AMG PRECONDITIONERS FOR LINEAR SOLVERS TOWARDS EXTREME SCALE[*]

PASQUA D'AMBRA[†], FABIO DURASTANTE[‡], AND SALVATORE FILIPPONE[§]

Abstract. Linear solvers for large and sparse systems are a key element of scientific applications, and their efficient implementation is necessary to harness the computational power of current computers. Algebraic MultiGrid (AMG) preconditioners are a popular ingredient of such linear solvers; this is the motivation for the present work where we examine some recent developments in a package of AMG preconditioners to improve efficiency, scalability and robustness on extreme scale problems. The main novelty is the design and implementation of a parallel coarsening algorithm based on aggregation of unknowns employing weighted graph matching techniques; this is a completely automated procedure, requiring no information from the user, and applicable to general symmetric positive definite (s.p.d.) matrices. The new coarsening algorithm improves in terms of numerical scalability at low operator complexity over decoupled aggregation algorithms available in previous releases of the package. The preconditioners package is built on the parallel software framework `PSBLAS`, which has also been updated to progress towards exascale. We present weak scalability results on one of the most powerful supercomputer in Europe for linear systems with sizes up to $O(10^{10})$ unknowns.

Key words. Algebraic MultiGrid, preconditioners, parallel scalability

AMS subject classifications. 65F08, 65F10, 65N55, 65Y05

1. Introduction. Solving algebraic linear systems of the form:

$$(1.1) \qquad\qquad Ax = b,$$

where $A \in \mathbb{R}^{n \times n}$ is a very large and sparse matrix and $x, b \in \mathbb{R}^n$ are vectors, is the most time consuming computational kernel in many areas of computational science and engineering, including more recent fields such as data and network analysis. Efficient and scalable methods and software libraries for solving such systems on the current generation of pre-exascale parallel computers are a key technology for high-resolution and high-fidelity simulations and analysis [3].

The Energy oriented Center of Excellence (EoCoE-II) is a project funded by the EU to enhance computing applications in the field of renewable energy generation and deployment to run on near future exascale parallel computers. In EoCoE-II, four out of the five flagship codes require efficient solution of systems of type (1.1) with sizes larger than 10^{10} to have reliable and predictive simulations of multi-physics phenomena; the size rapidly grows towards exascale as the models include more physics and ever more reliable mathematical formulations. Often the linear systems originate from time-dependent and/or non-linear models, so that the system (1.1) has to be solved multiple times within an iterative procedure; in this case the re-use of operators, such as preconditioners and/or Jacobian matrices, becomes crucial to reduce the overall time to solution.

[†]Institute for Applied Computing "Mauro Picone" (IAC), National Research Council (CNR), Napoli, Italy. (pasqua.dambra@cnr.it).

[‡]Department of Mathematics, University of Pisa, Pisa, Italy, and IAC-CNR. (fabio.durastante@unipi.it)

[§]Dept. of Civil Engineering and Computer Engineering, University of Rome Tor-Vergata, Rome, Italy, and IAC-CNR. E-mail salvatore.filippone@uniroma2.it

In this paper we will discuss our work on extensions of a package of Algebraic MultiGrid (AMG) preconditioners, motivated by EoCoE-II applications and built on top of the parallel software framework PSBLAS (Parallel Sparse Basic Linear Algebra Subroutines), which implements parallel basic linear algebra operations tailored for iterative sparse linear solvers [26, 28]. The main objective of the extensions was to improve efficiency, scalability and robustness of the preconditioners package, by employing novel AMG algorithms, to work better at extreme scale. The extensions' implementation also required a reworking of some basic kernels for parallel distributed matrix and data types management, for data communication, as well as the implementation of some additional Krylov and stationary iterative methods in the PSBLAS framework.

AMG methods are widely used as preconditioners in the iterative solution of large and sparse linear systems, in particular when the systems originate from the discretization of elliptic Partial Differential Equations (PDEs), on structured or unstructured meshes. Their success stems from their potential to achieve optimal complexity, i.e., a computational complexity that is linear in the size of the system, and on optimal convergence properties which ensure, in some cases, that the number of linear iterations to obtain a given accuracy remains almost constant for increasing system size [36, 40, 42]. The latter optimality property is often defined as *algorithmic or numerical scalability* [18], and is a necessary condition to have truly scalable AMG. Many more technological aspects have to be taken into account for efficient implementations on pre-exascale parallel machines, where many thousands of computational cores are distributed among parallel nodes connected by high-speed networks; often some nodes are equipped with accelerators exploiting fine-grained data parallelism and limited size high-speed memories, such as Graphics Processing Units (GPUs).

Since the introduction of AMG methods, substantial efforts have been put in the design and implementation of efficient packages for their use as parallel preconditioners. For a detailed description of the most important software projects aimed at preparing parallel linear solvers and preconditioners for the exascale challenge we refer the reader to [3]. For comparison purposes we consider here the Hypre library [24]; Hypre is very well known for its BoomerAMG multigrid preconditioner package [29], which includes many *de facto* standard AMG preconditioners, and is also interfaced by more comprehensive software frameworks, such as PETSc [1] and Trilinos [30]. In the early 2000s we proposed a package of parallel preconditioners, written in modern Fortran and exploiting basic concepts of object-oriented programming [15, 19]; previous versions of our package include AMG preconditioners based on the smoothed aggregation method introduced in [39]. In this paper we present the most recent activities for extending the package with new aggregation methods, smoothers and coarsest solvers; all these developments have been driven by the desire to achieve extreme scalability and flexibility with respect to the problem size and to the number and features of parallel processors.

The paper is organized as follows. In section 2 we describe the improvements in the PSBLAS framework for moving towards exascale. Section 3 introduces AMG preconditioners and presents the main features and computational kernels of the new parallel algorithm for coarsening; we focus on the convergence properties of the resulting AMG preconditioner for the different parallel smoothers available in the new version of the package. In section 4 we present the model problem and the parallel machines used in our study; section 5 provides an exhaustive analysis of the performance results. Finally section 6 addresses some concluding remarks.

2. PSBLAS towards exascale. Our development is based on the PSBLAS framework [28, 26, 16, 27]. Originally introduced for clusters that at the time were large-scale, it has gone through a number of revisions to keep up with the technology development of the past two decades, and the movement towards exascale is no exception. The software framework contains the computational building blocks for Krylov-type linear solvers on parallel computers, as well as support infrastructure to ease the writing of a parallel application using them. In particular, we introduced:

1. a framework for handling the mapping between the global index space of the problem and the local portions of the data structures [28];
2. the handling and optimization of the halo data exchange, also known as nearest-neighbour data exchange, the essential communication kernel;
3. an object-oriented architecture that enables choosing storage formats for sparse matrices and switching them at runtime to adapt to the application needs [26, 16];
4. a plugin for seamless integration of GPUs [16, 27].

During the development of the EoCoE project we have improved the handling of large index spaces requiring 8-byte integers, streamlined the process of setting up the data structures for halo data exchange, and also implemented some new computational kernels prompted by the extension of the preconditioners package.

In applications dealing with a large and sparse linear system, the system matrix is typically associated with a graph, examples being the discretization mesh of a PDE and the graph representing a complex network. All such applications handle the global numbering of the graph, which induces the global numbering of unknowns and matrix indices. In normal practice the global graph/matrix is partitioned and split among processes, and each portion local to a process is handled through a local numbering scheme. The solution in PSBLAS is to have an *index map* object contained in the *communication descriptor* to keep track of the correspondence between local and global indices. With the target of handling more than 10^{10} degrees of freedom (dofs), it is clear that global indices require 8-byte integers, but that does not necessarily mean that any individual portion will require the same; indeed, having the local portions of the matrices run over 4-byte integers enables memory savings that can be quite significant, especially when we consider accelerators such as NVIDIA GPUs which do not support virtual memory and for which memory management is a major concern for the developer. In the current development version of PSBLAS we can choose at configuration time the number of bytes for local and global numbering separately, with the default of using 4 bytes for local and 8 bytes for global indices.

One of the main design points of PSBLAS was to make it as easy as possible for the application developer to specify the distribution of the index space, with the only constraint that each global index/dof point is owned by one process; this is done at the time the descriptor for the index space is created. After this step, all processes need to figure out with whom they need to exchange data. In general there will be some mesh points whose value is needed to carry out the local part of the computation but are not locally owned, and are known as the *halo*; for each halo index, we need to know the owner process. This question would be easy to answer if we had available a vector mapping each index to a process; indeed, that is one of the possible ways to partition an index space, but for very large index spaces this would imply an excessive memory footprint. Instead, we normally keep an amount of auxiliary memory that is proportional to the number of local and halo indices on the current process, a solution scalable for increasing number of computational cores; this can be done in two main variants, with a set of hash tables, or by imposing the constraint that the

global indices owned by a process must be contiguous. Finding the process owning halo indices is equivalent to establishing a process topology mesh. To help with the construction of the data exchange lists:

1. we have devised a new iterative algorithm to identify the owner process for a given non-local index;
2. we defined a new interface for the user to provide additional information about the process topology, if available;
3. we create, when necessary, a copy of an existing index map employing a renumbering into a block-contiguous format, so as to speed up subsequent halo ownership identification.

The algorithm to identify index owners is based on the concept of *neighbouring processes*, i.e. processes that own indices needed by each other. The iterative algorithm alternates between probing, in which a global communication among processes establishes the ownership of a subset of the indices, and neighbour communication, in which the newly acquired neighbouring information is used to sweep through the remaining indices; the process is iterated until all index owners have been identified. Each of the communication steps has a limit on the amount of data to be exchanged, and potentially can be split into multiple sub-iterations; the data/memory limit can be adjusted by the user. In a block-contiguous numbering it is possible to determine index ownership by knowing the size of each contiguous block; this simplifies the operation, and therefore the library uses this numbering for auxiliary index maps internal to the preconditioner objects. All these aspects are handled internally by the software with minimal input by the user; they influence the setup time of the linear system and of the preconditioner, but have essentially no impact on the runtime of the solver methods.

3. AMG for PSBLAS. In [15, 19, 4] we proposed a package of AMG preconditioners built on top of the PSBLAS framework; the first version of the package implemented a multilevel version of some domain decomposition preconditioners of additive-Schwarz type and was based on a parallel decoupled version of the smoothed aggregation method described in [39, 37] to generate the multilevel hierarchy of coarser matrices. In this paper we will present a new version of the package, which inherits all the new features of the PSBLAS infrastructure and significantly extends the previous version in terms of algorithms and software modules. We are therefore improving flexibility, robustness and computational complexity, but we preserve the numerical scalability and concurrency of the preconditioners when tens of thousands cores are used, whilst at the same time including support for GPU accelerators. Given the significant number of changes and the increase in scope, we also decided to change the name of the package to AMG4PSBLAS. In the following we review the main features of the AMG methods, so that we can better describe the new parallel algorithms proposed in this paper and included in AMG4PSBLAS.

3.1. Introduction to AMG. In the sequel we will consider systems of the type (1.1) where A is symmetric and positive-definite (s.p.d.). An AMG method applied to such systems can be viewed as a particular instance of a general stationary iterative method:

$$x^{(k)} = x^{(k-1)} + B(b - Ax^{(k-1)}), \quad k = 1, 2, \ldots \text{ given } x^{(0)} \in \mathbb{R}^n,$$

where the matrix B is defined by recursion, as described in the following. Let A_l be a sequence of coarse matrices computed by the usual triple-matrix Galerkin product:

$$A_{l+1} = (P_l)^T A_l P_l, \quad l = 0, \ldots, nl - 1,$$

with $A_0 = A$ and P_l a sequence of prolongation matrices of size $n_l \times n_{l+1}$, with $n_{l+1} < n_l$ and $n_0 = n$. Let M_l be an A_l-convergent smoother for A_l, i.e., $\|I - M_l^{-1} A_l\|_{A_l} < 1$, where I is the identity matrix of size n_l and $\|\cdot\|_{A_l}$ indicates the A_l-energy-norm[1]. The preconditioner matrix B for the well known $V-$cycle, where 1 sweep of pre- and post-smoothing step is applied, is the linear operator corresponding to the multiplicative composition of the following error propagation matrices:

$$(3.1) \qquad I - B_l A_l = (I - (M_l)^{-T} A_l)(I - P_l B_{l+1}(P_l)^T A_l)(I - M_l^{-1} A_l) \quad \forall l < nl,$$

assuming that $B_{nl} \approx A_{nl}^{-1}$ is an approximation of the inverse of the coarsest-level matrix [40]. The main feature of AMG methods is that the preconditioner setup is completely algebraic, i.e. the prolongation matrices P_l and the smoothers M_l are defined using only information extracted from the corresponding matrices A_l. This is the main strength of AMG methods with respect to their geometric counterparts; indeed, in principle they may be applied to general linear systems not originating from PDEs. The AMG methods' convergence properties depend critically on the ability to define coarsening procedures and prolongation matrices which represent well at each new level the lower end of the spectrum of A. We refer the reader to [40, 42] for the general theory of AMG methods. Two main approaches are used to setup the hierarchy of prolongation matrices in a completely algebraic setting: *classical coarsening* and *coarsening by aggregation* [36, 42]. Classical coarsening separates the original index set into either coarse indices (C-indices), which form the coarse level, and fine indices (F-indices), whose unknowns will be interpolated by the C-indices values. The choice of interpolation formulas to transfer unknown values from coarser to finer levels defines the corresponding prolongation matrices. Coarsening by aggregation uses disjoint aggregates of fine unknowns to form the coarse unknowns and, in general, the prolongation matrices are piecewise-constant interpolation matrices (unsmoothed aggregation [33, 22, 12]) or a smoothed variant of them (smoothed aggregation [39]). In both cases, the way to select aggregates of fine level variables normally exploits heuristics measures of affinity (also known as *strength of connection*) among the variables; these measures have been constructed for systems arising from scalar elliptic PDEs, and often loose their robustness for more general systems. This paper deals with aggregation-based approaches; in particular, we present a parallel aggregation scheme exploiting maximum weight matching in the weighted adjacency graph of the sparse matrices at each level.

