{B}_{s,s'}^{\text{rcv}}$  from sub-domain  $s'$  as counterpart of the send operations in line 6
9   for  $k \in \mathcal{R}_{s,s'}^{\text{rcv}}$  do
10     $z \leftarrow Z_s(k)$ 
11    for  $a \in \Lambda(\tilde{T}_s^z)$  do
12      Get the next pair  $(x, g)$  in  $\mathcal{B}_{s,s'}^{\text{rcv}}$ 
13       $\tilde{x}_{s,z}^a \leftarrow x; \tilde{g}_{s,z}^a \leftarrow g$ 

```

$\mathcal{I}_s^{\text{in},0}$, where $\{\mathcal{I}_s^{\text{in},0}\}_{s \in \mathcal{S}}$ is a partition of $\mathcal{I}^{\text{in}}$. Note that we do not need to partition the set of constrained DOF ids $\mathcal{I}^{\text{out}}$ since, in our approach, these DOFs are removed from the global linear system. On the other hand, we consider linear algebra data structures that are *globally addressable*. That is, a given processor with sub-domain id $s \in \mathcal{S}$ may contribute to a global matrix row $[\mathbf{A}]_{gi}$, $i \in \mathcal{I}^{\text{in}}$ (or a global vector entry $[\mathbf{b}]_g$), which is not necessarily owned by it, namely $g \in \mathcal{I}^{\text{in}} \setminus \mathcal{I}_s^{\text{in},0}$. This is possible since, once the corresponding entries of the matrix and vector are set in the assembly process, a final communication stage transfers all (touched) values $[\mathbf{A}]_{gi}$, $[\mathbf{b}]_g$ such that $g \in \mathcal{I}^{\text{in}} \setminus \mathcal{I}_s^{\text{in},0}$ to the corresponding remote processor that owns the global id g . This communication phase is typically resolved by the linear algebra library and, thus, we do not cover it here for simplicity.

Even though globally addressable linear algebra objects allow one to assign entries that are not owned by the current processor, it is desirable to minimize the number of such assignments in order to reduce inter processor communication. To this end, we define the set of owned global DOF ids $\mathcal{I}_s^{\text{in},0}$ following the usual approach in FE applications. Let $\mathcal{I}_s^{\text{in}} \doteq \{i \in \mathcal{I}^{\text{in}} : i = g_{s,\ell}^a \text{ for some } a \in \Lambda(T_s^\ell) \text{ and } \ell \in \mathcal{L}_s^{\text{L}}\}$ be the set of interior global DOF ids touched only by local cells of sub-domain D_s and let $\mathcal{B}_s(i) \doteq \{\ell \in \mathcal{L}_s : i = g_{s,\ell}^a \text{ for some } a \in \Lambda(T_s^\ell)\}$ be the ids of the cells around a given id $i \in \mathcal{I}_s^{\text{in}}$. With this notation, we define the set of owned global DOF ids as $\mathcal{I}_s^{\text{in},0} \doteq \{i \in \mathcal{I}_s^{\text{in}} : s \leq S_s(\ell) \text{ for all } \ell \in \mathcal{B}_s(i)\}$. Note that each processor owns all DOF ids which are surrounded solely by cells in $\mathcal{T}_s^{\text{L}}$ plus some of the DOF ids that are on the interface between cells in $\mathcal{T}_s^{\text{L}}$ and $\mathcal{T}_s^{\text{G}}$. With this model, most of the cells in $\mathcal{T}_s^{\text{L}}$ contribute to locally owned DOF ids in $\mathcal{I}_s^{\text{in},0}$ during the assembly process (i.e., most of the values set from sub-domain D_s will correspond to values owned by D_s if we assume that the number of DOFs in the interior of the sub-domain D_s is larger than number of DOFs on its boundary). On the other hand, for DOF ids $i \in \mathcal{I}_s^{\text{in}}$ that are located at the interface between cells in $\mathcal{T}_s^{\text{L}}$ and $\mathcal{T}_s^{\text{G}}$ (i.e., at sub-domain boundaries), we define the owner sub-domain as the sub-domain with smaller id among all sub-domain that share the DOF id i . We have used this tie-breaker strategy in order to decide which sub-domain is the owner for interface DOFs, but other criteria can be also considered as long as they are consistent across all processors and require no communication.

Even though it is not strictly required, it is convenient, in order to facilitate the interaction with the parallel linear algebra package, that the sets $\{\mathcal{I}_s^{\text{in},0}\}_{s \in \mathcal{S}}$ contain contiguous, sequentially increasing DOF ids, namely $\mathcal{I}_1^{\text{in},0} = \{1, \dots, m_1^{\text{in}}\}$, $\mathcal{I}_2^{\text{in},0} = \{m_1^{\text{in}} + 1, \dots, m_1^{\text{in}} + m_2^{\text{in}}\}$, etc., where $m_s^{\text{in}} \doteq |\mathcal{I}_s^{\text{in},0}|$ for $s \in \mathcal{S}$. Methods for the generation of cell-wise global DOF ids $g_{s,\ell}^a$ that result in sets $\{\mathcal{I}_s^{\text{in},0}\}_{s \in \mathcal{S}}$ fulfilling this property are well known in distributed-memory FE analysis (see, e.g., [48] for the algorithm implemented in deal.II and [45] for the algorithm implemented in FEMPAR). These

techniques can readily be used to generate the cell-wise global DOF ids $g_{s,\ell}^a$ in the context of the AgFEM method by considering the following steps. First, the values $g_{s,\ell}^a$ are distributed over interior cell ids $\ell \in \mathcal{L}_s^{\text{in}}$ by applying the conventional algorithms [48, 45] restricted to the interior portion of the mesh (i.e., as though cut cells were removed from the mesh). In a second phase, global DOF ids $g_{s,\ell}^a$ are assigned to the vertices, edges and faces of cut cells that are in contact with an interior cell (if any) by taking the corresponding values of the neighbor interior cell. Note that, at the end of this process, the cell-wise global DOF ids $g_{s,\ell}^a$ are not defined in the interior of cut cells nor on the parts of their boundaries that are not in contact with interior cells. This is not a problem since this information is not required (see Sect. 3.7).

3.7 Distributed-memory finite element assembly

Finally, we discuss the assembly of the global matrix $\mathbf{A}$ and vector $\mathbf{b}$ in a distributed-memory context. The assembly at the current sub-domain D_s is performed with a loop in local cells in $\mathcal{T}_s^L$. At each cell $T_s^\ell \in \mathcal{T}_s^L$, we compute the corresponding elemental matrix and vector, namely $\mathbf{A}_s^\ell$ and $\mathbf{b}_s^\ell$, whose entries are defined as $[\mathbf{A}_s^\ell]_{ab} \doteq a_s^\ell(\phi_{s,\ell}^a, \phi_{s,\ell}^b)$ and $[\mathbf{b}_s^\ell]_a \doteq b_s^\ell(\phi_{s,\ell}^a)$, where $a_s^\ell(\cdot, \cdot)$ and $b_s^\ell(\cdot)$ represent the restriction of the bi-linear form $a(\cdot, \cdot)$ and the linear form $b(\cdot)$ to cell T_s^ℓ , resp. Then, the contributions of the elemental matrix and vector $\mathbf{A}_s^\ell$ and $\mathbf{b}_s^\ell$ are added to the corresponding entries of the global matrix and vector $\mathbf{A}$ and $\mathbf{b}$.

