



Original software publication

## Parallel Sparse Computation Toolkit

Pasqua D'Ambra<sup>a</sup>, Fabio Durastante<sup>b,a,\*</sup>, Salvatore Filippone<sup>c,a</sup><sup>a</sup> Institute for Applied Computing "Mauro Picone", National Research Council, Naples, Italy<sup>b</sup> Department of Mathematics, University of Pisa, Pisa, Italy<sup>c</sup> Department of Civil and Computer Engineering, University of Rome "Tor Vergata", Rome, Italy

### ARTICLE INFO

#### Keywords:

Linear solvers  
Algebraic preconditioners  
HPC  
GPU  
Heterogeneous computing

### ABSTRACT

This paper presents a new software framework for solving large and sparse linear systems on current hybrid architectures, from small servers to high-end supercomputers, embedding multi-core CPUs and Nvidia GPUs at the node level. The framework has a modular structure and is composed of three main components, which separate basic functionalities for managing distributed sparse matrices and executing some sparse matrix computations involved in iterative Krylov projection methods, eventually exploiting multi-threading and CUDA-based programming models, from the functionalities for setup and application of different types of one-level and multi-level algebraic preconditioners.

### Code metadata

Current code version	v1.0 (PSBLAS v3.8.0-2, AMG4PSBLAS v1.1.0, PSBLAS3-EXT v1.3.0.1)
Permanent link to code/repository used for this code version	<a href="https://github.com/SoftwareImpacts/SIMPAC-2022-282">https://github.com/SoftwareImpacts/SIMPAC-2022-282</a>
Permanent link to Reproducible Capsule	<a href="https://codeocean.com/capsule/3044824/tree">https://codeocean.com/capsule/3044824/tree</a>
Legal Code License	BSD 3.0
Code versioning system used	git
Software code languages, tools, and services used	Fortran, C, C++, MPI, CUDA, OpenMP
Compilation requirements, operating environments & dependencies	GNU Toolchain (v. $\geq$ 6.0.0), OpenMPI or MPICH or Spectrum MPI with OpenBLAS or ATLAS or BLAS/LAPACK.
If available Link to developer documentation/manual	PSBLAS <a href="https://psctoolkit.github.io/psblasguide/psblas-3.8.pdf">https://psctoolkit.github.io/psblasguide/psblas-3.8.pdf</a> ; AMG4PSBLAS <a href="https://psctoolkit.github.io/amg4psblasguide/amg4psblas_1.0-guide.pdf">https://psctoolkit.github.io/amg4psblasguide/amg4psblas_1.0-guide.pdf</a> ; PSBLAS-EXT <a href="https://psctoolkit.github.io/psblasextguide/psblas-ext-1.3.pdf">https://psctoolkit.github.io/psblasextguide/psblas-ext-1.3.pdf</a>
Support email for questions	<a href="mailto:psctoolkit@na.iac.cnr.it">psctoolkit@na.iac.cnr.it</a>

## 1. The parallel sparse computation Toolkit

Sparse linear algebra is essential for a wide variety of scientific applications, from the solution of a discretized system of either linear or nonlinear partial differential equations to the analysis of large complex networks, or the computation of a few eigenvalues and eigenvectors of an operator. For all these problems having parallel sparse solvers is pivotal and lies at the core of the quasi-totality of multi-physics and multi-scale simulations. The technological challenge we face nowadays is investing effort and research in efficiently using hybrid architectures, embedding multi-core CPUs, and many-core accelerators, such as GPUs. These architectures are widely available in large infrastructures hosting

petascale and pre-exascale platforms as well as in small data centers. The suite of libraries we present here tries to face the above challenge by offering a framework in which it is one hand possible to use scalable algorithmic solutions to make extreme scale computing possible, on the other one hand to leverage a set of tools for efficient exploitation of multi-core CPUs and GPU accelerators in sparse matrix computations at the node level.

The Parallel Sparse Computation Toolkit (PSCToolkit) is made of three main building blocks

**PSBLAS** was developed with the aim to facilitate the parallelization of computationally intensive scientific applications and is designed

The code (and data) in this article has been certified as Reproducible by Code Ocean: (<https://codeocean.com/>). More information on the Reproducibility Badge Initiative is available at <https://www.elsevier.com/physical-sciences-and-engineering/computer-science/journals>.

\* Corresponding author at: Department of Mathematics, University of Pisa, Pisa, Italy.

E-mail addresses: [pasqua.dambra@cnr.it](mailto:pasqua.dambra@cnr.it) (P. D'Ambra), [fabio.durastante@unipi.it](mailto:fabio.durastante@unipi.it) (F. Durastante), [salvatore.filippone@uniroma2.it](mailto:salvatore.filippone@uniroma2.it) (S. Filippone).