3.2. Parallel aggregation based on weighted graph matching. We present the first version of a parallel aggregation scheme for coarsening on large distributed-memory architectures. The method, named *coarsening based on compatible weighted matching* was first introduced in [22] and is already available in the sequential package described in [21]. A first parallel version of the method, exploiting fine-grained parallelism and specifically tailored for single GPU device is described in [7, 8]. The method is independent of any heuristics or a priori information on the *near kernel* of A, i.e., the lower part of the range of eigenvalues of the system matrix A which is generally

---

[1]Given a vector $\mathbf{v}$ and an s.p.d. matrix A, the A-energy-norm is defined as $\|\mathbf{v}\|_A = (\mathbf{v}^T A \mathbf{v})^{\frac{1}{2}}$.

used to obtain good-quality aggregates, and it is a completely automatic procedure applicable to general s.p.d. systems. Furthermore, the coupled coarsening based on compatible weighted matching has the advantage of building coarse matrices which are well-balanced among parallel processes; there is no need for special treatment of process-boundary dofs accounting for inter-processes coupling, as often happens in the coarsening procedures available in widely used software libraries. Finally, there is a significant flexibility in the choice of the size of aggregates: it is possible to have an almost arbitrarily aggressive coarsening.

The *coarsening based on compatible weighted matching* is a recursive procedure which starts from the adjacency graph $G = (\mathcal{V}, \mathcal{E})$ associated with the sparse matrix A, where the vertex set $\mathcal{V}$ consists of the row/column indices of A and the edge set $\mathcal{E}$ corresponds to the index pairs (i, j) of the nonzero entries in A. A *matching* $\mathcal{M}$ in the graph G is a subset of edges such that no two edges are incident on the same vertex. Graph matching is a general way to obtain partitions of graphs and was proposed for aggregation-based AMG methods in previous works [14, 13]. In our method we associate to the graph G a suitable edge weight matrix C, computed from the matrix A and an arbitrary vector w, and exploit the maximum product matching as a tool to obtain good quality aggregates for fast convergent AMG preconditioners.

Let C be the following weight matrix:

$$(3.2) \qquad (C)_{i,j} = c_{i,j} = 1 - \frac{2a_{i,j}w_i w_j}{a_{i,i}w_i^2 + a_{j,j}w_j^2},$$

where $a_{i,j}$ are the entries of A and $w = (w_i)_{i=1}^n$ is a given vector, and let $\mathcal{M}$ be a *maximum product matching* in the graph G with edge weight matrix C, i.e. $\mathcal{M} = \arg\max_{\mathcal{M}'} \prod_{(i,j)\in\mathcal{M}'} c_{ij}$. By applying a maximum product matching we can define the aggregates $\{\mathcal{G}_p\}_{p=1}^{n_p}$ for the row/column indices $\mathcal{I}$ of matrix A, consisting of pairs of indices, where $n_p = |\mathcal{M}|$ is the cardinality of the graph matching $\mathcal{M}$. Equivalently, we are decomposing the index set as

$$\mathcal{I} = \bigcup_{p=1}^{n_p} \mathcal{G}_p, \quad \mathcal{G}_p \cap \mathcal{G}_r = \emptyset \text{ if } p \neq r.$$

We observe that in case of sub-optimal matching, not all vertices will be endpoints of matched edges: we may have *unmatched* vertices. In this case, each unmatched vertex corresponds to a singleton G_s, and n_s is the total number of singletons. For each edge $e_{i\mapsto j} \in \mathcal{M}$ we can identify the vectors w_e as the orthonormal projection of the vector w on the aggregate G_p defined by the pair $\{i, j\}$. Given the w_e vectors, and an ordering of the indices which moves all the unknowns corresponding to unmatched vertices at the bottom, we can define a *prolongator*:

$$(3.3) \qquad P = \begin{pmatrix} \tilde{P} & 0 \\ 0 & W \end{pmatrix} \in \mathbb{R}^{n \times n_c},$$

where:

$$\tilde{P} = \text{blockdiag}(w_{e_1}, \ldots, w_{e_{n_p}}),$$

$W = \text{diag}(w_s/|w_s|)$, $s = 1, \ldots, n_s$, corresponds to unmatched vertices and $n_c = n_p + n_s$. The matrix P we have just built is a piecewise constant interpolation operator whose range includes, by construction, the vector w. The recursive application of the above procedure defines an unsmoothed-type aggregation coarsening whose quality and convergence analysis have been discussed in [20]; the analysis enables an a

posteriori evaluation of the quality of the aggregates and highlights the connection between the choice of the aggregates and the so-called compatible relaxation principle, originally introduced in [11] as a general way to obtain good quality coarsening in AMG. Good choices for the vector w are obtained by using *algebraically smooth vectors* with respect to A, i.e. performing a few iterations of an A-convergent smoother M on an arbitrary sample vector (see [20] for details). In the recursive application of the basic pairwise aggregation method, at each new level the input weighted graph $G = (\mathcal{V}, \mathcal{E}, C)$ corresponds to the adjacency graph of the computed coarse matrix whose weights are computed based on the restriction of the fine-level vector w onto the coarse space. It is possible to combine multiple steps of the basic pairwise aggregation by computing the product of m consecutive pairwise prolongators, to obtain a more aggressive coarsening; the resulting aggregates merge multiple pairs and have an almost arbitrary large size $n_c = 2^m$.

As already observed, the prolongation operator P we used for moving between coarse and fine levels corresponds to a piecewise constant interpolation operator. In such a case the V-cycle proves inadequate to obtain an optimal AMG; it is then necessary to employ more robust cycles such as general Algebraic Multilevel Iteration (AMLI) [40]. In this study we employ a Krylov-based MG cycle called the K-cycle, where at each level except the fine and the coarsest ones we apply two iterations of a Flexible Conjugate Gradient (FCG) method with the already defined AMG method starting on the current level as preconditioner [35].

An alternative to improve convergence while still employing a single V-cycle is to consider the use of a more regular interpolation operator obtained by applying one step of a weighted-Jacobi smoother to the basic piecewise constant interpolation, as in the smoothed aggregation AMG introduced in [39]. The actual prolongator $\overline{P}$ is then obtained from P as $\overline{P} = (I - \omega D^{-1} A)P$, where $D = \mathrm{diag}(A)$ and $\omega = 1/\|D^{-1}A\|_\infty \approx 1/\rho(D^{-1}A)$, with $\rho(D^{-1}A)$ the spectral radius of $D^{-1}A$. We will provide results for this choice in the experiments section.

The main kernels in the coarsening based on compatible weighted matching are therefore the computation of a maximum product matching in an edge-weighted graph, the sparse matrix-sparse matrix products needed to apply multiple steps of pairwise aggregation and to compute the Galerkin product, and, finally, the sparse matrix-vector product with matrix P^T needed to restrict the vector w from finer to coarser level.

The computation of a maximum product matching in an edge-weighted graph can be transformed into the classic *maximum weight matching* by maximizing the following additive weight function:

$$ - \sum_{(i,j)\in\mathcal{M}} (\log \max_i |c_{ij}| - \log |c_{ij}|), \ c_{ij} \neq 0. $$

The computation of an optimal maximum weight matching in a parallel setting is not a trivial task because the corresponding algorithm is intrinsically sequential. Indeed, practical approaches are based on a relaxation of the optimality requirement and look for solutions which approximate the optimal weight. In our parallel coarsening we use the `MatchBox-P` software library, which implements the parallel algorithm for the computation of half-approximate maximum weight matching described in [17]. This algorithm has a complexity $\mathcal{O}(|E|\Delta)$, where $|E|$ is the cardinality of the graph edge set and Δ is the maximum vertex degree[2], and guarantees a solution that is at

---

[2]The degree of a vertex is the number of edges that are incident on it.

least half of the optimal weight. The `MatchBox-P` algorithm is based on the idea of identifying *locally dominant* edges, i.e., edges with largest weight for both end-vertices. `MatchBox-P` supports data distributions arising from general row-block sparse matrix parallel distribution and implements message aggregation and overlapping between communication and computation to reduce the impact of the data communication on the parallel efficiency. It uses MPI asynchronous communication functionalities and was demonstrated to have good weak scalability properties for large graphs resulting from 2D PDE discretization on up to tens of thousands of parallel cores.

The construction of the multilevel hierarchy requires the computation of a number of parallel matrix-matrix products where the matrices are sparse. The implementation of this kernel can logically be split in two parts: the data exchange and the local computations. If we are computing the product $A \times B$, any given process will own a subset of the rows of A; to compute the product it will need not only the same subset of rows of B, but also the rows of B corresponding to the column indices of non-zero entries in A, including *halo* indices. The data exchange is thus essentially the same as the halo data exchange for vectors in a matrix-vector product, keeping in mind that items being exchanged are rows of sparse matrices, and in general the rows will be of differing lengths, and this must be accounted for in the use of the basic `PSBLAS` data structures for halo data exchange. An additional twist applies here: during the process of collecting remote matrix rows, the set of column indices on each process will in general expand, and hence there will be an impact on the support data structures. Once the matrix pieces have been converted to local numbering we can use any serial sparse-matrix by sparse-matrix product code to complete the computation.

3.3. Parallel Smoothers and Coarsest Solvers. Scalable AMG relies on smoothers that are both highly parallel and robust. A smoother is a convergent iterative method which is able to swiftly reduce the error components associated with the large eigenvalues of the system matrix. For s.p.d. matrices, we consider A-convergent smoothers, represented by a matrix operator M, such that $\|I - M^{-1}A\|_A < 1$. A common choice with a good smoothing factor is the Gauss-Seidel (GS) method, which is based on the convergent splitting of the system matrix $A = M - N$, with $M = L + D$ and $N = -L^T$, where $D = \text{diag}(A)$ and L is the lower triangular part of A. However, the GS method is intrinsically sequential; in parallel we typically use an inexact block-Jacobi version of it, where in the portion of the row-block local to each process the method acts as the GS method. This version is commonly known as *Hybrid GS (HGS)* and has been demonstrated to be a convergent smoother, with better smoothing properties than the block-Jacobi method when the local diagonal block of the matrix is sufficiently large with respect to the off-diagonal portion [5]. This dependence of the smoothing properties of HGS on the size of local diagonal blocks of the distributed matrix may have an important impact for AMG when very large numbers of parallel cores have to be employed. Indeed, for coarser-level matrices, where the size of local diagonal block is reduced and the matrix becomes denser, the off-diagonal part becomes heavier and the smoothing properties of HGS may degrade. In this setting a weighted version of the method, named ℓ_1-HGS, can be useful. ℓ_1-HGS is A-convergent and its smoothing properties appear less sensitive to the number of parallel cores and to the size of local diagonal blocks of the matrix, as discussed in [5]. Let A be divided in np blocks of size $nb \times n$, as in a general row-block parallel distribution involving np processes, and let A_{pp} be the corresponding diagonal block of A. Starting from the usual decomposition of the block as $A_{pp} = L_{pp} + D_{pp} + L_{pp}^T$, where $D_{pp} = \text{diag}(A_{pp})$ and L_{pp} is the lower triangular part of

A_{pp}, the ℓ_1-HGS method is an inexact block-Jacobi method defined by the following diagonal block matrix:

$$(3.4) \qquad M_{\ell_1-HGS} = \text{diag}((M_{\ell_1-HGS})_p)_{p=1,\ldots np},$$

where $(M_{\ell_1-HGS})_p = L_{pp} + D_{pp} + D_{\ell_1 p}$, with $D_{\ell_1 p}$ the diagonal matrix whose entries are:

$$(3.5) \qquad ((d_{\ell_1})_i)_{i=1}^{nb} = \sum_{j \in \Omega_p^{nb}} |a_{ij}|.$$

In the above definition Ω_p^{nb} is the set of indices of the p-th block of A outside the diagonal block A_{pp}. Note that the matrix blocks do not necessarily have the same size, i.e. nb can different for each of the np blocks. For $nb = 1$, ℓ_1-HGS becomes the so-called ℓ_1-Jacobi smoother which is $A-$convergent and generally has good smoothing properties. It is of particular interest for our aims, since its application allows to exploit the high level of parallelism of GPU accelerators, whereas the application of the ℓ_1-HGS method is inefficient due to need of apply the intrinsically sequential triangular solve involved in the inversion of the diagonal blocks in (3.4).

In the present study we also experiment with sparse approximate inverses since they do not require solving an auxiliary linear systems and have as their main kernel a sparse matrix-vector product, whose efficient implementation is available in the PSBLAS GPU plugin. There exist several different algorithms for computing a sparse approximate inverse; we focus here on the inversion and sparsification of an incomplete factorization introduced in [38]. This strategy is based on the application of a sparse inversion technique for the triangular factors of an existing incomplete factorization in the form $M = LDL^T$, where, as usual, D is a diagonal matrix and L is lower triangular with an all ones main diagonal. In this way an expression for

$$M^{-1} = L^{-T} D^{-1} L^{-1} = Z D^{-1} Z^T$$

is obtained, and the application of the smoother is reduced to the computation of a matrix-vector product. To have sparse expressions for the incomplete factorization of A^{-1} it is necessary to employ a sparsification process during the computation of the matrix M^{-1}, i.e., a sparsification process for the matrix Z. As discussed in [38], and analyzed in detail in [10], the sparsification can be based on either a thresholding procedure or a positional dropping; we refer to [9, Chapter 3.5] for a complete discussion. As for the other smoothers, we will always consider the case of computing the decomposition M for the diagonal blocks A_{pp} of A in a parallel block-Jacobi setting, by means of the positional sparsification technique, i.e., the INVK variant from [10, Algorithm 2].

We can also consider the ℓ_1-modification of this smoother by using again the modified block decomposition with the $D_{\ell_1,p}$ matrix in (3.5), that is, by computing the approximate inverse of the $A_{pp} + D_{\ell_1,p}$ blocks.

It is possible to define damped or weighted versions of all these methods, that is, to introduce the modified iteration

$$x^{(k+1)} = (I - \omega M^{-1} A) x^{(i)} + \omega M^{-1} b;$$

a sufficient condition for convergence is that $\omega \in (0, 2/\rho(M^{-1}A))$. When all the eigenvalues of $M^{-1}A$ are real and positive the best convergence ratio is obtained at the midpoint of the extreme eigenvalues $\omega_{\text{opt}} = 2/(\lambda_{\min}(M^{-1}A) + \lambda_{\max}(M^{-1}A))$.