The assembly operations to be performed are closely related to the ones detailed in Sect. 2.5 for the serial case. In order to grasp the main differences between the serial and the distributed memory implementations, we rewrite the definition of $\mathcal{V}_s^{\text{agg}}$ in Eq. (6) using the notation introduced in Sect. 3.5, namely

$$\mathcal{V}_s^{\text{agg}} \doteq \{v_s \in \mathcal{V}_s^{\text{std}} : v_s^j = \sum_{g \in \mathcal{M}_s(j)} C_s^{jg} [v^g]_s \text{ for any } j \in \mathcal{J}_s^{\text{out}}\}, \quad (7)$$

where $\mathcal{M}_s(j) \subset \mathcal{I}^{\text{in}}$ denotes the set containing the global DOF ids that constrain the value associated with a given sub-domain local id $j \in \mathcal{J}_s^{\text{out}}$. The values C_s^{jg} , for $g \in \mathcal{M}_s(j)$ and $j \in \mathcal{J}_s^{\text{out}}$, are the coefficients of the cell aggregation constraints that are needed to be computed in sub-domain D_s , and the value $[v^g]_s$ represents a variable that stores (in the local scope of sub-domain D_s) the value of v^g for each master global DOF id $g \in \mathcal{M}_s(j)$ and each $j \in \mathcal{J}_s^{\text{out}}$. These quantities are formally defined as

$$\begin{aligned} \mathcal{M}_s(j) &\doteq \{g_{s,\ell}^a\}_{a \in \Lambda(T_s^\ell)}, & C_s^{jg} &\doteq \phi_{s,\ell}^b(x_s^j), & [v^g]_s &\doteq v_{s,\ell}^b & \text{ if } k \in \mathcal{K}_s \\ \mathcal{M}_s(j) &\doteq \{\tilde{g}_{s,z}^a\}_{a \in \Lambda(\tilde{T}_s^z)}, & C_s^{jg} &\doteq \tilde{\phi}_{s,z}^b(x_s^j), & [v^g]_s &\doteq \tilde{v}_{s,z}^b & \text{ otherwise} \end{aligned}$$

with $\ell = L_s(k)$, $z = Z_s(k)$, $k = R_s(L_s^{\text{own}}(j))$, and being b such that $g = g_{s,\ell}^b$ or $g = \tilde{g}_{s,z}^b$. Note that the definition of these quantities takes one form or another depending on whether the global cell id $k = R_s(L_s^{\text{own}}(j))$ belongs to $\mathcal{K}_s$ or not (i.e., if the root cell T^k associated with j is a locally relevant cell in D_s or not). If $k \in \mathcal{K}_s$, the quantities can be computed from data stored in the local scope of the current processor. Otherwise, if $k \notin \mathcal{K}_s$ (which implies $k \in \mathcal{K}_s^R$), the quantities need to be computed from the specially allocated variables $\tilde{g}_{s,z}^b$, $\tilde{x}_{s,z}^b$, and $\tilde{v}_{s,z}^b$ presented in Sect. 3.5.

The main differences between the serial and distributed-memory assembly are readily grasped by comparing formulas (1) and (7). The first main difference is the following. Note that, in both the serial and distributed implementations, the master DOFs are represented with global DOF ids. However, the constrained DOFs are represented with global ids in the serial implementation and with sub-domain local ids in the distributed-memory version. For this reason, it is not needed to generate a global numbering for the constrained DOFs. The second major difference is that, in the distributed memory implementation, one has to import data from remote sub-domains in order to compute the quantities $\mathcal{M}_s(j)$, C_s^{jg} , and $[v^g]_s$, for $g \in \mathcal{M}_s(j)$, if the root cell associated with the constrained DOF id $j \in \mathcal{J}_s^{\text{out}}$ is not a locally relevant cell. The assembly operations to be performed in the distributed-memory case are obtained by modifying the operations of the serial case (Table 1) in accordance with these

two main differences and taking into account that cells are represented by sub-domain local ids in the parallel implementation. The result is detailed in Table 2.

Case	Assembly operation
Right hand side vector	
$j_{s,\ell}^a \in \mathcal{J}_s^{\text{in}}$	$[\mathbf{b}]_i += [\mathbf{b}_s^\ell]_a, i = g_{s,\ell}^a \quad (\text{standard case})$
$j_{s,\ell}^a \notin \mathcal{J}_s^{\text{in}}$	$[\mathbf{b}]_i += C_s^{ni} [\mathbf{b}_s^\ell]_a, n = j_{s,\ell}^a, \forall i \in \mathcal{M}_s(n)$
System matrix	
$j_{s,\ell}^a \in \mathcal{J}_s^{\text{in}}$ and $j_{s,\ell}^b \in \mathcal{J}_s^{\text{in}}$	$[\mathbf{A}]_{ij} += [\mathbf{A}_s^\ell]_{ab}, i = g_{s,\ell}^a, j = g_{s,\ell}^b \quad (\text{standard case})$
$j_{s,\ell}^a \notin \mathcal{J}_s^{\text{in}}$ and $j_{s,\ell}^b \in \mathcal{J}_s^{\text{in}}$	$[\mathbf{A}]_{ij} += C_s^{ni} [\mathbf{A}_s^\ell]_{ab}, n = j_{s,\ell}^a, j = g_{s,\ell}^b, \forall i \in \mathcal{M}_s(n)$
$j_{s,\ell}^a \in \mathcal{J}_s^{\text{in}}$ and $j_{s,\ell}^b \notin \mathcal{J}_s^{\text{in}}$	$[\mathbf{A}]_{ij} += C_s^{mj} [\mathbf{A}_s^\ell]_{ab}, i = g_{s,\ell}^a, m = j_{s,\ell}^b, \forall j \in \mathcal{M}_s(m)$
$j_{s,\ell}^a \notin \mathcal{J}_s^{\text{in}}$ and $j_{s,\ell}^b \notin \mathcal{J}_s^{\text{in}}$	$[\mathbf{A}]_{ij} += C_s^{ni} C_s^{mj} [\mathbf{A}_s^\ell]_{ab}, n = j_{s,\ell}^a, m = j_{s,\ell}^b, \forall i \in \mathcal{M}_s(n), \forall j \in \mathcal{M}_s(m)$

Table 2: Summary of the assembly operations related with the agFE space $\mathcal{V}^{\text{agg}}$ in the distributed-memory implementation.

4 NUMERICAL EXAMPLES

4.1 Experimental environment

The numerical experiments are run at the Marenstrum-IV supercomputer [50], hosted by the Barcelona Supercomputing Center. The Marenstrum-IV is a petascale machine equipped with 3,456 compute nodes interconnected with the Intel OPA HPC network. Each node has 2x Intel Xeon Platinum 8160 multi-core CPUs, with 24 cores each (i.e. 48 cores per node) and 96 GBytes of RAM.

With respect to the software, we used the MPI-parallel implementation of the AgFEM method available at FEMPAR [35] linked against p4est v2.0 [41] as the Cartesian grid manipulation engine, and PETSc v3.9.0 [33] for distributed-memory linear algebra data structures and solvers. These software were compiled with Intel v18.0.5 compilers using system recommended optimization flags and linked against the Intel MPI Library (v2018.4.057) for message-passing and the BLAS/LAPACK available on the Intel MKL library for optimized dense linear algebra kernels. All floating-point operations were performed in IEEE double precision.