<https://doi.org/10.1016/j.simpa.2022.100463>

Received 30 November 2022; Received in revised form 21 December 2022; Accepted 28 December 2022

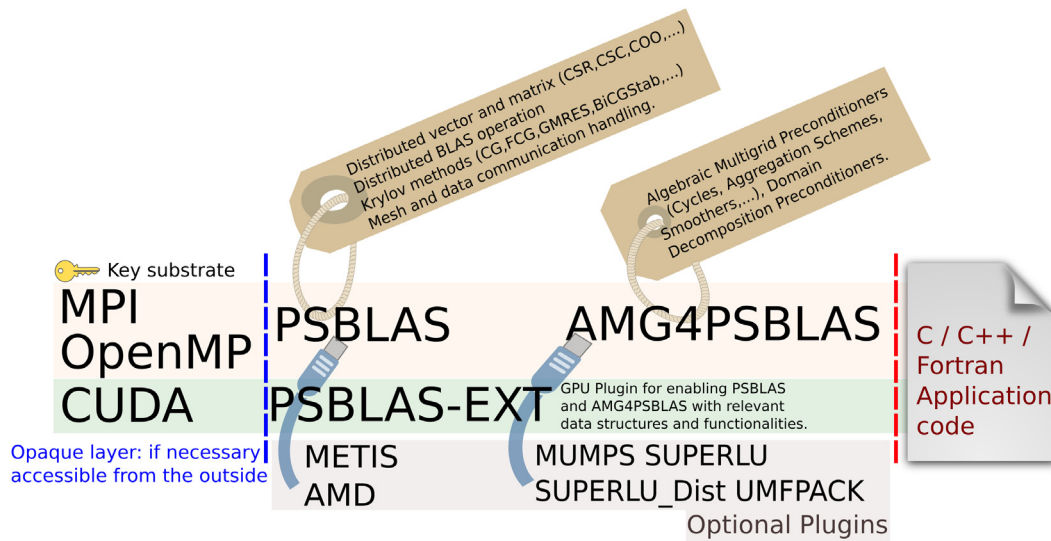


Fig. 1. A graphical depiction of the structure of the library. The base layer is represented by an implementation of MPI (we have run against several open and proprietary implementations: OpenMPI, MPICH, Spectrum MPI, Intel oneAPI), OpenMP, and a version of the CUDA Compilation Toolkit to run on NVIDIA GPUs. Optional functionalities can be added from third-party libraries such as the Metis graph partitioner [1], SuiteSparse <http://suitesparse.com> (AMD [2], UMFPACK [3]), MUMPS [4], SuperLU [5] and SuperLU\_dist [6]. The functionalities introduced by the toolkit (mesh and data handling, distributed data structure for vector and matrices, Krylov method, purely algebraic preconditioners) can be accessed by code written in both Fortran, C, and C++. In cases of a more elaborate integration with pre-existing code that needs to directly access the MPI/OpenMP/CUDA substrate, the toolkit is provided with routines to do so.

to address the parallel implementation of iterative solvers for sparse linear systems through the distributed-memory paradigm with message passing [7,8]. It includes routines for multiplying sparse matrices by dense matrices, solving block diagonal systems with triangular diagonal entries, and preprocessing sparse matrices. It contains also additional routines for dense matrix operations. In order to address the solution of extreme-scale problems, in the last versions of PSBLAS, effective handling of large index spaces requiring 8-byte integers has been implemented.

**AMG4PSBLAS** is a package of parallel algebraic preconditioners. It is a progress of a software development project started in 2007, named MLD2P4 [9–11], which originally implemented a multilevel version of some domain decomposition preconditioners of additive-Schwarz type, and was based on a parallel decoupled version of the well known smoothed aggregation method [12,13] to generate multilevel hierarchies of coarser matrices. The current version of the library extends the previous one by including many new smoothers and more general multigrid cycles and relies on a parallel implementation of the Coarsening based on Compatible Weighted Matching introduced in [14,15] and described in detail in [16].

**PSBLAS-EXT** contains a set of extensions to the PSBLAS library. The extensions provide additional storage matrix formats beyond the ones already contained in PSBLAS [8,17], as well as interfaces to two external libraries called SPGPU, for computations on NVIDIA GPUs and LIBRSB [18], for computations on multi-core parallel machines.

The PSCToolkit is implemented in the Fortran 2003 [19] programming language, with reuse and/or adaptation of existing Fortran 77 and Fortran 95 software, plus a handful of C and C++ routines. At the user level, the library can be called from both Fortran and C/C++. In Fig. 1 we give a graphical depiction of the whole structure.