Following the line of reasoning in [5], to evaluate the smoothing properties of the selected smoothers we can perform a smoothing analysis, for a fixed projector P, by looking at the spectral radius of the error propagation matrix in (3.1) for the two grid method using only a pre-smoothing step:

$$\|I - BA\|_A^2 = \|(I - P(P^T AP)^{-1} P^T A)(I - M^{-1}A)\|_A^2 \le 1 - \frac{1}{K};$$

for the convergence constant K it holds

$$K = \sup_e \frac{\|(I - PR)e\|_{\overline{M}}^2}{\|e\|_A^2} \ge 1, \qquad \overline{M} = M^T(M^T + M - A)^{-1}M,$$

where R is any matrix such that PR is a projection onto $\mathcal{R}ange(P)$. If we use the prolongator P in (3.3), we have that $R = (P^T DP)^{-1} P^T D$, with $D = \mathrm{diag}(A)$, defines the coarse grid variables, and is such that $RP = I_{n_c}$, with I_{n_c} the identity matrix of size n_c. We can then write the associated full-rank matrix $S \in \mathbb{R}^{(n-n_c) \times n}$ such that $RS = 0$ by the D-orthogonal vectors to the diagonal blocks of P. Finally, we can explicitly compute K from [25]:

$$(3.6) \qquad K = \frac{1}{\lambda_{\min}\left[(S^T \overline{M}S)^{-1}(S^T AS)\right]},$$

and use it to compare the smoothing properties for different choices of M while keeping fixed the remaining parts of the hierarchy. In Table 1 we report the constant K relative to the application of the parallel smoothers we have discussed in this section in a two-level method built by our matching-based parallel coarsening, for a 3D Laplace equation with homogeneous Dirichlet boundary conditions on a unit cube.

We can observe that all the smoothers considered here have good smoothing properties, leading to good convergent two-level AMG. HGS appears to have the best behavior, with no significant deterioration for decreasing block size. HINVK also produces good convergence constants for decreasing block size. Using the optimal value for ω, as expected, improves the convergence, however using $\omega = 1$ does not degrade K in a significant way. Finally, we note that ℓ_1 versions of the smoothers do not seem to improve K: the constant value is generally worse than that of the corresponding standard smoother.

In `AMG4PSBLAS` we provide interfaces to some of the widely used parallel direct solvers, such as `SuperLU` [32] and `MUMPS` [2]. However, using direct solvers at the coarsest level of an AMG method on many thousands of parallel cores can be very expensive, because the coarsest-level matrix tends to have a small size, therefore the cost of data communication dominates the local arithmetic computations. Sometimes, as in the `Hypre` library, the coarsest-level system is solved by a subset of parallel cores, leaving many processes idle, then the solution is redistributed to all the processes in order to proceed. In this study we focus on a distributed coarsest solver involving all the parallel cores and limiting the maximum size of the coarsest-level matrix to a sufficiently large size. In order to have an effective parallel solution of the corresponding coarsest system we exploit a preconditioned version of the Conjugate Gradient (CG) Method, coupled with block-Jacobi preconditioners, as available in `PSBLAS`.

4. Model Problem and Computational Environments. To benchmark our code, we consider the solution of the Laplace equation in 3D, i.e., the self-adjoint elliptic partial differential equation of the form

$$(4.1) \qquad -\nabla \cdot (K \nabla u) = \mathbf{f}, \quad \text{in } [0,1]^3,$$

Table 1: Convergence constant K from (3.6) for various smoothers for the homogeneous 3D Laplacian problem with parallel coarsening based on compatible weighted matching. The global size of the problem is $m = 2^4 \times 2^4 \times 2^4$ distributed over np processes using a 3D block distribution. By INVK we denote here the approximate inverse computed from the incomplete LU factorization with 0 levels of fill-in, and admitting a single level of fill-in in the inversion procedure.

			Unsmoothed prolongator with 3 sweeps of pairwise aggregation				
m	np	HGS	ℓ_1-HGS	HINVK	ℓ_1-INVK	HINVK	ℓ_1-INVK
				$\omega = 1$		ω_{opt}	
4096	1	1.3766	1.3766	1.5562	1.5562	1.0024	1.0024
2048	2	1.4194	1.5270	1.5273	1.7196	1.0356	1.0812
1024	4	1.4587	1.6621	1.6093	2.2149	1.2593	1.3638
512	8	1.4744	1.7803	1.8284	2.6713	1.5345	1.6147
256	16	1.4945	1.8230	1.8608	2.7307	1.5589	1.5816
128	32	1.5149	1.8682	1.8977	2.7972	1.5875	1.5669
64	64	1.5335	1.9162	1.9390	2.8715	1.6198	1.5736
32	128	1.5880	2.0343	2.0272	3.0707	1.7469	1.6461
16	256	1.6406	2.1594	2.1440	3.3688	1.9108	1.7678
8	512	1.6665	2.3088	2.3137	3.7280	2.1317	1.9254
l_1-Jacobi				5.6220			

with homogeneous Dirichlet boundary conditions, $K = 1$ and unitary right-hand side. The discretization of this problem is obtained by the usual 7-points finite difference stencil.

The choice of the above model problem satisfies two requirements: on the one hand it allows for comparisons with results available in the literature [34, 6], and on the other hand it is of interest because it is representative of the computational kernel at the core of two of the flagship applications included in the EoCoE-II project. Indeed, the Large Eddy Simulation of wind-turbines models in Alya [41], and the simulation of 3D hydrologic models, simulating subsurface flows based on the solution of Darcy's equation in Parflow [31], both reduce to elliptic equations similar in nature to (4.1).

For the performance results in section 5 we investigate weak scalability: we aim to show that our solvers scale (almost) linearly with the amount of resources. For this task we consider two computational frameworks. The first one is a purely distributed context in which the resources to be considered for the weak scaling are the number of MPI ranks. The second one is a hybrid context in which we exploit both the distributed MPI architecture and the use of GPU accelerators; therefore the resources in this case are CUDA devices, i.e., the number of accelerators. In both cases, the resulting linear system is distributed by means of a block 3D distribution: the MPI ranks are arranged in a Cartesian grid similar to the cubic grid induced by the discretization, and the unknowns in each discretized sub-block are then assigned to the corresponding process. The weak scalability tests is then performed by a number of subsequent refinements of the mesh in such a way that, as the number of processes is increased, the number of unknowns per process is constant.

The machine we used for our experiments is the **Piz Daint** supercomputer from

the Partnership for Advanced Computing in Europe (PRACE). Piz Daint is operated by the Swiss National Supercomputing Center. It is based on the Cray Model XC40/Cray XC50 architecture with 5704 hybrid compute nodes (Intel Xeon E5-2690 v3 with Nvidia Tesla P100 accelerator), and 1431 multicore compute nodes (Intel Xeon E5-2695 v4), using the Cray Aries routing and communications ASIC with Dragonfly network topology. The gross total of cores is 387,872 and it is ranked at the twelfth position in the November 2020 *TOP500*[3].

5. Performance Results. To measure the overall performance of the parallel AMG preconditioners introduced in this paper we focus on both *algorithmic* and *implementation* scalability. Perfect algorithmic scalability is achieved when $\rho(B^{-1}A) \approx 1$ independently of the global size n of the linear system; implementation scalability on the other hand corresponds to having an optimal application cost for B^{-1} that is $\mathcal{O}(n)$ flops per iteration while achieving parallel speedup proportional to the number of processes employed. In all cases we use the AMG methods as preconditioners for a Flexible Conjugate Gradient (FCG) algorithm. Algorithmic scalability can be analyzed by looking at the number of iterations needed by FCG to achieve a relative residual norm of 10^{-6} as the size n grows; implementation scalability is analyzed by considering both the total solve time for the procedure as well as the average time per iteration. We also analyze the timings for the preconditioner setup with increasing number of processes; all the timings reported in the figures are in seconds.

Regarding the memory footprint of the proposed multigrid hierarchies as well as the cost of the application of a V-cycle, a quantitative measure is given by the operator complexity

$$\text{opc} = \frac{\sum_{l=0}^{\text{nl}-1} \text{nnz}(A_l)}{\text{nnz}(A_0)} > 1.$$

Note that for both AMG cycles we employed (K- and V-cycle), we always apply 1 pre-smoothing and 1 post-smoothing step. As coarsest solver we use a version of the CG method coupled with a block-Jacobi preconditioner, with incomplete LU factorization (ILU(0)) on the diagonal blocks; the iterative solution of the coarsest system is stopped when the relative residual is less than 10^{-4} or the number of iterations is larger than 30. In all experiments discussed in this paper we applied the parallel coarsening based on compatible weighted matching starting from a vector w of all ones, since this vector is in the near kernel of our model test case. Furthermore, we used the library default approach which sets a target size of the coarsest-level matrix and stops the coarsening procedure as soon as the coarse matrix size is less than or equal to the target; in the experiments with the pure MPI version of the library we set $maxsize = 200 \times np$, where np is the number of cores.

To identify the different algorithmic variants implemented in `AMG4PSBLAS` and analyzed in our experiments, we use the labeling convention described in Table 2: the overall name reported in the figures and in the analysis is obtained by combining together the labels of the various components, e.g., KA1S1CS1 is the preconditioner employing a K-cycle, a hierarchy built with three sweeps of parallel coarsening based on matching and unsmoothed prolongator (unsmoothed parallel matching, for brevity), one sweep of the forward/backward Hybrid Gauss-Seidel smoother and the preconditioned CG method as coarsest solver.

5.1. Performance of the smoothers. In this section we analyze performance results obtained by using the smoothers discussed in Table 1 on CPU cores of Piz

---

Table 2: Strings identifying each preconditioner are built by combination of the strings identifying the various algorithmic variants.

```
        ┌──────────────────────┐
        ↓  K  A1  S2  CS1  ◄───┤
        └──────────↑───↑───────┘
```

Cycle	**A**ggregation	**S**moother	**C**oarsest **S**olver
K	1) Unsmoothed Parallel Matching: 3 sweeps	1) Hybrid Gauss-Seidel	1) Preconditioned CG
V	2) Unsmoothed Parallel Matching: 4 sweeps	2) l_1–Hybrid Gauss-Seidel	Hypre **C**oarsest **S**olver
	3) Smoothed Parallel Matching: 3 sweeps	3) INVK	2) Direct Solver
	4) Smoothed Parallel Matching: 4 sweeps	4) l_1–INVK	
	5) VBM	5) l_1–Jacobi	

Hypre **C**oarsening

1) Falgout

2) HMIS

3) HMIS1

Daint. We compare the different versions of the HGS smoothers and of HINVK smoothers (with $\omega = 1$) within the K-cycle coupled with A1 aggregation (see Table 2); we refer to this type of preconditioners as KA1 types. We consider weak scalability with 256×10^3 unknowns (dofs: degrees of freedom) per core and up to 16384 cores, i.e., we reach an overall number of $\sim 4.2 \times 10^9$ dofs.

In Fig. 1a we observe that for all smoothers we obtain a good algorithmic scalability, with the number of iterations ranging from 11 to 18. The HGS smoother is generally a little better than the $\ell1$-HGS, which sometimes requires one or two additional iterations. For a wide range of number of cores, $\ell1$-HINVK requires fewer iterations than HGS, but when the largest core count is used HGS is better. In Fig. 1b we report the solve time per FCG iteration for the different preconditioners. We observe that all preconditioners show a very good implementation scalability with a maximum time per iteration of 0.22 seconds on 16384 cores, corresponding to a solve time of 5×10^{-11} seconds per dof.

HGS shows a consistently good behavior with increasing number of cores which is reflected in the total solve time, as shown in Fig. 2a. In Fig. 2b we show the setup time for the different preconditioners, including the setup of the AMG hierarchy, the setup of the smoother operators at each level of the hierarchy, and the setup of the coarsest-level solver. The setup of the preconditoners with the HINVK smoothers requires a somewhat larger time; this is due to the additional (local) computations needed to perform the approximate inversion, whereas the HGS smoothers only require memory copies to generate the matrix splitting. The smoother setup time is completely flat, since it only depends on the size of the local matrix, whereas the hierarchy setup time tends to grow with the number of cores/size of the systems, hence the gap between
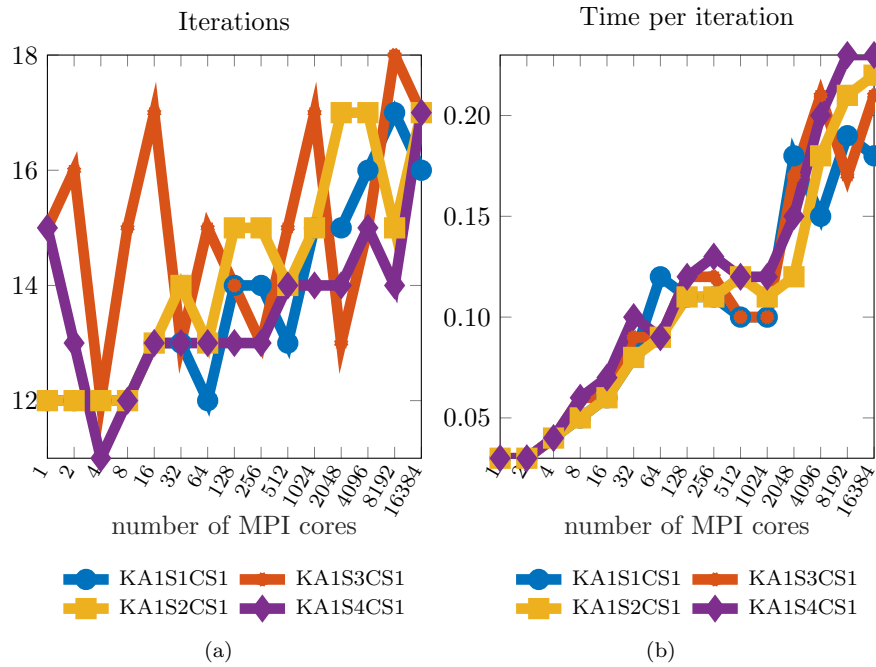
Fig. 1: Weak scaling results 256k dofs per core. Number of iterations and time per iteration for different smoothers when KA1-type preconditioners is used.
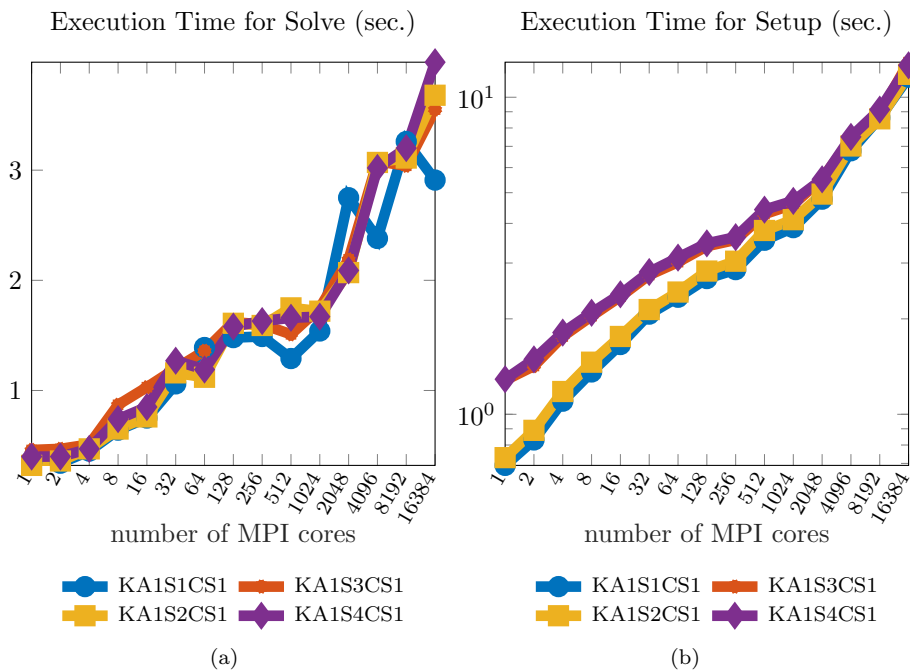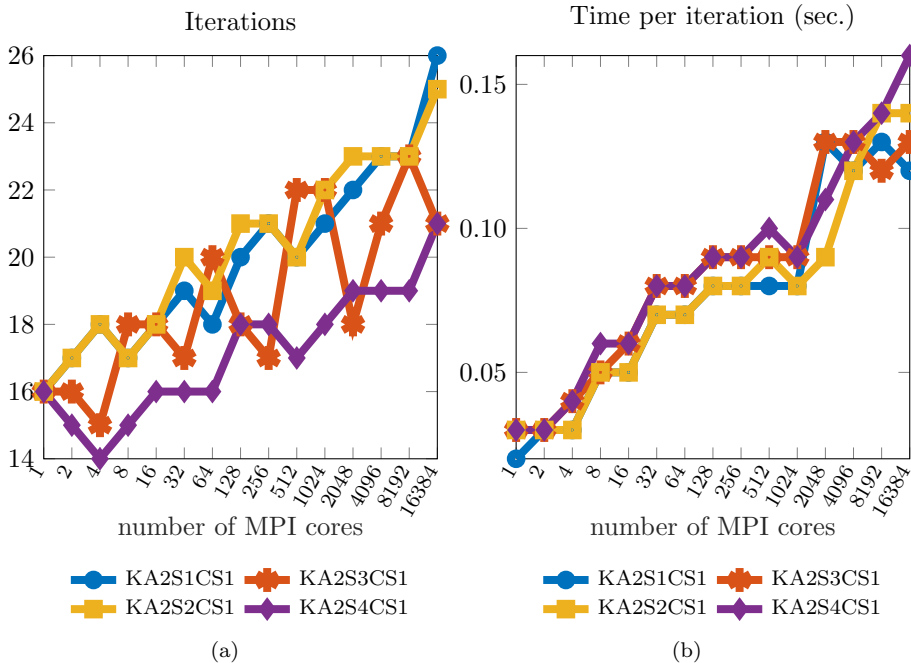