4.2 Model problem

As a model problem for the numerical examples, let us consider the Poisson equation with Dirichlet boundary conditions. After scaling with the diffusion term, the equation reads: find $u \in H^1(\Omega)$ such that

$$-\Delta u = f \quad \text{in } \Omega, \quad u = g \quad \text{on } \Gamma \doteq \partial\Omega, \quad (8)$$

where $f \in H^{-1}(\Omega)$ is the source term and $g \in H^{1/2}(\Gamma)$ is the prescribed value on the Dirichlet boundary. We denote by $\mathcal{V}^x$ any of the FE spaces $\mathcal{V}^{\text{std}}$, or $\mathcal{V}^{\text{agg}}$, when it is not necessary to distinguish between them. We approximate problem (8) in $\mathcal{V}^x$ with the following variational equation: find $u_h \in \mathcal{V}^x$ such that $a(v_h, u_h) = b(v_h)$ for all $v_h \in \mathcal{V}^x$, where

$$\begin{aligned} a(v, u) &\doteq \int_{\Omega} \nabla v \cdot \nabla u \, d\Omega + \int_{\Gamma} (\tau v u - v(n \cdot \nabla u) - u(n \cdot \nabla v)) \, d\Gamma, \text{ and} \\ b(v) &\doteq \int_{\Omega} v f \, d\Omega + \int_{\Gamma} (\tau v g - (n \cdot \nabla v) g^D) \, d\Gamma, \end{aligned} \quad (9)$$

with n being the outward unit normal on Γ . Note that the forms $a(\cdot, \cdot)$ and $b(\cdot)$ include the usual terms resulting from the integration by parts of (8) plus additional terms associated with the weak

imposition of Dirichlet boundary conditions with Nitsche's method [51, 52]. The coefficient $\tau > 0$ is a mesh-dependent parameter that has to be large enough to ensure the coercivity of $a(\cdot, \cdot)$. We consider a weak imposition of boundary conditions since it is not straightforward to include prescribed values in the approximation space in a strong manner, when dealing with an unfitted grid. Nitsche's method is commonly used in the embedded boundary community for circumventing this problem (see, e.g., [7, 9]) since it provides a consistent numerical scheme with optimal converge rates (also for high-order FEs).

When using the agFE space (i.e., $\mathcal{V}^x = \mathcal{V}^{\text{agg}}$), we compute the mesh-dependent parameter τ as

$$\tau = \frac{\beta}{h}, \quad (10)$$

where h is the cell size in the computational grid and β is a user defined constant parameter. In the numerical examples, we take two different values of β , namely $\beta = 10.0$ and $\beta = 100.0$. The choice of $\beta = 10.0$ is common in the literature of unfitted methods (see, e.g., reference [53]). We also have considered $\beta = 100.0$ in order to study the sensibility of the method with respect to this parameter. We note that using the space $\mathcal{V}^{\text{agg}}$ and formula (10) for the computation of τ , the bilinear form in (9) is coercive independently of the position of the cuts. We refer to [10] for the mathematical proof of this statement. However, when using the standard FE space $\mathcal{V}^{\text{std}}$, formula (10) does not necessarily lead to a coercive bi-linear form. For the simulations with the standard space $\mathcal{V}^{\text{std}}$, we compute τ following an alternative approach in order recover a coercive problem independently of the position of the cuts (see [15] for details). For the sake of completeness, we briefly present this methodology here. It consists in building the penalty parameter τ as an element-wise constant function, where the value in a generic cut cell, namely $\tau_k \doteq \tau|_{T^k}$, $k \in \mathcal{K}^{\text{cut}}$, is computed as $\tau_k = \beta \lambda_k^{\text{max}}$. In this context, β is also a user-defined constant (with values $\beta = 10.0$ and $\beta = 100.0$ in the numerical examples). On the other hand, λ_k^{max} is the maximum eigenvalue of the following generalized eigenvalue problem: find $\eta_k \in \mathcal{V}(T^k)$ and $\lambda_k \in \mathbb{R}$ such that

$$\int_{T^k \cap \Omega} \nabla \eta_k \cdot \nabla \xi_k \, d\Omega = \lambda_k \int_{T^k \cap \Gamma} (\nabla \eta_k \cdot n)(\nabla \xi_k \cdot n) \, d\Gamma \text{ for all } \xi_k \in \mathcal{V}(T^k) \text{ for each cut cell id } k \in \mathcal{K}^{\text{cut}}.$$

Note that this problem is local at each cut cell, and its definition only depends on how the physical domain intersects the cut cell and the functional space used within the cell. We refer to [15] for the specific details on the solution of this problem. As indicated in [15], one of the main drawbacks of using the standard FE space $\mathcal{V}^{\text{std}}$ is that the optimal value of τ needed for ensuring coercivity of the bilinear form $a(\cdot, \cdot)$ (i.e., λ_k^{max}) might be arbitrarily large if the active part of a cell $T \in \mathcal{T}$, namely $\Omega \cap T$, has a volume that tends to zero. This fact leads to strongly ill-conditioned systems of linear equations when considering the standard space $\mathcal{V}^{\text{std}}$. This issue is solved when considering the agFE space $\mathcal{V}^{\text{agg}}$ as the optimal value of τ , which leads to a coercive bi-linear form, is independent of the position of the cuts (see [10] for further details).

4.3 Linear solver setup

The FE discretization of the weak problem represented by the bi-linear and linear forms in (9) leads to a system of linear algebraic equations, namely $\mathbf{Ax} = \mathbf{b}$, which, at large scales, requires efficient parallel linear solvers. Our goal is to show that the usage of the agFE spaces allows one to efficiently solve the linear systems using solvers that are well established for conventional FE analysis based on body-fitted grids. To this end, we consider the widely available linear solvers available in the PETSc library [33]. In particular, we use a conjugate gradient method from the KSP module of PETSc, preconditioned with a smoothed-aggregation AMG preconditioner called GAMG (see [54] for specific details).

We set up the linear solver by relying as much as possible on the default configuration given by GAMG to effectively show that the agFE spaces lead to linear systems that are efficiently solved using standard multigrid tools. The modifications that we have introduced in the default parameters are,


```

1 -ksp_type cg
2 -ksp_rtol 1.0e-6
3 -ksp_converged_reason
4 -ksp_max_it 500
5 -ksp_norm_type unpreconditioned
6 -pc_type gamg
7 -pc_gamg_type agg
8 -mg_coarse_sub_pc_type cholesky
9 -mg_levels_esteig_ksp_type cg
10 -pc_gamg_square_graph 0
11 -build_twosided redscatter

```

Listing 1: Contents of the used PETSc configuration file.

basically, in order to build a preconditioner tailored for symmetric definite positive matrices (since the default configuration of GAMG is designed for the non-symmetric case). We report the configuration file used to set up the PETSc solver and preconditioner in Listing 1. The file is written using a domain specific language provided by the PETSc library (see the PETSc users' manual [33] for further details on the file syntax).