Fig. 2: Weak scaling results 256k dofs per core. Execution times for the solve and setup for different smoothers when KA1-type preconditioners are used.

the curves for the total setup time in Fig. 2b tends to close.

In all cases the setup times show good scalability with a sub-linear increase for increasing number of cores and dofs. We did the same analysis with a more aggressive coarsening using 4 sweeps of basic pairwise aggregation, i.e., when the size of aggregates is at most 16, which we denote as KA2-type preconditioners.   In Fig. 3a-3b

Fig. 3: Weak scaling results 256k dofs per core. Number of iterations and time per iteration for different smoothers when KA2-type preconditioners are used.

we report number of iterations and solve time per iterations, respectively, when the different preconditioners are applied. As expected, in this case the number of iterations is generally larger than the case of KA1-type preconditioners, however we still get a fairly good numerical scalability, with a moderate increase in the number of iterations for increasing number of cores. The best convergence behavior is obtained when HINVK smoothers are employed, while the minimum number of iterations is generally obtained by $\ell1$-HINVK. This good convergence behavior balances the small increase in solve time per iteration required by the HINVK smoothers, resulting in a total solve time which is generally better when $\ell1$-HINVK is employed, except on 16384 cores, where HINVK and HGS obtain a slightly better solve time (see Fig. 4a). For the sake of completeness we report in Fig. 4b the setup time for the different preconditioners based on the more aggressive coarsening, showing a behavior similar to that of the KA1-type preconditioners.

5.2. Comparison with Hypre. We now compare our preconditioners with some of those available in the Hypre library [24], using default algorithmic parameters for best practice. Specifically, we compare with three different coarsening approaches, i.e., Falgout [29], [43, section 3.2], HMIS [23], and HMIS1 [44, section 3]; we re-

Fig. 4: Weak scaling results 256k dofs per core. Execution times for the solve and setup for different smoothers when KA2-type preconditioners are used.

fer to section 2 in the supplementary materials for a more extensive discussion of their features. All these coarsening schemes are used to generate an AMG hierarchy which is applied as a V-cycle preconditioner for a CG method, with one sweep of forward/backward HGS being applied as pre/post-smoother. Default choices are used for the coarsest system, where a direct method is employed; using the same notation applied for our preconditioners, the three `Hypre` preconditioners are denoted, respectively, VC1S1CS2, VC2S1CS2, and VC3S1CS2.

We compare these preconditioners from `Hypre` with the KA1S1CS1 and VA3S1CS1 methods in `AMG4PSBLAS` and with the VA5S1CS1 preconditioner, based on the parallel decoupled version of the smoothed aggregation strategy from [39, 37], implemented in the previous version of the library [15, 19], i.e., `MLD2P4` v. 2.2.

For these experiments we used the Piz Daint machine up to 8192 CPU cores. We detail first, in Fig. 5, the operator complexity obtained by the different coarsening approaches; we can observe that the new strategy based on the A1 aggregation consistently produces the smallest operator complexity, thus reducing the memory footprint for the preconditioner. On the other hand, the smoothed version A3 produces a small increase in the operator complexity, whilst still remaining much smaller than VC1S1CS2 and VC2S1CS2 and smaller than 2, therefore ensuring a linear cost in the application of the preconditioner. Note that our VA3S1CS1 shows the best algorithmic scalability, requiring a number of iterations which is almost constant (between 7 and 9) for increasing number of MPI cores except on 2048 where a small increase to 12 is observed. This is essentially due to a small decrease in the ability of
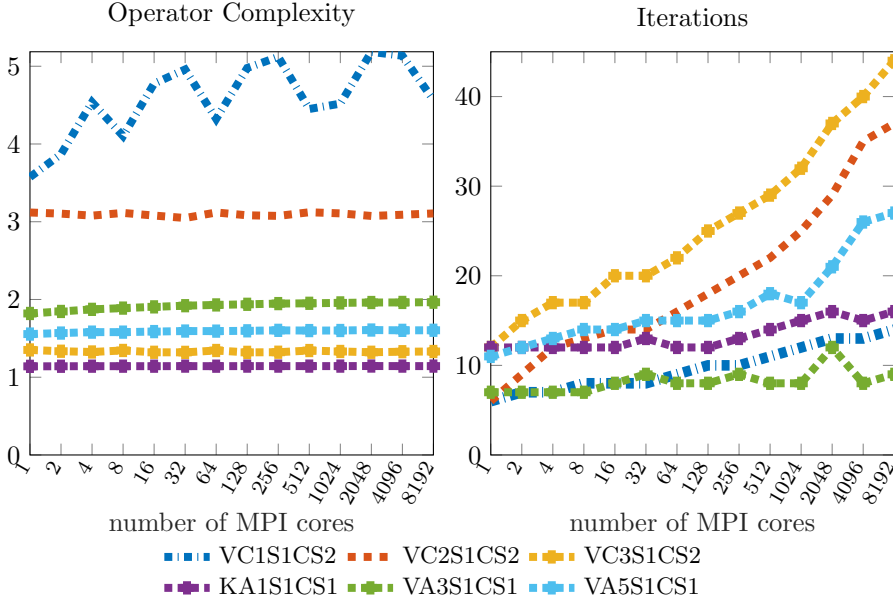
Fig. 5: Comparison with `Hypre` 256k dofs per core. Operator complexity and number of iterations for different preconditioners.

obtaining a matching of maximum cardinality, hence producing a small deterioration of the quality of the final aggregates at the coarser levels during the hierarchy setup. This aspect is currently under investigation: we are studying new parallel approximate matching algorithms aiming to realize the best trade-off between maximum weight and cardinality in a way that is independent of both problem size and data partitioning. VC1S1CS2 shows a comparable number of iterations ranging from 6 to 14: this is the benefit accrued by the large operator complexity of the corresponding AMG hierarchy. On the other hand, KA1S1CS1 also shows a very good algorithmic scalability, with number of iterations ranging from 12 to 16, especially considering the fact that both operator complexity and memory requirements are about 5 times smaller than those of VC1S1CS2. All the other variants show a more significant increase in the number of iterations with respect to the problem size. To complement this information we look also at the setup time for all the preconditioners. From the results in Fig. 6 we observe that using KA1S1CS1 tends to have a similar cost to that of VC1S1CS2, while VA3S1CS1 shows a small increase to the setup time due to the application of 1 step of the weighted Jacobi smoother to the hierarchy prolongators, as described in Sec. 3.2. VA5S1C1, based on a decoupled smoothed aggregation, has a clear advantage in the setup cost due to the absence of communication in the aggregation algorithm, and, as expected, obtains the best speedup for the setup phase. Very good speedups are also shown by KA1S1CS1 and VA3S1CS1, which confirms the effectiveness of the parallel implementations of all the computational kernels described in sections 3.2-3.3. If we look at the solve phase in Fig.7, we can see that VA3S1CS1 has generally the best solve time with respect to the other preconditioners, with a very small increase with increasing number of cores. This efficiency compensates the small increase in the setup time with respect to the `Hypre` preconditioners when applied to
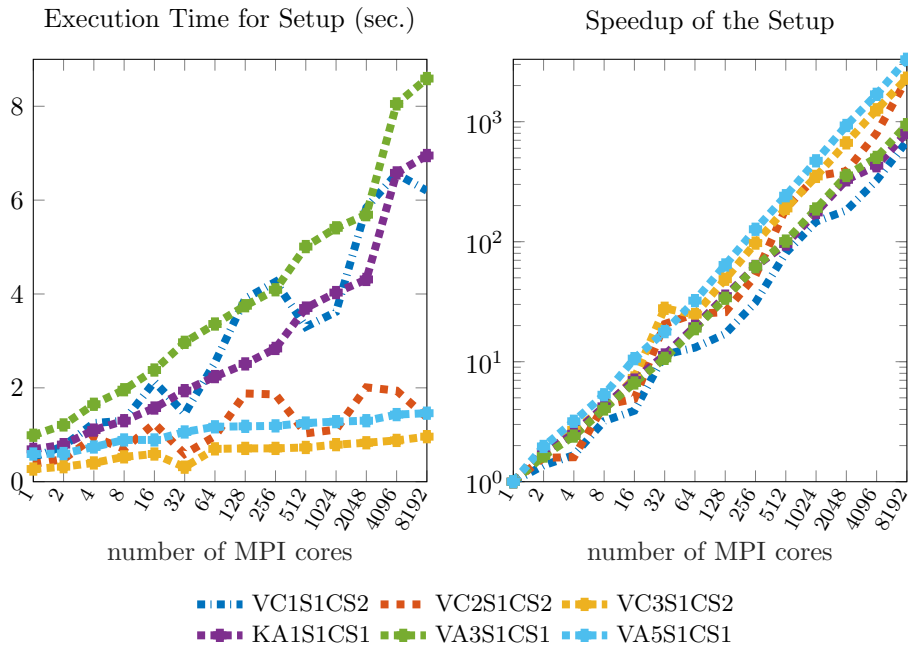
Fig. 6: Comparison with `Hypre` 256k dofs per core. Preconditioners setup: execution time (left), speedup (right).
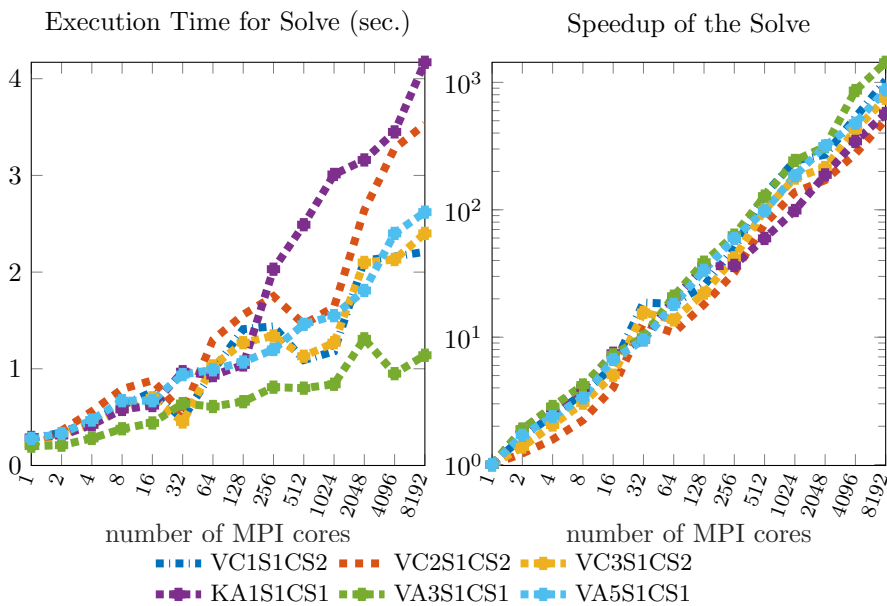

Fig. 7: Comparison with `Hypre` 256k dofs per core. Solve: execution time (left), speedup (right).

time-dependent or non-linear problems, where the same preconditioner can be reused for multiple external iterations. We observe that our KA1S1CS1, although showing a very good algorithmic scalability, has a rapid increase in the solve time for increasing number of cores, due to the K-cycle application. Indeed this cycle requires $2^{nl-2} + 1$ coarsest level visit and solutions per iteration, leading to a worse ratio between computation and communication with respect to the simpler V-cycle. Therefore, when the number of levels grows with the problem size so as to keep the coarsest matrix reasonably small, the above approach is poorly scalable. The speedups of the solve phase for all preconditioners are broadly comparable.

We conclude this comparison by claiming that our new coarsening based on weighted matching coupled with a smoothed version of the transfer operators, and with a robust iterative solver at the coarsest level shows very good scalability results and outperforms the preconditioners included in `Hypre` both in terms of memory requirements and in terms of solve times. Algorithmic scalability for Poisson-type problems of VA3S1CS1 appears optimal and comparable with the standard geometric multigrid approach.