In line 1 of Listing 1, we select a conjugate gradient method to compute an approximation $\mathbf{x}^{\text{cg}} \approx \mathbf{x}$ of the solution of the symmetric definite positive system $\mathbf{Ax} = \mathbf{b}$. We declare convergence when $\|\mathbf{r}\|_2 / \|\mathbf{b}\|_2 < 10^{-6}$ within the first 500 iterations (see lines 2-5), where $\mathbf{r} \doteq \mathbf{b} - \mathbf{Ax}^{\text{cg}}$ is the unpreconditioned residual and $\|\cdot\|_2$ denotes the Euclidean norm. In lines 6-7, we select the smoothed aggregation AMG preconditioner called GAMG and in line 8 we select a Cholesky factorization in order to solve the problem on the coarsest level of the AMG hierarchy. The GAMG preconditioner uses Chebyshev smoothers in order to approximate the solution at the finest and intermediate levels of the AMG hierarchy. In order to set up a Chebyshev smoother one needs to estimate the maximum eigenvalue of the corresponding system matrix (see, e.g., [55] for details). To compute this eigenvalue estimate, GAMG uses internally an iterative Krylov sub-space method. In line 9, we instruct GAMG to use a conjugate gradient iteration for this purpose. In particular, we are using the default number of iterations provided by PETSc (i.e., 10 iterations) in the conjugate gradient method used to estimate the eigenvalue. In the process of generating the prolongation and restriction operators associated with the AMG hierarchy, GAMG allows one to use the graph of the matrix $\mathbf{A}^T \mathbf{A}$ instead of the graph of the original matrix $\mathbf{A}$. By default, GAMG uses the graph of $\mathbf{A}^T \mathbf{A}$ in the finest level and the graph of (the coarse representation of) $\mathbf{A}$ at the other levels. In line 10, we instruct GAMG to use the graph of $\mathbf{A}$ also in the finest level. We have observed that this setting leads to a faster solver (since the product $\mathbf{A}^T \mathbf{A}$ is not computed) even though the complexity of the AMG preconditioner is slightly increased. Finally, we have included the option in line 11 in order to bypass a software BUG in the Intel MPI implementation used in the experiments. Please refer to the GAMG package documentation [39] for the default methods and parameters that are not covered in Listing 1.

4.4 Problem setup

In this section, we detail the rest of the parameter values and methods used in the numerical examples. In order to show that the methods are robust enough to deal with different geometrical data, we have performed all the experiments for three different complex geometries. On the one hand, we have considered a bulky 3D body whose shape reminds the one of a popcorn flake (see Fig. 7a). This "popcorn-flake" geometry is often used in the literature to study the performance of unfitted FE methods (see, e.g., [7]). The popcorn flake geometry considered here is obtained by taking the one defined in [7], scaling it by a factor of 0.5 and translating it by a value of 0.5 in each direction such that the body fits in the unit cube $[0, 1]^3$. The second geometry is a massive spiral pipe (see Fig. 7b). The radius of its tubular cross section is 0.1, whereas the radius of the spiral central axis is 0.875. The third geometry, called "Swiss cheese", is a complex body often used in the literature of unfitted FE

methods. The precise definition of this geometry is found in [7]. In all cases, the bounding box used to define the background mesh is the unit cube $[0, 1]^3$, which is also depicted in Fig. 7.

(a) Popcorn flake.

(b) Spiral.

(c) Swiss Cheese.

Figure 7: Views of the three geometries and the bounding boxes considered in the numerical examples.

The source term function f and the Dirichlet condition function g of the Poisson problem (8) are defined using the method of the manufactured analytic solutions so that problem (8) has the following exact solution: $u(x, y, z) = x + y + z$. This exact solution is used to check the quality of the computed FE approximations. Since the exact solution u belongs to the considered FE spaces, the computed FE approximations u_h have to coincide with u up to the tolerance used to declare convergence in the

linear solver. In order to check that this is true in our experiments, we have computed the (relative) error norms $|u - u_h|_{H^1}/|u|_{H^1}$ and $\|u - u_h\|_{L^2}/\|u\|_{L^2}$, where

$$|u|_{H^1} \doteq \left(\int_{\Omega} \nabla u \cdot \nabla u \, d\Omega \right)^{1/2} \quad \text{and} \quad \|u\|_{L^2} \doteq \left(\int_{\Omega} u^2 \, d\Omega \right)^{1/2}$$

are the H^1 semi-norm and L^2 norm, resp.

As indicated in Sect. 2.1, we consider background meshes of hexahedral cells generated and partitioned in parallel with the `p4est` library. We generate the meshes by recursive sub-division of the unit cube $[0, 1]^3$. That is, in the first refinement level, the cube is decomposed into $(2^1)^3 = 8$ cells, in the second refinement level into $(2^2)^3 = 64$ cells, etc. For the weak scaling test shown below, we generate three families of partitions such that the average number of cells per sub-domain is constant as the mesh is refined (see Table 3). The first one (referred as load 1) has approximately 1,000 cells per sub-domain, and the second and the third ones (load 2 and load 3) have approximately 8,000 and 64,000 cells per sub-domain, resp. The Swiss cheese geometry is not simulated with the coarsest mesh since it is unable to capture the main geometrical features. The finest mesh considered in the examples has 1,073,741,824 cells (corresponding to refinement level 10), and it is partitioned into 16,777 sub-domains (which are mapped to the same number of MPI tasks). We used the default MPI task placement policy of Intel MPI (v2018.4.057) with partially filled nodes. The number of nodes N is selected as $N = \text{ceil}(|S|/48)$, where $\text{ceil}(x)$ rounds x to the nearest integer that is greater or equal to x . If $|S|$ is not multiple of 48, the placement policy fully populates the first $N - 1$ nodes with 48 MPI tasks per node. The remaining $\text{mod}(|S|, 48)$ MPI tasks are placed on the last node, with mod denoting integer division remainder.

S	Number of cells		
	load 1	load 2	load 3
4	4,096 (34% 40% -)	32,768 (29% 29% 56%)	262,144 (26% 24% 47%)
32	32,768 (29% 29% 56%)	262,144 (26% 24% 47%)	2,097,152 (25% 22% 41%)
262	262,144 (26% 24% 47%)	2,097,152 (25% 22% 41%)	16,777,216 (24% 20% 38%)
2,097	2,097,152 (25% 22% 41%)	16,777,216 (24% 20% 38%)	134,217,728 (24% 20% 36%)
16,777	16,777,216 (24% 20% 38%)	134,217,728 (24% 20% 36%)	1,073,741,824 (24% 20% 35%)

Table 3: Number of sub-domains ($|S|$) and number of total cells in the background mesh for the three families of partitioned meshes considered in the examples. In parentheses, the percentage of active cells (over the total number of background cells) for each of the three geometries is given. Within a parenthesis, the three given quantities correspond to the "pop-corn flake", the spiral, and the "Swiss cheese" geometries, respectively.

Each mesh is decomposed using the load-balancing mechanism of the `p4est` library. By default, the method splits the mesh such that the number of cells in each of the resulting sub-domains is approximately equally balanced. `p4est` also includes the option to provide a user-defined weight to each of the cells. In this case, the algorithm leads to partitions, where the sum of the weights in each sub-domain is equally balanced across the entire partition. We use this functionality in order to give different weights to active and exterior cells. In this context, one has to find a trade-off between two conflicting goals. On the one hand, it is desirable to load balance the storage of the background mesh and the cost of the geometrical computations that run on top of it. In this case, one has to assign the same weight to active and exterior cells. On the other hand, it is also desirable to load balance the assembly, storage, and solution of the underlying systems of linear algebraic equations. In that case, exterior cells do not play any role, and one has to balance the number of active cells per sub-domain. This is achieved in practice by assigning a much smaller weight to exterior cells than to active ones. In the experiments, we have used a ratio $w = 10/1$ (i.e., weight 10 for active cells and 1 for exterior ones). We observed experimentally (see Figs. 8 and 9) that this value of w leads to the best trade-off

among the two aforementioned factors for the considered geometries. The wall clock time of the main phases of the AgFEM method (Fig. 8) and the wall clock time of the linear solver (Fig. 9) decreases significantly when the weight factor w moves from $w = 1$ to $w = 10$. Beyond this point, the timings tend to stagnate or slightly increase for $w \geq 10$. An important exception is the assembly of the algebraic linear system, whose wall clock time significantly increases again for $w \geq 10$ (see the orange curve with pentagonal markers in Fig. 8). For this reason, we have chosen a value of $w = 10$ for the subsequent numerical experiments.

Figure 8: Wall clock time for the major phases of the AgFEM method versus the weighting factor w . Results computed with agFE spaces, $\beta = 10$, and load 3 on 262 CPU cores.

We build the FE spaces $\mathcal{V}^{\text{std}}$ and $\mathcal{V}^{\text{agg}}$ using hexahedral elements with continuous piece-wise trilinear shape functions. We consider both agFE and conventional spaces in order to evaluate the benefits of using the AgFEM method. A summary of the main parameters and methods considered in the numerical examples is given in Table 4.

4.5 Weak scaling test

We start this section devoted to the weak scaling test by checking that the problems were accurately solved in the experiments. To this end, Fig. 10 shows the relative L^2 norm of the error associated with the computed numerical solutions versus problem size for the three loads per processor detailed in Table 3. The results in Figs. 10a, 10b, and 10c are for the "popcorn flake", the spiral and the "Swiss cheese" geometries. Note that the results are qualitatively equivalent for the three geometries. In Fig. 10, some points of the red lines (the ones obtained with standard FE spaces) are missing since the linear solver was not able to provide a solution. Either the setup of the preconditioner failed or the computed preconditioner was not positive definite (which makes it unusable in a conjugate gradients iteration). These problems are related with the strong ill-conditioning that results from the usage of standard FE spaces. Note, however, that the linear solvers properly worked in all cases, when

Figure 9: Wall clock time for the main linear solver phases versus the weighting factor w . Results computed with agFE spaces, $\beta = 10$, and load 3 on 262 CPU cores.

Description	Considered methods/values
Model problem	Poisson equation
Problem geometry	"Popcorn flake", massive spiral, and "Swiss cheese"
Parallel mesh generation and partitioning	p4est library
FE spaces	Aggregated $\mathcal{V}^{\text{agg}}$ and standard $\mathcal{V}^{\text{std}}$
Cell type	Hexahedral cells
Interpolation	Piece-wise trilinear shape functions
Linear solver	Preconditioned conjugate gradients (PETSc)
Preconditioner	Smoothed-aggregation AMG (GAMG preconditioner from PETSc)
Weighting factor for space filling-curve	$w = 10$
Coefficient used in Nitsche's penalty terms	$\beta = 10.0$ and $\beta = 100.0$

Table 4: Summary of the main parameters and computational strategies used in the numerical examples.

considering the agFE spaces. Note that (as expected) the relative errors of the computed numerical solution are of the same order of magnitude as the tolerance imposed to declare convergence of the linear solver. Moreover, the computed errors in the L^2 norm are reduced about three orders of magnitude, if the problem is solved by setting a linear solver tolerance three orders of magnitude smaller, i.e., 10^{-9} instead of 10^{-6} (see Fig. 11). This clearly demonstrates that the solution quality is dictated by the linear solver tolerance as one would expect. Results for the H^1 norm of the error are qualitatively equivalent as the ones showed in Fig. 10 and they are not included for the sake of brevity.

We start the reporting of the weak scaling test, by discussing the impact of the different parameters on the algorithmic performance of the considered AMG preconditioners. First, we consider the impact

Figure 10: Weak scaling test: Quality check of the computed numerical solution. The figure shows the (relative) L^2 norm of the error versus problem size for three different loads per processor. Results computed with parameters $\beta = 10$ and $w = 10$.

of using agFE or standard FE spaces in the number of solver iterations. As the results in Fig. 12 show, it is clear that the usage of AgFEM is beneficial (and essential) for the good performance of the linear solver. As it was already discussed for Fig. 10, some results in Fig. 12 (the ones associated with the standard FE spaces) are missing since the linear solver was not able to provide a converged solution. This shows that standard FE spaces are unusable in practice since they do not lead to a robust discretization method. In contrast, the usage of agFE spaces allows one to effectively solve the underlying linear systems. In particular, note that the results obtained with agFE spaces (blue lines) are very close to the expected optimal performance of multigrid methods (i.e., number of linear solver iterations is asymptotically independent of the problem size). This optimal behavior is observed for all considered geometries, which further demonstrates the usability and robustness on the computational strategy based on agFE spaces.

Next, we study the influence the user-defined coefficient β associated with Nitsche's method (see Sect. 4.2) has in the performance of the solver. Since it is clear from Figs. 10 and 12 that the method based on standard FE spaces is unusable in practice, from now on, we only provide results associated with the agFE spaces. The results in Fig. 13 show the influence of β in the solver performance for the two values of β previously detailed in Table 4 and for the three loads per processor detailed in Table 3. As it is seen in the figure, the number of linear solver iterations are weakly scalable for the two values of β (i.e., iterations asymptotically independent of the problem size). However, the absolute number of iterations is significantly larger for $\beta = 100$ than for $\beta = 10$, which suggests that β has to be chosen large enough to ensure coercivity of the underlying bi-linear form, but cannot be chosen arbitrary large since it has a negative impact on the solver performance. This requirement is well known in the literature of Nitsche-based FE methods (see, e.g., [7]).

Figure 11: Weak scaling test: Quality check of the computed numerical solution. The figure shows the (relative) L^2 norm of the error versus problem size for three different loads per processor and for two different values of the linear solver tolerance. Results computed with aggregated FE spaces and parameters $\beta = 10$ and $w = 10$.

We conclude the numerical examples by reporting the wall clock time consumed in the main phases of the AgFEM method (see Fig. 14) and the one in the linear solver step (see Fig. 15). As it is observed in Fig. 14, all phases have remarkable weak scaling. For the linear solver (see Fig. 15), we report the wall clock times of the solver setup (i.e., the setup of the AMG preconditioner) and the wall clock time of the solver run (i.e., the preconditioned conjugate gradient iterations). Note that the scaling of the linear solver is not as optimal as for the different phases reported in Fig. 14. However, it can still be considered excellent since the solver wall clock times increase only moderately as the problem is scaled, and also if one takes into account that we have customized the linear solver as little as possible. As an illustrative example, for the popcorn flake body and for load 3, the total solver wall clock time (setup plus iterations) scales from 0.36 to 4.33 seconds, while the problem size scales from 67,192 to 298,911,096 DOFs. That is to say, the total solver time increases by a factor of $12x$, while the problem size increases by a factor of $4,449x$, which are excellent results.