5.3. Performance results towards extreme scale. In this section we discuss scalability results obtained on Piz Daint, running tests with 512×10^3 dofs per core up to 27000 cores, i.e., we reach an overall number of $\sim 1.4 \times 10^{10}$ dofs; in the same vein, we also analyze results obtained with GPU accelerators (in the solve phase). In the latter case we run tests with $12 \times 512 \times 10^3 \sim 6.2 \times 10^6$ dofs per GPU, and up to 2048 GPUs, i.e, we reach an overall number dofs of more than 1.2×10^{10}; in these experiments on GPUs we keep the same amount of memory per node, hence each GPU will handle the same number of dofs as 12 CPU cores. In the same vein, we stop the coarsening process for the setup of the multilevel hierarchy when the maximum size of the coarsest matrix is $12 \times 200 \times np$, where np is the number of GPUs, i.e. with the same size of coarsest matrix per node.

We begin with results on pure MPI; for the sake of space, we limit our discussion to preconditioners using the HGS smoother. We compare the KA1S1CS1 and KA2S1CS1 preconditioners (KA1/A2-types), using the K-cycle in the application of an AMG hierarchy where unsmoothed prolongators are employed, with the VA3S1CS1 and VA4S1CS1 preconditioners (VA3/A4-types), which use the V-cycle coupled with the smoothed version of the prolongators. In Fig. 8a we show the operator complexity of the multilevel hierarchies corresponding to all preconditioners. As expected, the operator complexity of the VA3/A4-type preconditioners, with smoothed prolongators, is larger than that of the corresponding KA1/A2-type preconditioners, since the coarse matrices are denser than those built with the unsmoothed prolongators. Nevertheless, when aggregates of size 8 are built, the operator complexity is about 1.9, while when aggregates of size 16 are employed, the operator complexity is about 1.3. This indicates that even for VA3/A4-type preconditioners, the memory requirements for the AMG hierarchies are less than double the memory needed for the system matrix. Moreover, despite the small operator complexity, the numerical scalability of all the preconditioners is very satisfactory, even optimal for some of them; in Fig. 8b we see that KA1S1CS1 requires a number of iterations ranging from 12 to 17, while KA2S1CS1 requires a number of iterations ranging from 17 to 24, showing a small increase in iterations despite a reduction in operator complexity, and preserving numerical scalability for increasing number of cores. The VA3/A4-type preconditioners, employing the smoothed version of the prolongators and a less expensive cycle, require a smaller number of iterations than the corresponding KA1/A2-type preconditioners.
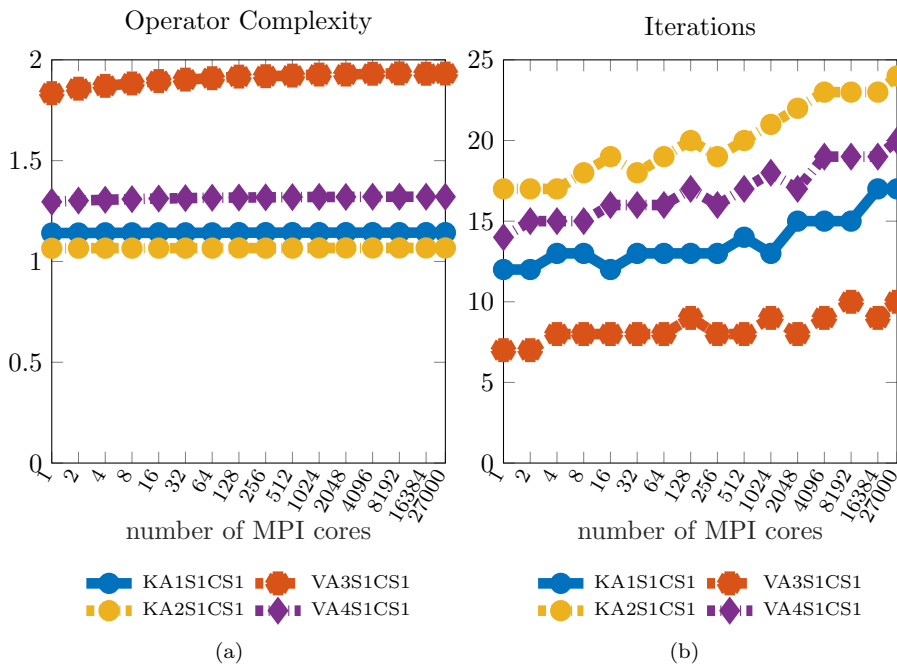
Fig. 8: Weak Scaling results for 512k dofs per core. Operator complexity and number of iterations.

In more details, VA3S1CS1 continues to show an almost perfect algorithmic scalability with a number of iterations ranging from 7 to 10 going from 1 core to 27700. A very good algorithmic scalability is also observed for VA4S1CS1, where the number of iterations ranges from 14 to 20. This behavior confirms that the new parallel coarsening based on weighted matching is able to detect, in a completely automatic way, good quality aggregates of variable size and, exploiting smoothed operators allows to obtain optimality also using V-cycle at low operator complexity. VA3S1CS1 always obtains the best solve time per each number of cores, with a very slow increase for increasing number of cores; it solves the biggest size problem on 27000 cores in less than 2 seconds. In Fig. 9a the application of the more expensive K-cycle shows its effect on the solve time, especially when increasing the number of cores. Better solve times than the KA1/A2-type preconditioners are also obtained by VA4S1CS1 for increasing number of cores, showing that the use of V-cycle coupled with A3/A4 aggregation types is the best choice for extreme scalability, albeit at a small increase in the setup cost (see Fig. 9b). For the sake of completeness we report in Figs. 9d-9c the speedups obtained; we observe a very similar behavior for all methods, displaying a smooth increase with increasing number of cores, and demonstrating the good implementation scalability of all computational kernels.

The GPU plugin of PSBLAS allows us to run the solve phase on a cluster of GPUs, implementing the preconditioned FCG method coupled with some of the preconditioners included in AMG4PSBLAS. The plugin implements efficient GPU versions of the sparse matrix-vector products and vector-vector operations, such as vector updates and scalar products, including the necessary MPI communications [16, 27]. In the
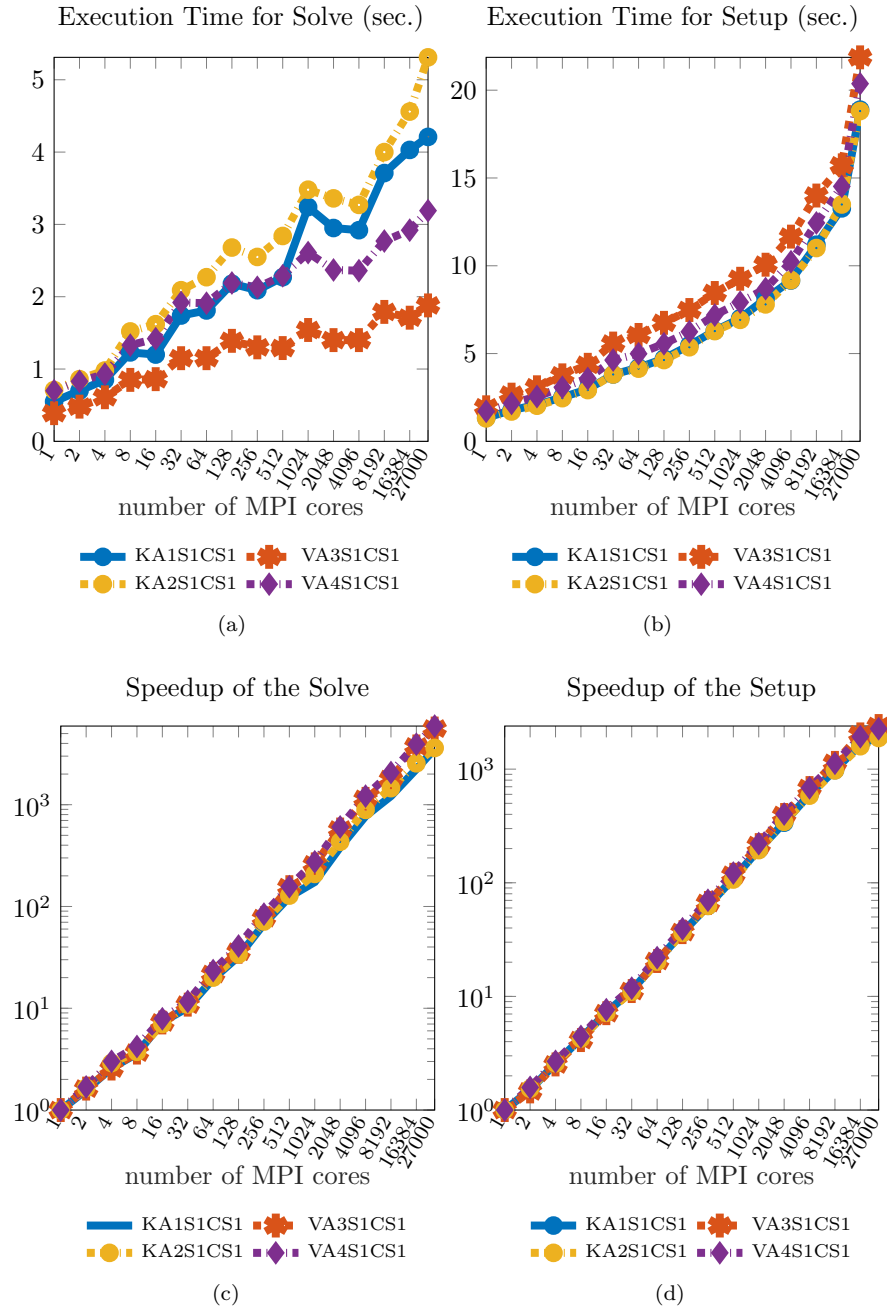
Fig. 9: Weak scaling results 512k dofs per core. Setup and solve time with the relative speed-up.

following we discuss results obtained by employing VA3-type preconditioners coupled with both the HINVK and ℓ_1-Jacobi smoothers, whose implementation on GPU users the sparse matrix-vector kernel. We compare VA3S3CS1 and VA3S5CS1, when they are coupled with the FCG iterative solver; in the case of VA3S3CS1, we apply HINVK also as preconditioner of the CG method at the coarsest level, while for VA3S5CS1 we use ℓ_1-Jacobi both as smoother and as preconditioner for CG at the coarsest level. Note that, due to the larger constant value of the ℓ_1-Jacobi smoother (see Tab. 1), we apply 4 pre/post-smoothing sweeps of the method at each V-cycle application, whereas for HINVK we only apply 1 sweep. From Fig. 10a we see that both methods have a very similar behavior, showing a number of iterations ranging from 7 to 23 for VA3S3CS1 and from 7 to 26 for VA3S5CS1. The cost per iteration of VA3S3CS1 is generally better, as shown in Fig. 10b: we can observe a cost per iteration ranging from 0.04 to 0.14 seconds, corresponding to a solve time per dof in the range $[10^{-12} : 10^{-9}]$.
  For the total solve time reported in Fig.11a, we observe that the two preconditioners
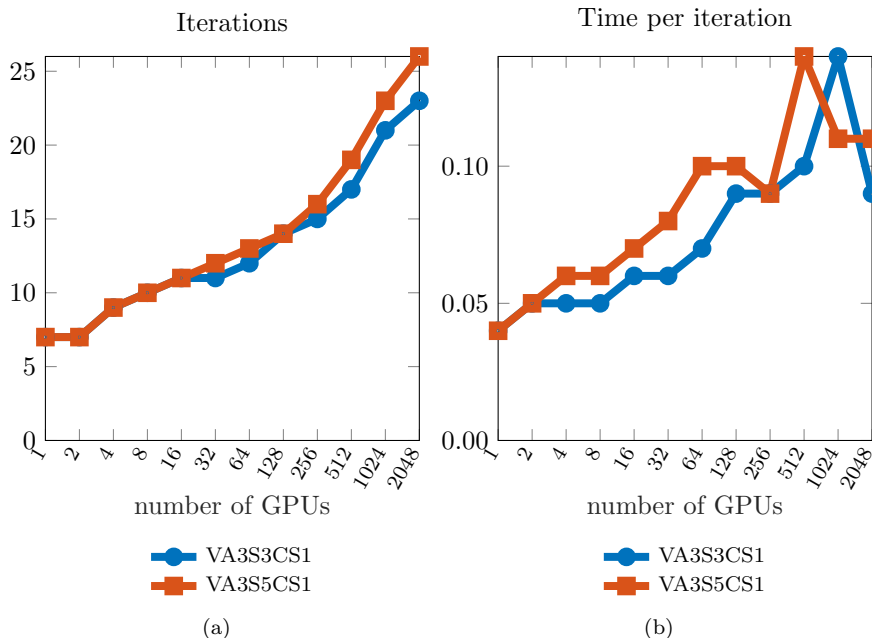

Fig. 10: Weak scaling results 6M dofs per GPU. Number of iterations and time per iteration on GPUs.

have a similar behavior but VA3S3CS1 is generally better. In Fig. 11b we also show the setup time for both preconditioners. In the current version of the library the setup of the preconditioners is not yet implemented on the GPU: indeed, it is carried out on a single core of the CPU host device. We see that, as expected, VA3S3CS1 shows larger setup costs with respect to VA3S5CS1 due to the larger cost for the setup of the HINVK smoother. The smoother setup cost, as expected, is about constant for all numbers of CPU cores, with the ℓ_1-Jacobi smoother being less expensive by about 2 orders of magnitude. Further development activities planned for future releases of the library include the implementation of hybrid OpenMP-MPI and possibly CUDA
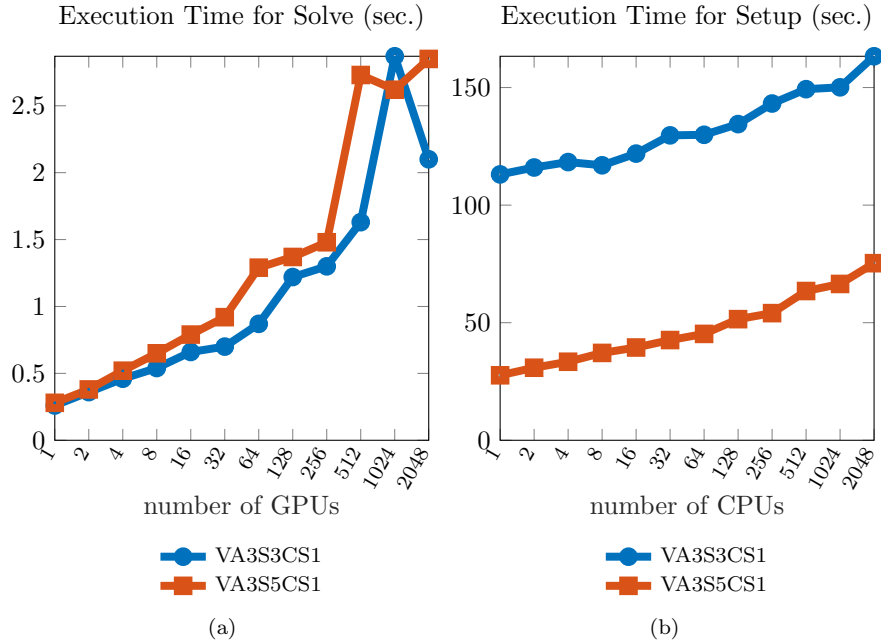
Fig. 11: Weak scaling results 6M dofs per GPU. Solve time on GPUs and Setup time on CPU.

versions of the HINVK setup phase to make better use of hybrid computing nodes. Let us observe that in time-dependent problems, such as in most of the test cases from the EoCoE-II project, large setup times are generally well-tolerated by the computational procedure if the solve phase is very efficient, since the same preconditioner is applied in a very large number of time steps. An extended comparison of the behavior of the various preconditioners in Figures 9, and 11, is available in the supplementary materials; in particular we have illustrated in Fig. 16 the solve time for CPU and GPU variants for weak scaling up to 10^{10} dofs.