4.6 Weak scaling test for the Poisson equation on the unit cube

In order to demonstrate that the (moderate) degradation of the linear solver time observed in Fig. 15 is not caused by the AgFEM method, we have performed (as a reference) a weak scaling test with a Poisson problem defined on the unit cube, with a body-fitted mesh optimally partitioned into sub-domains by means of a Cartesian partition. The test is done for three different local loads (referred to as "load 1", "load 2" and "load 3") with 20^3 , 30^3 , and 40^3 cells per sub-domain respectively. The largest mesh considered has 1,053,696,000 cells which are distributed over 16,464 sub-domains (see Table 5). The linear solver considered here is the same as the one used before in Section 4.5 (i.e., we have used the configuration file displayed in Listing 1). The results of this weak scaling test are

Figure 12: Weak scaling test: Influence of standard or agFE spaces on the linear solver behavior. The figure shows the conjugate gradient (CG) iterations versus problem size for three different loads per processor. Results computed parameters $\beta = 10$ and $w = 10$.

reported in Fig. 16. Note that the number of iterations of the conjugate gradient solver exhibits the optimal scaling behavior expected for this problem type (see Fig. 16a). However, even though this is the most favorable setup one can consider for an AMG preconditioner (i.e., Poisson problem, Cartesian body fitted meshes, optimal load balance, no AgFEM constraints), the obtained weak scaling of the solver time is not the expected optimal one (constant time as the problem size increases in the same proportion as the number of processors). This can be clearly seen in Figs. 16b and 16c. These results show that the moderate loss of scalability previously observed also for the computations based on the AgFEM method in Fig. 15 cannot be directly attributed to the AgFEM method itself. The degradation of the linear solver time could be caused by other factors such as the underlying parallel system (i.e., the HPC cluster architecture plus the software stack that runs on top of it), but confirming this is out of the scope of the current work. In any case, by performing the weak scaling test on the unit cube with a body-fitted mesh, we have discarded that the loss of efficiency is caused by the AgFEM method.

5 CONCLUSIONS

In this work, we have presented the distributed-memory implementation and performance of the so-called AgFEM method for the solution of large-scale problems with unfitted FE techniques. The method is based on removal of basis functions associated with badly cut cells by introducing carefully designed constraints, which result in well-conditioned systems of linear algebraic equations independently of the position of the cuts. As the main contributions of the paper, (a) we have proposed a parallel implementation of AgFEM, (b) shown that the proposed parallel setup is scalable, and (c) shown that, by using AgFEM, the underlying system of linear equations can be efficiently solved using preconditioned iterative linear solvers designed for conventional FE analysis. In this case,

Figure 13: Weak scaling test: Influence of the coefficient β on the linear solver behavior. The figure shows the conjugate gradient (CG) iterations versus problem size for three different loads per processor. Results computed with agFE spaces and $w = 10$.

S	Number of cells		
	load 1	load 2	load 3
48	384,000	1,296,000	3,072,000
1,296	10,368,000	34,992,000	82,944,000
3,072	24,576,000	82,944,000	196,608,000
6,000	48,000,000	162,000,000	384,000,000
10,368	82,944,000	279,936,000	663,552,000
16,464	131,712,000	444,528,000	1,053,696,000

Table 5: Weak scaling test for the Poisson equation on the unit cube: Number of sub-domains ($|S|$) and number of cells of the considered meshes.

we have considered well-known AMG preconditioners available in PETSc. This is in contrast to previous works, which considered highly customized preconditioners in order to handle the severely ill-conditioned operators resulting from the discretization based on unfitted grids. Also in contrast to previous works, which mainly consider non-scalable serial algorithms, we have studied the performance and scalability of the method with weak scaling tests up to 16K processors and up to 300M DOFs. The proposed parallel AgFEM method can be easily incorporated in existing large-scale codes since (a) it allows one to use standard parallel linear solvers, and (b) the algorithmic phases of the method can be implemented using standard functionality available in distributed-memory FE codes. In this case, the parallel implementation of the AgFEM method has been done in the large-scale FE package FEMPAR.

Figure 14: Weak scaling test: Wall clock time for the major phases of the AgFEM method versus problem size. Results computed with agFE spaces, $\beta = 10$, $w = 10$, and load 3.

ACKNOWLEDGMENTS

Financial support from the European Commission under the FET-HPC ExaQute project (Grant agreement ID: 800898) within the Horizon 2020 Framework Programme is gratefully acknowledged. This work has been partially funded by the project MTM2014-60713-P from the “Ministerio de Economía, industria y Competitividad” of Spain. S. Badia gratefully acknowledges the support received from the Catalan Government through the ICREA Acadèmia Research Program. F. Verdugo gratefully acknowledges the support received from the *Secretaria d’Universitats i Recerca* of the Catalan Government in the framework of the Beatriu Pinós Program (Grant Id.: 2016 BP 00145). The authors thankfully acknowledge the computer resources at Marenstrum-IV and the technical support provided by the Barcelona Supercomputing Center (RES-ActivityID: FI-2018-1-0014, FI-2018-2-0009, FI-2018-3-0029, FI-2019-1-0007).

REFERENCES

- [1] N. Sukumar, N. Moës, B. Moran, and T. Belytschko. Extended finite element method for three-dimensional crack modelling. *International Journal for Numerical Methods in Engineering*, 48(11):1549–1570, 2000. doi:[10.1002/1097-0207\(20000820\)48:11<1549::AID-NME955>3.0.CO;2-A](https://doi.org/10.1002/1097-0207(20000820)48:11<1549::AID-NME955>3.0.CO;2-A).
- [2] A. Massing, M. G. Larson, A. Logg, and M. E. Rognes. A Nitsche-based cut finite element method for a fluid-structure interaction problem. *Communications in Applied Mathematics and Computational Science*, 10(2):97–120, 2015. doi:[10.2140/camcos.2015.10.97](https://doi.org/10.2140/camcos.2015.10.97).
- [3] H. Sauerland and T. P. Fries. The extended finite element method for two-phase and free-surface flows: A systematic study. *Journal of Computational Physics*, 230(9):3369–3390, 2011. doi:[10.1016/j.jcp.2011.01.033](https://doi.org/10.1016/j.jcp.2011.01.033).
- [4] E. Burman, D. Elfverson, P. Hansbo, M. G. Larson, and K. Larsson. Shape optimization using the cut

Figure 15: Weak scaling test: Wall clock time for the two main linear solver phases (setup and run) versus problem size. Results computed with agFE spaces, $\beta = 10$, $w = 10$, and load 3.

finite element method. *Computer Methods in Applied Mechanics and Engineering*, 328:242–261, 2018. doi:[10.1016/j.cma.2017.09.005](https://doi.org/10.1016/j.cma.2017.09.005).

- [5] M. Chiumenti, E. Neiva, E. Salsi, M. Cervera, S. Badia, J. Moya, Z. Chen, C. Lee, and C. Davies. Numerical modelling and experimental validation in Selective Laser Melting. *Additive Manufacturing*, 18:171–185, 2017. doi:[10.1016/j.addma.2017.09.002](https://doi.org/10.1016/j.addma.2017.09.002).
- [6] L. Nguyen, S. Stoter, T. Baum, J. Kirschke, M. Ruess, Z. Yosibash, and D. Schillinger. Phase-field boundary conditions for the voxel finite cell method: Surface-free stress analysis of CT-based bone structures. *International Journal for Numerical Methods in Biomedical Engineering*, 33(12):e2880, 2017. doi:[10.1002/cnm.2880](https://doi.org/10.1002/cnm.2880).
- [7] E. Burman, S. Claus, P. Hansbo, M. G. Larson, and A. Massing. CutFEM: Discretizing Geometry and Partial Differential Equations. *International Journal for Numerical Methods in Engineering*, 104(7): 472–501, 2015. doi:[10.1002/nme.4823](https://doi.org/10.1002/nme.4823).
- [8] D. Elfverson, M. G. Larson, and K. Larsson. CutIGA with basis function removal. *Advanced Modeling and Simulation in Engineering Sciences*, 5(1):6, 2018. doi:[10.1186/s40323-018-0099-2](https://doi.org/10.1186/s40323-018-0099-2).
- [9] D. Schillinger and M. Ruess. The Finite Cell Method: A Review in the Context of Higher-Order Structural Analysis of CAD and Image-Based Geometric Models. *Archives of Computational Methods in Engineering*, 22(3):391–455, 2015. doi:[10.1007/s11831-014-9115-y](https://doi.org/10.1007/s11831-014-9115-y).
- [10] S. Badia, F. Verdugo, and A. F. Martín. The aggregated unfitted finite element method for elliptic problems. *Computer Methods in Applied Mechanics and Engineering*, 336:533–553, 2018. doi:[10.1016/j.cma.2018.03.022](https://doi.org/10.1016/j.cma.2018.03.022).
- [11] E. Nadal, J. J. Ródenas, J. Albelda, M. Tur, J. E. Tarancón, and F. J. Fuenmayor. Efficient Finite Element Methodology based on Cartesian grids: Application to structural shape optimization. *Abstract and Applied Analysis*, 2013:1–19, 2013. doi:[10.1155/2013/953786](https://doi.org/10.1155/2013/953786).
- [12] R. Mittal and G. Iaccarino. Immersed Boundary Methods. *Annual Review of Fluid Mechanics*, 37(1):

(a) Weak scaling of CG iterations.

(b) Weak scaling of linear solver (iteration phase).

(c) Weak scaling of linear solver (setup phase).

Figure 16: Weak scaling test for the Poisson equation on the unit cube: Performance of the linear solver.

239–261, 2005. doi:[10.1146/annurev.fluid.37.061903.175743](https://doi.org/10.1146/annurev.fluid.37.061903.175743).

- [13] R. A. K. Sanches, P. B. Bornemann, and F. Cirak. Immersed b-spline (i-spline) finite element method for geometrically complex domains. *Computer Methods in Applied Mechanics and Engineering*, 200(13–16): 1432–1445, 2011. doi:[10.1016/j.cma.2010.12.008](https://doi.org/10.1016/j.cma.2010.12.008).
- [14] N. Sukumar, D. L. Chopp, N. Moës, and T. Belytschko. Modeling holes and inclusions by level sets in the extended finite-element method. *Computer Methods in Applied Mechanics and Engineering*, 190(46–47): 6183–6200, 2001. doi:[10.1016/S0045-7825\(01\)00215-8](https://doi.org/10.1016/S0045-7825(01)00215-8).
- [15] F. de Prenter, C. V. Verhoosel, G. J. van Zwieten, and E. H. van Brummelen. Condition number analysis and preconditioning of the finite cell method. *Computer Methods in Applied Mechanics and Engineering*, 316:297–327, 2017. doi:[10.1016/j.cma.2016.07.006](https://doi.org/10.1016/j.cma.2016.07.006).
- [16] T. A. Davis. *Direct methods for sparse linear systems*. Society for Industrial and Applied Mathematics, 2006.
- [17] Y. Saad. *Iterative Methods for Sparse Linear Systems, Second Edition*. Society for Industrial and Applied Mathematics, 2 edition, 2003.
- [18] W. L. Briggs, V. E. Henson, and S. F. McCormick. *A Multigrid Tutorial, Second Edition*. Society for Industrial and Applied Mathematics, 2000. doi:[10.1137/1.9780898719505](https://doi.org/10.1137/1.9780898719505).
- [19] E. Chow, R. D. Falgout, J. J. Hu, R. S. Tuminaro, and U. M. Yang. 10. A Survey of Parallelization Techniques for Multigrid Solvers. In *Parallel Processing for Scientific Computing*, pages 179–201. Society for Industrial and Applied Mathematics, 2006. doi:[10.1137/1.9780898718133.ch10](https://doi.org/10.1137/1.9780898718133.ch10).
- [20] P. Wesseling and C. W. Oosterlee. Geometric multigrid with applications to computational fluid dynamics. *Journal of Computational and Applied Mathematics*, 128(1-2):311–334, 2001. doi:[10.1016/S0377-0427\(00\)00517-3](https://doi.org/10.1016/S0377-0427(00)00517-3).
- [21] J. W. Ruge and K. Stüben. 4. Algebraic Multigrid. In *Multigrid Methods*, pages 73–130. Society for Industrial and Applied Mathematics, 1987. doi:[10.1137/1.9781611971057.ch4](https://doi.org/10.1137/1.9781611971057.ch4).
- [22] P. Vanek, M. Brezina, and J. Mandel. Convergence of algebraic multigrid based on smoothed aggregation.

- Numerische Mathematik*, 88(3):559–579, 2001. doi:[10.1007/s211-001-8015-y](https://doi.org/10.1007/s211-001-8015-y).
- [23] S. Badia, A. F. Martín, and J. Principe. Multilevel Balancing Domain Decomposition at Extreme Scales. *SIAM Journal on Scientific Computing*, 38(1):C22—C52, 2016. doi:[10.1137/15M1013511](https://doi.org/10.1137/15M1013511).
 - [24] A. Toselli and O. B. Widlund. *Domain Decomposition Methods — Algorithms and Theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. doi:[10.1007/b137868](https://doi.org/10.1007/b137868).
 - [25] L. Berger-Vergiat, H. Waisman, B. Hiriyyur, R. Tuminaro, and D. Keyes. Inexact Schwarz-algebraic multigrid preconditioners for crack problems modeled by extended finite element methods. *International Journal for Numerical Methods in Engineering*, 90(3):311–328, 2012. doi:[10.1002/nme.3318](https://doi.org/10.1002/nme.3318).
 - [26] B. Hiriyyur, R. Tuminaro, H. Waisman, E. Boman, and D. Keyes. A Quasi-algebraic Multigrid Approach to Fracture Problems Based on Extended Finite Elements. *SIAM Journal on Scientific Computing*, 34(2):A603—A626, 2012. doi:[10.1137/110819913](https://doi.org/10.1137/110819913).
 - [27] J. N. Jomo, F. D. Prenter, M. Elhaddad, D. D. Angella, C. V. Verhoosel, S. Kollmannsberger, J. S. Kirschke, E. H. V. Brummelen, and E. Rank. Robust and parallel scalable iterative solutions for large-scale finite cell analyses. *Arxiv*, pages 1–32, 2018.