6. Conclusions and future work. We presented a new version of a parallel preconditioners package implementing AMG methods, addressing scalability, flexibility and robustness for high-performance scientific computing at extreme scale. We focused on the design and implementation of a new parallel coarsening algorithm to be used in conjunction with highly parallel smoothers within AMG cycles for iterative solution of s.p.d. linear systems with size larger than billions on a current pre-exascale computer. Scalability results both on CPU and on GPU accelerators are very promising and a comparison with available software demonstrates the validity of our approaches both in terms of algorithms and in terms of software development.

Currently we are investigating alternative OpenMP and CUDA implementations of the kernels involved in the AMG setup, including graph matching, sparse matrix by sparse matrix product and approximate sparse matrix inversion, as well as other kernels employed in various phases of the matrix and communication setup. We plan to include these alternative OpenMP and/or CUDA versions in future releases of the package.

REFERENCES

[1] *PETSc users manual.* http://www.mcs.anl.gov/petsc. Technical Report ANL-95/11 - Revision 3.111, Argonne National Laboratory.

[2] P. Amestoy, I. Duff, and J.-Y. L'Excellent, *Multifrontal parallel distributed symmetric and unsymmetric solvers*, Comput. Methods in Appl. Mech. Eng., 184 (2000), pp. 501–520, https://doi.org/10.1016/S0045-7825(99)00242-X.

[3] H. Anzt, E. Boman, R. Falgout, P. Ghysels, M. Heroux, X. Li, L. Curfman McInnes, R. Tran Mills, S. Rajamanickam, K. Rupp, B. Smith, I. Yamazaki, and U. M. Yang, *Preparing sparse solvers for exascale computing*, Philos. Trans. R. Soc. A, 378 (2020), p. 20190053, https://doi.org/10.1098/rsta.2019.0053.

[4] A. Aprovitola, P. D'Ambra, F. Denaro, D. di Serafino, and S. Filippone, *Scalable algebraic multilevel preconditioners with application to CFD*, in Parallel computational fluid dynamics 2008, vol. 74 of Lect. Notes Comput. Sci. Eng., Springer, Heidelberg, 2010, pp. 15–27, https://doi.org/10.1007/978-3-642-14438-7_2.

[5] A. H. Baker, R. D. Falgout, T. V. Kolev, and U. M. Yang, *Multigrid smoothers for ultraparallel computing*, SIAM J. Sci. Comput., 33 (2011), pp. 2864–2887, https://doi.org/10.1137/100798806.

[6] A. H. Baker, R. D. Falgout, T. V. Kolev, and U. M. Yang, *Scaling Hypre's Multigrid Solvers to 100,000 Cores*, in High-Performance Scientific Computing: Algorithms and Applications, M. W. Berry, K. A. Gallivan, E. Gallopoulos, A. Grama, B. Philippe, Y. Saad, and F. Saied, eds., Springer London, London, 2012, pp. 261–279, https://doi.org/10.1007/978-1-4471-2437-5_13.

[7] M. Bernaschi, P. D'Ambra, and D. Pasquini, *AMG based on compatible weighted matching for GPUs*, Parallel Comput., 92 (2020), pp. 102599, 13, https://doi.org/10.1016/j.parco.2019.102599.

[8] M. Bernaschi, P. D'Ambra, and D. Pasquini, *BootCMatchG: An adaptive Algebraic Multi-Grid linear solver for GPUs*, Software Impacts, 6 (2020), https://doi.org/10.1016/j.simpa.2020.100041.

[9] D. Bertaccini and F. Durastante, *Iterative methods and preconditioning for large and sparse linear systems with applications*, Monographs and Research Notes in Mathematics, CRC Press, Boca Raton, FL, 2018.

[10] D. Bertaccini and S. Filippone, *Sparse approximate inverse preconditioners on high performance GPU platforms*, Comput. Math. Appl., 71 (2016), pp. 693–711, https://doi.org/10.1016/j.camwa.2015.12.008.

[11] A. Brandt, *General highly accurate algebraic coarsening*, Electronic Transactions on Numerical Analysis, 10 (2000), pp. 1–20.

[12] J. Brannick, Y. Chen, X. Hu, and L. Zikatanov, *Parallel unsmoothed aggregation algebraic multigrid algorithms on GPUs*, in Numerical solution of partial differential equations: theory, algorithms, and their applications, vol. 45 of Springer Proc. Math. Stat., Springer, New York, 2013, pp. 81–102, https://doi.org/10.1007/978-1-4614-7172-1_5.

[13] J. Brannick, Y. Chen, J. Kraus, and L. Zikatanov, *Algebraic multilevel preconditioners for the graph Laplacian based on matching in graphs*, SIAM J. Numer. Anal., 51 (2013), pp. 1805–1827, https://doi.org/10.1137/120876083.

[14] J. Brannick, Y. Chen, and L. Zikatanov, *An algebraic multilevel method for anisotropic elliptic equations based on subgraph matching*, Numer. Linear Algebra Appl., 19 (2012), pp. 279—-295, https://doi.org/doi.org/10.1002/nla.1804.

[15] A. Buttari, P. D'Ambra, D. di Serafino, and S. Filippone, *2LEV-D2P4: a package of high-performance preconditioners for scientific and engineering applications*, Appl. Algebra Engrg. Comm. Comput., 18 (2007), pp. 223–239, https://doi.org/0.1007/s00200-007-0035-z.

[16] V. Cardellini, S. Filippone, and D. Rouson, *Design patterns for sparse-matrix computations on hybrid CPU/GPU platforms*, Sci. Programming, 22 (2014), pp. 1–19, https://doi.org/10.3233/SPR-130363.

[17] U. V. Catalyürek, F. Dobrian, A. Gebremedhin, M. Halappanavar, and A. Pothen,

Distributed-memory parallel algorithms for matching and coloring, in PCO'11 New Trends in Parallel Computing and Optimization, IEEE International Symposium on Parallel and Distributed Processing Workshops, IEEE CS, 2011, https://doi.org/10.1109/IPDPS.2011.360.

[18] A. J. CLEARY, R. D. FALGOUT, V. E. HENSON, J. E. JONES, T. A. MANTEUFFEL, S. F. MCCORMICK, G. N. MIRANDA, AND J. W. RUGE, *Robustness and scalability of algebraic multigrid*, SIAM J. Sci. Comput., 21 (2000), pp. 1886–1908, https://doi.org/10.1137/S1064827598339402. Iterative methods for solving systems of algebraic equations (Copper Mountain, CO, 1998).

[19] P. D'AMBRA, D. DI SERAFINO, AND S. FILIPPONE, *MLD2P4: a package of parallel algebraic multilevel domain decomposition preconditioners in Fortran 95*, ACM Trans. Math. Software, 37 (2010), pp. Art. 30, 23, https://doi.org/10.1145/1824801.1824808.

[20] P. D'AMBRA, F. DURASTANTE, AND S. FILIPPONE, *On the quality of matching-based aggregates for algebraic coarsening of SPD matrices in AMG*, 2020, https://arxiv.org/abs/arXiv:2001.09969.

[21] P. D'AMBRA, S. FILIPPONE, AND P. S. VASSILEVSKI, *BootCMatch: a software package for bootstrap AMG based on graph weighted matching*, ACM Trans. Math. Software, 44 (2018), pp. Art. 39, 25, https://doi.org/10.1145/3190647.

[22] P. D'AMBRA AND P. S. VASSILEVSKI, *Adaptive AMG with coarsening based on compatible weighted matching*, Comput. Vis. Sci., 16 (2013), pp. 59–76, https://doi.org/10.1007/s00791-014-0224-9.

[23] H. DE STERCK, U. YANG, AND J. J. HEYS, *Reducing complexity in parallel algebraic multigrid preconditioners*, SIAM J. Matrix Anal. Appl., 27 (2006), pp. 1019–1039, https://doi.org/10.1137/040615729.

[24] R. D. FALGOUT, J. E. JONES, AND U. M. YANG, *The design and implementation of hypre, a library of parallel high performance preconditioners*, in Numerical solution of partial differential equations on parallel computers, vol. 51 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2006, pp. 267–294, https://doi.org/10.1007/3-540-31619-1_8.

[25] R. D. FALGOUT AND P. S. VASSILEVSKI, *On generalizing the algebraic multigrid framework*, SIAM J. Numer. Anal., 42 (2004), pp. 1669–1693, https://doi.org/10.1137/S0036142903429742.

[26] S. FILIPPONE AND A. BUTTARI, *Object-oriented techniques for sparse matrix computations in Fortran 2003*, ACM TOMS, 38 (2012), pp. 23:1–23:20.

[27] S. FILIPPONE, V. CARDELLINI, D. BARBIERI, AND A. FANFARILLO, *Sparse matrix-vector multiplication on GPGPUs*, ACM Trans. Math. Software, 43 (2017), pp. Art. 30, 49.

[28] S. FILIPPONE AND M. COLAJANNI, *PSBLAS: a library for parallel linear algebra computations on sparse matrices*, ACM TOMS, 26 (2000), pp. 527–550.

[29] V. E. HENSON AND U. YANG, *BoomerAMG: a parallel algebraic multigrid solver and preconditioner*, Appl. Numer. Math., 41 (2002), pp. 155–177, https://doi.org/10.1016/S0168-9274(01)00115-5. Developments and trends in iterative methods for large systems of equations—in memoriam Rüdiger Weiss (Lausanne, 2000).

[30] M. HEROUX AND ALT., *An overview of the trilinos project*, ACM Trans. Math. Software, 31 (2015), pp. 397–423, https://doi.org/10.1145/1089014.1089021.

[31] S. J. KOLLET AND R. M. MAXWELL, *Integrated surface–groundwater flow modeling: A free-surface overland flow boundary condition in a parallel groundwater flow model*, Advances in Water Resources, 29 (2006), pp. 945 – 958, https://doi.org/https://doi.org/10.1016/j.advwatres.2005.08.006.

[32] X. S. LI, *An overview of SuperLU: algorithms, implementation, and user interface*, ACM Trans. Math. Software, 31 (2005), pp. 302–325, https://doi.org/10.1145/1089014.1089017.

[33] Y. NOTAY, *An aggregation-based algebraic multigrid method*, Electron. Trans. Numer. Anal., 37 (2010), pp. 123–146.

[34] Y. NOTAY AND A. NAPOV, *A massively parallel solver for discrete Poisson-like problems*, J. Comput. Phys., 281 (2015), pp. 237–250, https://doi.org/10.1016/j.jcp.2014.10.043.

[35] Y. NOTAY AND P. S. VASSILEVSKI, *Recursive Krylov-based multigrid cycles*, Numer. Linear Algebra Appl., 15 (2008), pp. 473–487, https://doi.org/10.1002/nla.542.

[36] K. STÜBEN, *Algebraic multigrid (AMG): an introduction with applications*, in Multigrid, U. Trottenberg, C. Oosterlee, and A. Schüller, eds., Frontiers in Applied Mathematics, Academic Press, 2001, pp. Appendix A, 413–532.

[37] R. TUMINARO AND C. TONG, *Parallel smoothed aggregation multigrid: aggregation strategies on massively parallel machines*, in SC '00: Proceedings of the 2000 ACM/IEEE Conference on Supercomputing, IEEE CS, 2000, https://doi.org/10.1109/SC.2000.10008.

[38] A. C. N. VAN DUIN, *Scalable parallel preconditioning with the sparse approximate inverse of*

triangular matrices, SIAM J. Matrix Anal. Appl., 20 (1999), pp. 987–1006, https://doi.org/10.1137/S0895479897317788. Sparse and structured matrices and their applications (Coeur d'Alene, ID, 1996).

[39] P. VANĚK, J. MANDEL, AND M. BREZINA, *Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196, https://doi.org/10.1007/BF02238511. International GAMM-Workshop on Multi-level Methods (Meisdorf, 1994).

[40] P. VASSILEVSKI, *Multilevel block factorization preconditioners: matrix-based analysis and algorithms for solving finite element equations*, Springer, New York, USA, 2008.

[41] M. VÁZQUEZ, G. HOUZEAUX, S. KORIC, A. ARTIGUES, J. AGUADO-SIERRA, R. ARÍS, D. MIRA, H. CALMET, F. CUCCHIETTI, H. OWEN, A. TAHA, E. D. BURNESS, J. M. CELA, AND M. VALERO, *Alya: Multiphysics engineering simulation toward exascale*, Journal of Computational Science, 14 (2016), pp. 15 – 27, https://doi.org/https://doi.org/10.1016/j.jocs.2015.12.007. The Route to Exascale: Novel Mathematical Methods, Scalable Algorithms and Computational Science Skills.

[42] J. XU AND L. ZIKATANOV, *Algebraic multigrid methods*, Acta Numer., 26 (2017), pp. 591–721.

[43] U. YANG, *Parallel algebraic multigrid methods—high performance preconditioners*, in Numerical solution of partial differential equations on parallel computers, vol. 51 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2006, pp. 209–236, https://doi.org/10.1007/3-540-31619-1_6.