 - [28] A. Menk and S. P. A. Bordas. A robust preconditioning technique for the extended finite element method. *International Journal for Numerical Methods in Engineering*, 85(13):1609–1632, 2011. doi:[10.1002/nme.3032](https://doi.org/10.1002/nme.3032).
 - [29] S. Badia and F. Verdugo. Robust and scalable domain decomposition solvers for unfitted finite element methods. *Journal of Computational and Applied Mathematics*, 344:740–759, 2018. doi:[10.1016/j.cam.2017.09.034](https://doi.org/10.1016/j.cam.2017.09.034).
 - [30] J. Mandel and C. R. Dohrmann. Convergence of a balancing domain decomposition by constraints and energy minimization. *Numerical Linear Algebra with Applications*, 10(7):639–659, 2003. doi:[10.1002/nla.341](https://doi.org/10.1002/nla.341).
 - [31] M. A. Heroux, E. T. Phipps, A. G. Salinger, H. K. Thornquist, R. S. Tuminaro, J. M. Willenbring, A. Williams, K. S. Stanley, R. A. Bartlett, V. E. Howle, R. J. Hoekstra, J. J. Hu, T. G. Kolda, R. B. Lehoucq, K. R. Long, and R. P. Pawlowski. An overview of the Trilinos project. *ACM Transactions on Mathematical Software*, 31(3):397–423, 2005. doi:[10.1145/1089014.1089021](https://doi.org/10.1145/1089014.1089021).
 - [32] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp, D. Kaushik, M. Knepley, D. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Web page. <http://www.mcs.anl.gov/petsc>, 2018.
 - [33] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp, D. Kaushik, M. Knepley, D. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Users Manual. Technical Report ANL-95/11 - Revision 3.10, Argonne National Laboratory, 2018.
 - [34] S. Badia, A. F. Martín, and F. Verdugo. Mixed aggregated finite element methods for the unfitted discretization of the Stokes problem. *SIAM Journal on Scientific Computing*, 40(6):B1541–B1576, 2018. doi:[10.1137/18M1185624](https://doi.org/10.1137/18M1185624).
 - [35] S. Badia, A. F. Martín, and J. Principe. FEMPAR: An Object-Oriented Parallel Finite Element Framework. *Archives of Computational Methods in Engineering*, 25(2):195–271, 2018. doi:[10.1007/s11831-017-9244-1](https://doi.org/10.1007/s11831-017-9244-1).
 - [36] C. Helzel, M. J. Berger, and R. J. Leveque. A High-Resolution Rotated Grid Method for Conservation Laws with Embedded Geometries. *SIAM Journal on Scientific Computing*, 26(3):785–809, 2005. doi:[10.1137/S106482750343028X](https://doi.org/10.1137/S106482750343028X).
 - [37] A. Johansson and M. G. Larson. A high order discontinuous Galerkin Nitsche method for elliptic problems with fictitious boundary. *Numerische Mathematik*, 123(4):607–628, 2013. doi:[10.1007/s00211-012-0497-1](https://doi.org/10.1007/s00211-012-0497-1).
 - [38] F. Kummer. Extended discontinuous Galerkin methods for two-phase flows: the spatial discretization. *International Journal for Numerical Methods in Engineering*, 109(2):259–289, 2017. doi:[10.1002/nme.5288](https://doi.org/10.1002/nme.5288).
 - [39] GAMG online documentnation. <https://www.mcs.anl.gov/petsc/petsc-current/docs/manualpages/PC/PCGAMG.html>.

- [40] Barcelona Supercomputing Center home page. <https://www.bsc.es/>.
- [41] C. Burstedde, L. C. Wilcox, and O. Ghattas. p4est: Scalable Algorithms for Parallel Adaptive Mesh Refinement on Forests of Octrees. *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011. doi:[10.1137/100791634](https://doi.org/10.1137/100791634).
- [42] O. Marco, R. Sevilla, Y. Zhang, J. J. Ródenas, and M. Tur. Exact 3D boundary representation in finite element analysis based on Cartesian grids independent of the geometry. *International Journal for Numerical Methods in Engineering*, 103(6):445–468, 2015. doi:[10.1002/nme.4914](https://doi.org/10.1002/nme.4914).
- [43] M. Olm, S. Badia, and A. F. Martín. On a general implementation of h - and p -adaptive curl-conforming finite elements. 2018.
- [44] M. S. Shephard. Linear multipoint constraints applied via transformation as part of a direct stiffness assembly process. *International Journal for Numerical Methods in Engineering*, 20(11):2107–2112, 1984. doi:[doi:10.1002/nme.1620201112](https://doi.org/10.1002/nme.1620201112).
- [45] S. Badia, A. F. Martín, E. Neiva, and F. Verdugo. A generic finite element framework on parallel tree-based adaptive meshes. *Arxiv*, 2019.
- [46] M. Bader. *Space-Filling Curves: An Introduction with Applications in Scientific Computing (Texts in Computational Science and Engineering)*. 2012.
- [47] G. Karypis. METIS and ParMETIS. In D. Padua, editor, *Encyclopedia of Parallel Computing*, pages 1117–1124. Springer US, Boston, MA, 2011. doi:[10.1007/978-0-387-09766-4_500](https://doi.org/10.1007/978-0-387-09766-4_500).
- [48] W. Bangerth, C. Burstedde, T. Heister, and M. Kronbichler. Algorithms and data structures for massively parallel generic adaptive finite element codes. *ACM Trans. Math. Softw.*, 38(2):14:1–14:28, 2012. doi:[10.1145/2049673.2049678](https://doi.org/10.1145/2049673.2049678).
- [49] S. Badia, A. F. Martín, and J. Principe. Implementation and scalability analysis of balancing domain decomposition methods. *Archives of Computational Methods in Engineering*, 20(3):239–262, 2013. doi:[10.1007/s11831-013-9086-4](https://doi.org/10.1007/s11831-013-9086-4).
- [50] MareNostrum4 User’s Guide. Technical report, Barcelona Supercomputing Centre, 2018.
- [51] R. Becker. Mesh adaptation for Dirichlet flow control via Nitsche’s method. *Communications in Numerical Methods in Engineering*, 18(9):669–680, 2002. doi:[10.1002/cnm.529](https://doi.org/10.1002/cnm.529).
- [52] J. Nitsche. Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 36(1):9–15, 1971. doi:[10.1007/BF02995904](https://doi.org/10.1007/BF02995904).
- [53] A. Massing, M. G. Larson, A. Logg, and M. E. Rognes. A Stabilized Nitsche Fictitious Domain Method for the Stokes Problem. *Journal of Scientific Computing*, 61(3):604–628, 2014. doi:[10.1007/s10915-014-9838-9](https://doi.org/10.1007/s10915-014-9838-9).
- [54] D. A. May, P. Sanan, K. Rupp, M. G. Knepley, and B. F. Smith. Extreme-scale Multigrid Components within PETSc. *Proceedings of the Platform for Advanced Scientific Computing Conference on - PASC ’16*, pages 1–12, 2016. doi:[10.1145/2929908.2929913](https://doi.org/10.1145/2929908.2929913).
- [55] J. Hu, R. Tuminaro, M. F. Adams, and M. Brezina. Parallel Multigrid Smoothing: Polynomial versus Gauss-Seidel. *Journal of Computational Physics*, 188(2):593–610, 2003.