[44] U. YANG, *On long-range interpolation operators for aggressive coarsening*, Numer. Linear Algebra Appl., 17 (2010), pp. 453–472, https://doi.org/10.1002/nla.689.

SUPPLEMENTARY MATERIALS: AMG PRECONDITIONERS FOR LINEAR SOLVERS TOWARDS EXTREME SCALE

This supplementary document provides further information on the experiments described in the paper. The sections to which the information refer have the same name of the corresponding sections in the paper.

Appendix 1. Parallel Smoothers and Coarsest Solvers. To describe the effect of using different smoothers in the AMG algorithm, we have used the two-level convergence analysis based on the computation of the constant K to bound the spectral radius of the error propagation matrix of the method, i.e:

$$\|I - BA\|_A^2 = \|(I - P(P^T A P)^{-1} P^T A)(I - M^{-1} A)\|_A^2 \leq 1 - \frac{1}{K},$$

for which it holds

$$K = \sup_e \frac{\|(I - PR)e\|_{\overline{M}}^2}{\|e\|_A^2} \geq 1, \quad \overline{M} = M^T (M^T + M - A)^{-1} M,$$

for a fixed projector P, and R any matrix such that PR is a projection onto $\mathcal{R}ange(P)$. As we mentioned in section 3.3, to improve the convergence properties of a single V-cycle, one can consider a smoothed version of such a prolongator, i.e., a more accurate interpolation operator obtained by applying one step of a weighted-Jacobi smoother to the P obtained from the matching strategy. The actual prolongator $\overline{P}$ is obtained from P as $\overline{P} = (I - \omega D^{-1} A)P$, where $D = \mathrm{diag}(A)$ and $\omega = 1/\|D^{-1}A\|_\infty \approx 1/\rho(D^{-1}A)$, with $\rho(D^{-1}A)$ the spectral radius of $D^{-1}A$. In this case we no longer have an explicit expression for the P and R needed for computing the constant K inherited from the matching construction in section 3.2. Nevertheless, we can analyze the behavior on the same test case by numerically computing an R such that PR is a projection onto $\mathcal{R}ange(P)$. In Tab. 3 we report the results of these computations; if we compare them with those obtained for the unsmoothed aggregation in Tab. 1, we observe the same behavior with respect to the smoothers, but we achieve better constants in absolute terms. This confirms that increasing the regularity of the projection operator leads to better convergence properties in this framework.

Appendix 2. Comparison with `Hypre`. We compare our preconditioners with some preconditioners available in the `Hypre` library [24], using default algorithmic parameters for best practices. For our choice of coarsening algorithms we consider:

Falgout This is a combination of the classic Ruge-Stüben and the parallel independent set, named CLJP, coarsening algorithms. An initial RS0 coarsening requiring no communication is built; this generates a set of coarse variables (C-indices) for each processor, and the the variables associated with the internal points of each process are subsequently used as the first independent set in a CLJP coarsening sweep (see [29], and [43, section 3.2]).

HMIS The Hybrid Maximal Independent Set algorithm uses one pass of the Ruge-Stüben coarsening algorithm on each processor independently. The interior coarse variables (C-indices) are then used as first independent sets for a parallel coarsening algorithm [23].

HMIS1 This is a variant of the HMIS algorithm in which 1 level of *aggressive* coarsening is used, i.e., instead of the classical Ruge-Stüben coarsening algorithm a

Table 3: Convergence constant K for various smoothers for the homogeneous 3D Laplacian problem with parallel coarsening based on compatible weighted matching. The global size of the problem is $m = 2^4 \times 2^4 \times 2^4$ distributed over np processes using a 3D block distribution. By INVK we denote here the approximate inverse computed from the incomplete LU factorization with 0 levels of fill-in, and admitting a single level of fill-in in the inversion procedure.

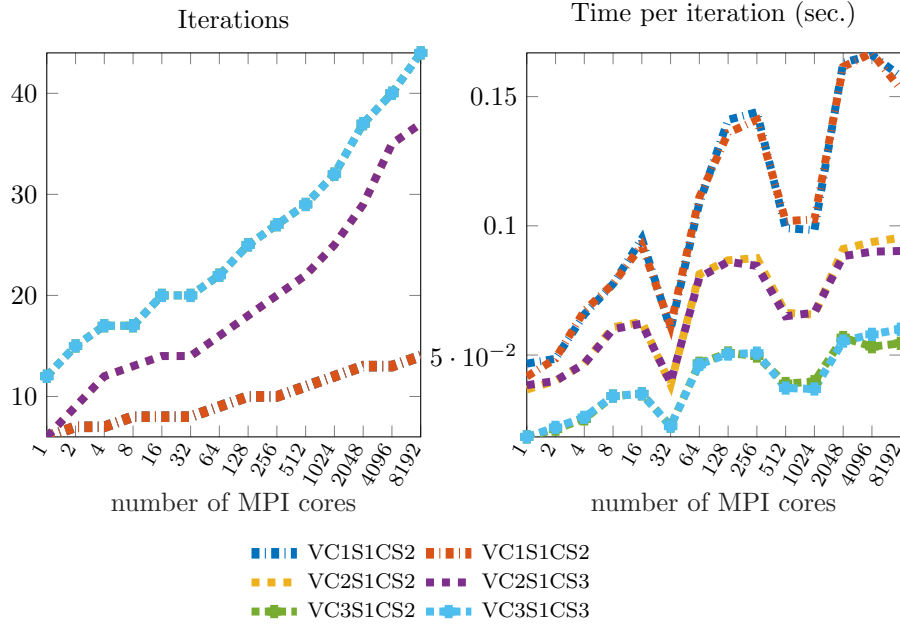| | | Smoothed prolongator 3 sweeps of matching | | | | | |
| m | np | HGS | L1-HGS | HINVK | L1-INVK | HINVK | L1-INVK |
				$\omega = 1$		ω_{opt}	
4096	1	1.3686	1.3686	1.5611	1.5611	1.0056	1.0056
2048	2	1.4108	1.5166	1.5293	1.5150	1.0390	1.0751
1024	4	1.4496	1.6500	1.5681	1.6318	1.2541	1.3553
512	8	1.4632	1.7443	1.7362	1.7450	1.4971	1.5554
256	16	1.4832	1.7938	1.7796	1.7966	1.5310	1.5400
128	32	1.5035	1.8435	1.8172	1.8413	1.5663	1.5371
64	64	1.5218	1.8948	1.8533	1.8836	1.6034	1.5523
32	128	1.5757	2.0123	1.9392	2.0133	1.7287	1.6242
16	256	1.6278	2.1369	2.0326	2.1491	1.8903	1.7451
8	512	1.6523	2.2833	2.1858	2.2921	2.1073	1.8999

stricter definition of connectivity is enforced resulting in sparser coarse grids; see [44, section 3] for the details.

We also include here the comparison between the `Hypre` preconditioners included in the paper and the same versions of the preconditioners, when the direct solution at the coarsest level is substituted by an iterative solution. `Hypre` makes available at the coarsest level the same (stationary) smoothers available at the intermediate levels; we selected the Symmetrized Gauss-Seidel method, which is the most robust among the available methods, and applied 10 iterations of this method to solve the coarsest system. All the other parameters are left as in the default. We can observe that there is no significant difference between the use of the iterative coarsest solver (see the preconditioner's names with the CS3 suffix) and the direct one (whose preconditioner's names have the CS2 suffix) both in terms of iterations and in terms of time for the setup and the solve phases.
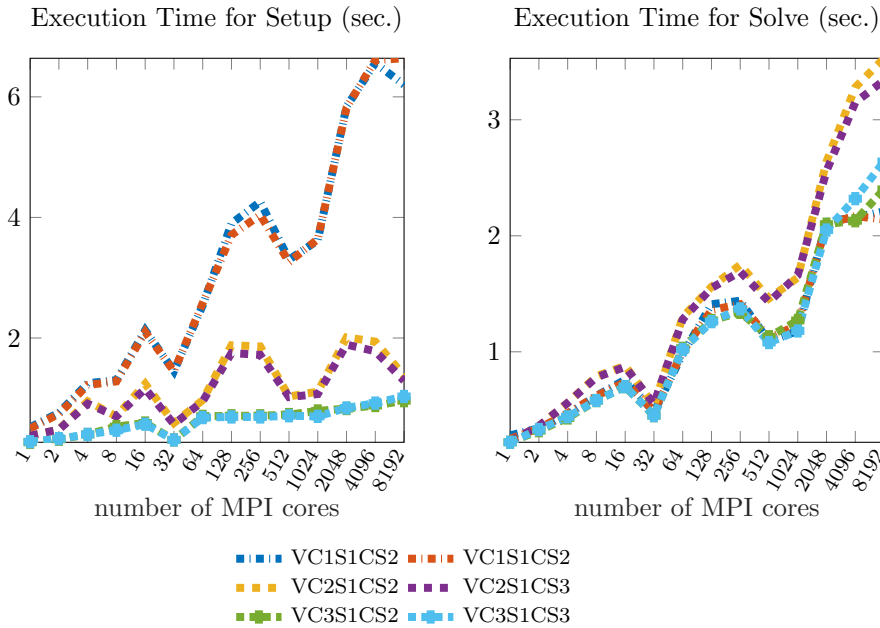
2.1. Anisotropic Poisson problem. To complement the experiment in Section 4, we consider here also the case of an anisotropic Poisson problem on a structured grid, i.e., we replace (4.1) with

$$\sum_{i=1}^{3} k_i \frac{\partial^2 u}{\partial x_i^2} = \mathbf{f}, \quad \text{in } [0, 1]^3.$$

Let us observe that in general the weighted matching algorithm is able of recovering the aggregates one would expect for the different type of anisotropies, consider, e.g., the case of z-aligned anisotropy, in which $k_3 = 100$ and $k_1 = k_2 = 1$ represented in Figure 13. We investigate then the behavior of the preconditioners on the test case for $k_1 = k_2 = 100$, and $k_3 = 1$ while using again $512k$ dofs per node. In Fig. 14a we show the operator complexity of the multilevel hierarchies corresponding to all the consid-

Iterations

Time per iteration (sec.)

number of MPI cores

number of MPI cores

VC1S1CS2 VC1S1CS2
VC2S1CS2 VC2S1CS3
VC3S1CS2 VC3S1CS3

(a) Number of iterations and Time per iteration of the `Hypre` preconditioners

Execution Time for Setup (sec.)

Execution Time for Solve (sec.)

number of MPI cores

number of MPI cores

VC1S1CS2 VC1S1CS2
VC2S1CS2 VC2S1CS3
VC3S1CS2 VC3S1CS3

(b) Setup Time and Solve Time of the `Hypre` Preconditioners

Fig. 12: Weak scaling results for 256k dofs per core. Comparison of the convergence behavior and execution time of the `Hypre` preconditioners, when direct or iterative coarsest solver is applied.
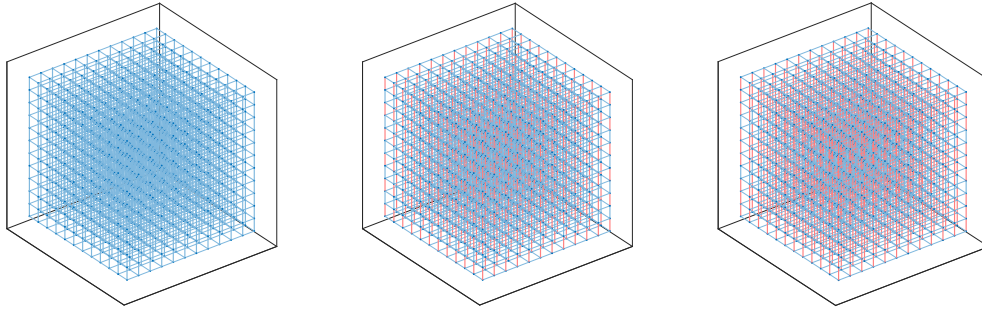
Fig. 13: Anisotropic Poisson problem with z-oriented axial anisotropy. The width of the edges is proportional to the weight of the corresponding edge. Edges that are part of the matching are highlighted in red, while unmatched edges are reported in cyan. The graph in the central panel corresponds to the case of a single sweep of matching, while the one in the last panel corresponds to the case of two matching sweeps.

ered preconditioners. What we observe is that, consistently with what happens with
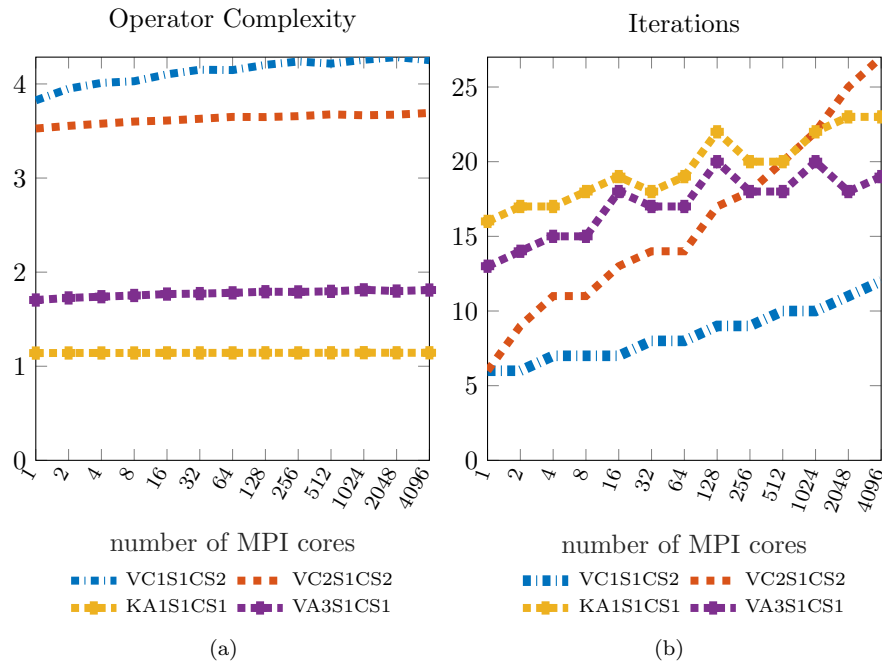

Fig. 14: Weak Scaling results for 512k dofs per core. Operator complexity and number of iterations.

the isotropic case, both the HMIS and Falgout coarsening schemes produce hierarchies with a significantly higher operator complexity. As seen in Fig. 14b, for Falgout this corresponds to the smaller number of overall iterations for the preconditioned system; however the number of iterations in the HMIS case tends to increase with the number

of cores. Let us also stress that we have stopped the analysis at 4096 cores for this case since both `Hypre` AMG schemes produce out of memory errors for the next step. On the other hand, both of our schemes have a smaller operator complexity and a better behavior with respect to the number of iterations if compared with the HMIS case in `Hypre`.   When we turn to the solve times in Fig. 15a, we observe that the KA1S1CS1 and VC2S1CS2 have comparable solve time; indeed what is happening in this case is that the interplay between the number of iterations and the cost of applying the K-Cycle in our KA1S1CS1 balances the higher operator complexity and the cost of using a V-Cycle in VC2S1CS2. On the other hand, as it already seen in the isotropic case, our VA3S1CS1 achieves a solve time that is comparable with the `Hypre` scheme based on the Falgout coarsening strategy VC1S1CS2, especially for increasing number of cores. Clearly, the smaller solve time of the latter for a reduced number of cores comes at the cost of a higher setup time, as it can be seen by comparing the two relative lines in Fig. 15b. The KA1S1CS1 and VC2S1CS2 have again comparable setup time, with the first one being marginally smaller. It is interesting to note that the speedups obtained in the solve phase are very similar for all methods, and behave consistently with the number of iterations and application cost; see Fig. 15c. On the other hand, in Fig. 15d, the speedup for the setup phase displays a smooth increase with increasing number of cores for all methods, with a slight better behavior for our proposed method, demonstrating again the good implementation scalability of all computational kernels.

Appendix 3. Performance results towards extreme scale. The resulting performance of the multigrid preconditioner in term of implementation scalability depends also on how effective the coarsening procedure is, and on how well balanced is the distribution of the coarsest matrix. We report in Tab. 4a the average coarsening ratio $\overline{cr}$ obtained for the Weak Scaling test with 512k dofs per node, that is, the average of the coarsening ratio cr_l at each level l defined as $\frac{1}{nl} \sum_{l=0}^{nl-1} cr_l = \frac{1}{nl} \sum_{l=0}^{nl-1} n_l/n_{l+1}$, with n_l the matrix size at the level l. Moreover, in Tab. 4b we report the size of the coarsest matrix together with the minimum and maximum block size. The size of the blocks and the minimum/maximum block sizes for the coarse matrix are the same for both the parallel matching with three (PMC3) and four sweeps (PMC4) aggregation strategies. The difference is that we reach this target dimension with a hierarchy with $nl = 4$ instead of $nl = 5$, and this is reflected in the average coarsening ratios in Tab. 4a. As we can observe, we achieve a good load-balancing among the different cores, i.e., there is a small variation between the size of the largest and smallest block size on each processor. Observe moreover that we can exactly compute an upper bound of the cost of the application of 1 K-cycle, so that, as also described in [35], the complexity for our K-cycle can be defined in terms of its operator complexity: K-opc $= \frac{\sum_{l=0}^{nl-1} 2^l \, \mathrm{nnz}(A_l)}{\mathrm{nnz}(A_0)}$. This takes into account the number of iterations ($it = 2$) at each level and the maximum number of times each level is visited in the recursion (2^{nl-l-1}). This complexity is then generally larger than that of the V-cycle although a still linearly increasing function in the non-zero entries of the system matrix (while the increase factor of the V-cycle is $1/(1-q)$, where q is the inverse of the coarsening ratio, the increase factor of the K-cycle is then $1/(1-2q)$.

Finally, we consider Fig. 16a and Fig. 16b in which we summarize the result that are discussed in detail in section 5.3. We can observe again how the smoothed version of the preconditioner using a V-cycle consistently achieves faster solution times with respect to the corresponding unsmoothed version. Moreover, when we compare these results with the GPU version of the VA3S5CS1 preconditioner we
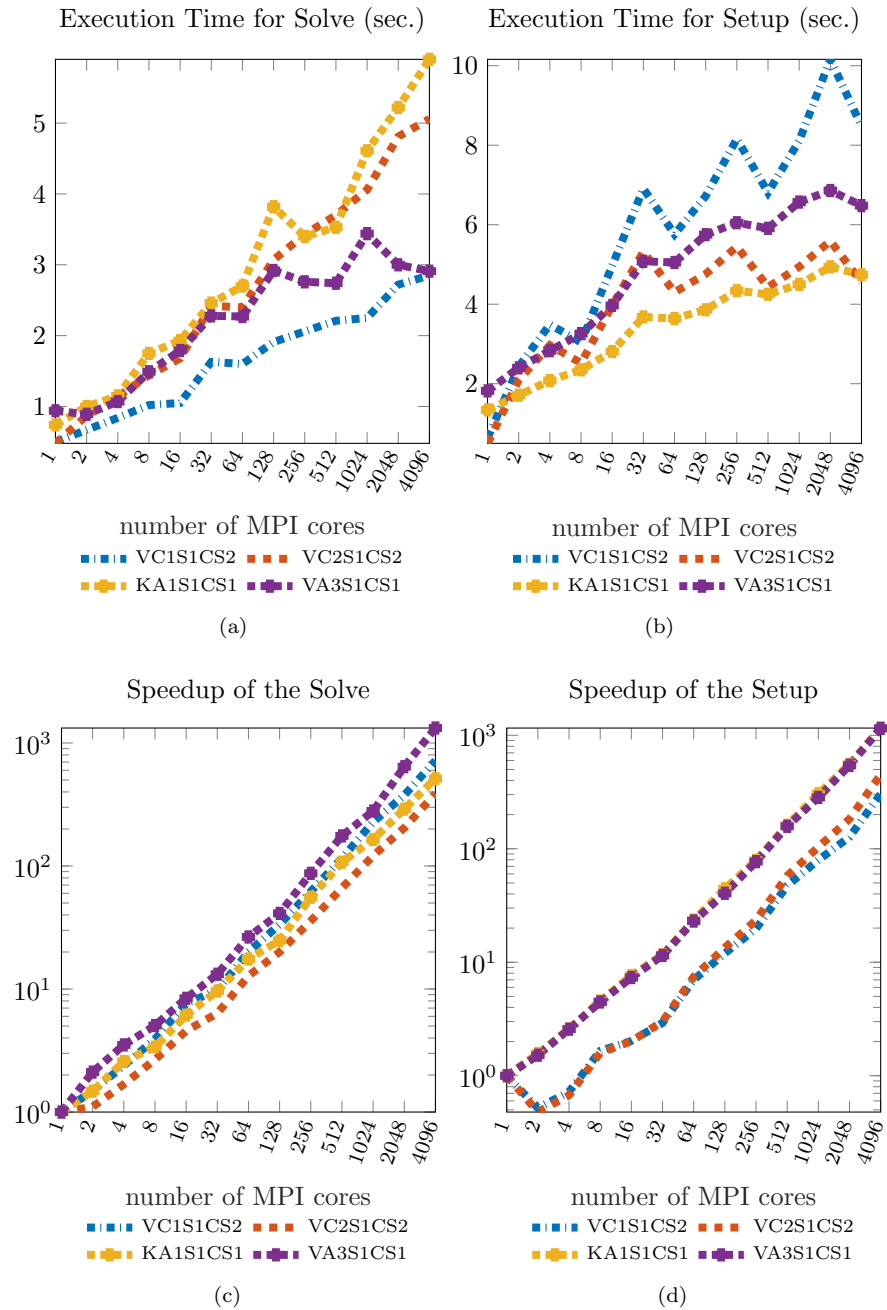
Fig. 15: Weak scaling results 512k dofs per core. Setup and solve time with the relative speed-up.

Execution Time for Solve (sec.)

(a) VA3S1CS1 preconditioner on MPI and VA3S5CS1 on MPI-CUDA

Execution Time for Solve (sec.)

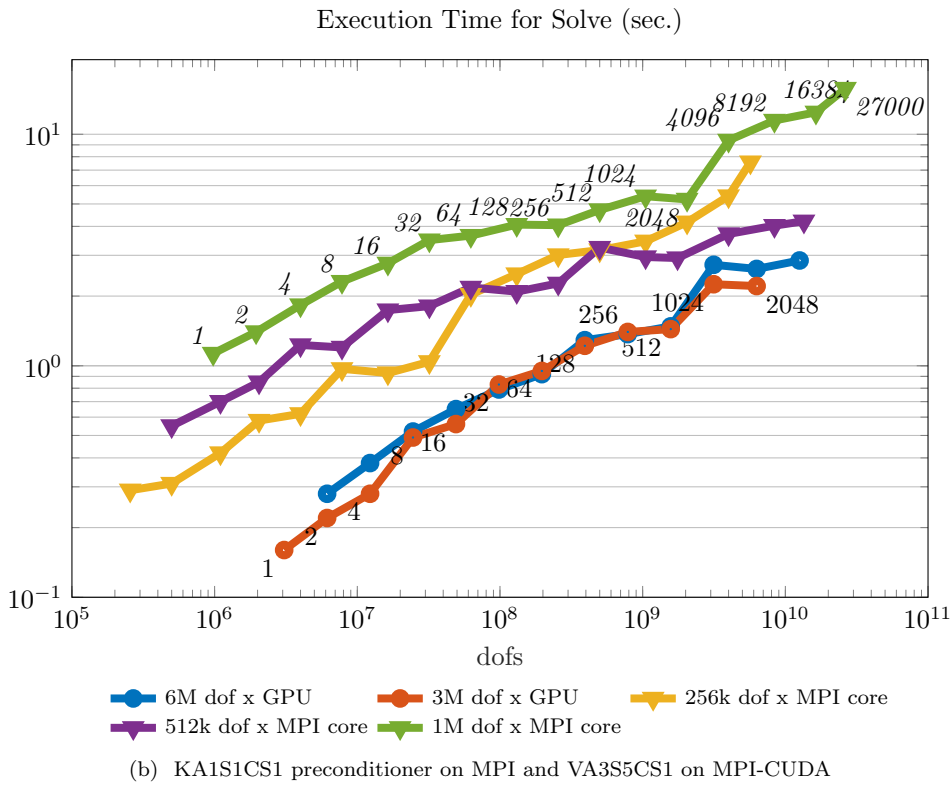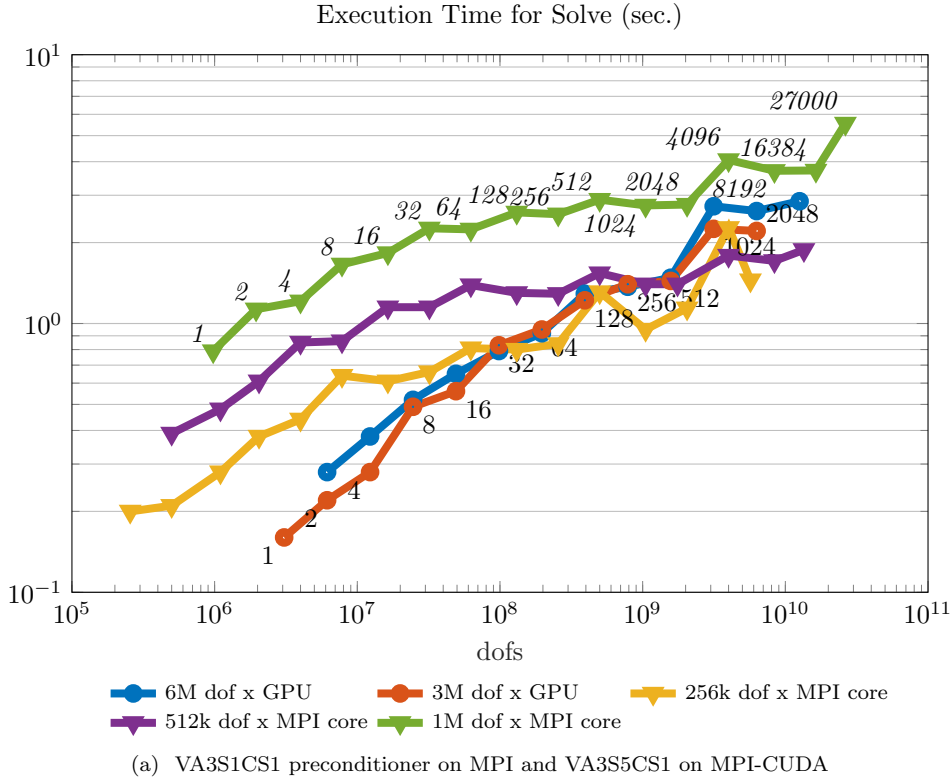(b) KA1S1CS1 preconditioner on MPI and VA3S5CS1 on MPI-CUDA

Fig. 16: Weak scaling results. Comparison of the execution time for solve between the pure MPI implementation of the VA3S1CS1 and KA1S1CS1 preconditioners, and the hybrid MPI-CUDA implementation of the VA3S5CS1 preconditioner.

Table 4: Average coarsening ratio $\overline{cr}$ for the weak scalability test with 512k dofs per node while using the Parallel Matching with 3 and 4 sweeps coarsening strategy, and information on the block-distribution of the coarsest matrix.

(a) average coarsening ratio						(b) Coarsest matrix block distribution			
Aggregation	PMC3		PMC4				block size		coarsest
np	nl	$\overline{cr}$	nl	$\overline{cr}$		np	min	max	size
1	5	8.00	4	16.00		1	125	125	125
2	5	7.98	4	15.95		2	138	139	277
4	5	8.00	4	16.00		4	128	128	512
8	5	8.00	4	16.00		8	125	125	1000
16	5	7.95	4	15.86		16	121	129	2004
32	5	8.00	4	16.00		32	128	128	4096
64	5	8.00	4	16.00		64	125	125	8000
128	5	7.93	4	15.81		128	119	139	16206
256	5	8.00	4	16.00		256	128	128	32768
512	5	8.00	4	16.00		512	125	125	64000
1024	5	7.96	4	15.90		1024	123	162	127466
2048	5	8.00	4	16.00		2048	128	128	262144
4096	5	7.94	4	15.84		4096	106	117	452335
8192	5	7.96	4	15.90		8192	115	133	1019452
16384	5	8.00	4	16.00		16384	128	128	2097152
27000	5	8.00	4	16.00		27000	125	125	3375000

observe that the hybrid approach permits savings in the solve time and, possibly even more importantly, large savings in the energy consumption, since to reach the same number of dofs and very similar execution time per linear iteration we need to use less CPU nodes.

Appendix 4. Software. The development versions of the `PSBLAS`, `PSBLAS-GPU` and `AMG4PSBLAS`, can be obtained from the website https://psctoolkit.github.io/. The new extension, containing the aggregation strategy based on graph matching will be made available by the end of the EoCoE-II project at the same address, together with the releases of the new versions of the libraries composing the whole software toolkit.