## 2. Impact and use cases

The prototypical problem that PSCToolkit permits to solve is the solution of a large linear system of the form

$$Ax = b, \quad A \in \mathbb{K}^{n \times n}, \quad x, b \in \mathbb{K}^n, \quad \mathbb{K} = \{\mathbb{R}, \mathbb{C}\}, \quad (1)$$

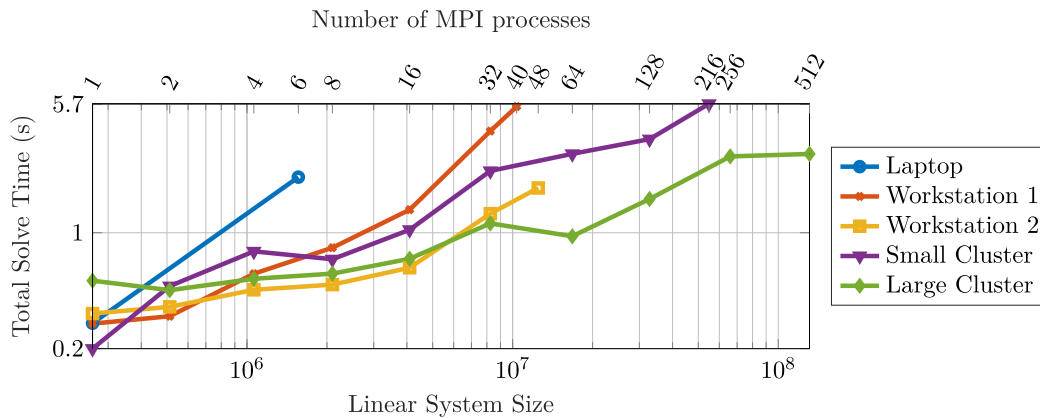
with  $A$  a *sparse matrix*, i.e., a matrix with a number of *nonzero entries* of the order of  $n$ , by using a Krylov method [20] – CG [21], FCG [22], GMRES [23], BiCGStab [24], BiCGStab( $\ell$ ) [25] – in conjunction with a purely algebraic preconditioner  $P$  transforming (1) into an equivalent one

$$P^{-1}Ax = P^{-1}b, \quad \text{or} \quad AP^{-1}(Px) = b, \quad \text{or} \quad P_1^{-1}AP_2^{-1}(P_2x) = P_1^{-1}b \quad \text{with} \quad P = P_1P_2,$$


with more favorable convergence properties. The toolkit permits to combine together classical stationary iterative methods – Jacobi, Gauss–Seidel, and their  $\ell_1$  and block-variants – together with incomplete matrix factorizations [26,27] – Incomplete LU factorizations [28], Approximate Inverses [29–31], sparse inverse LU factors [32] – and domain decomposition and algebraic multilevel strategies – Additive Schwarz, Restricted Additive Schwarz, Decoupled Smoothed Aggregation [12,13], Parallel Coupled smoothed and unsmoothed aggregation based on Compatible Weighted Matching [14,15] with  $V$ -,  $W$ - and  $K$ -cycles [33].

The various libraries offer a simplicity of extension that allows users to leverage the distributed environment management tools on the one hand to directly solve the linear systems that come from their application, on the other hand, to try new algorithmic solutions focusing mostly on mathematical aspects while delegating to the library many of the bearings related to the parallel environment. Furthermore, the solution of linear systems is not the only problem that can be faced with the tools available in the toolkit. More generally, problems that can be described in terms of BLAS for sparse arrays in a distributed environment can be implemented coherently with these tools, consider, e.g., the computation of eigenpairs, the computation of matrix functions, the solution of matrix equations using projection methods and various model order reduction techniques employing Krylov subspaces methods.

In this spirit, within the EoCoE (Energy-oriented Center of Excellence: towards Exascale for Energy) European project, the PSC-Toolkit has been recently interfaced and integrated into the Alya code [34], this is a simulation code for high-performance computational mechanics, which is part of the Unified European Applications Benchmark Suite (UEBAS) of scalable, currently relevant and publicly available application codes and datasets, of a size which can realistically be run on large systems, and maintained into the future. It



(a) Weak scaling. The size per computing element is set to  $n = 256000$  unknowns, and solve time is measured in seconds.



(b) Exploiting a full node, the number on top of the bars is the number of computing units. (c) Weak scaling. The size per computing element is set to  $n = 256000$  unknowns, and solve time is measured in seconds. We use from 1 to 16 nodes either with flat MPI (32 tasks per node) or using 4 MPI tasks per node with 4 GPUs.

Fig. 2. Numerical examples run with the machines described in Table 1.

enables the solution of coupled multiphysics problems using distributed and shared memory supercomputers and now permits the direct use of linear solvers and preconditioners from PSCToolkit. To test the use of the library on some example problems in an HPC environment it is available the Software as a Service (SaaS) portal provided by the Poznan Supercomputing and Networking Center (PSNC) at the URL [eocoe.psnc.pl](http://eocoe.psnc.pl). Furthermore, some interfacing modules for the KINSOL package available from the Sundials project [35] have been developed, that enable KINSOL to use the solvers and preconditioners from PSCToolkit inside its Newton-based non-linear procedure. These modules and the documentation which details their implementations are available from the GitHub repository. The interfacing modules extend the functionalities of PSCToolkit for dealing with algebraic non-linear systems and impacts on KINSOL so that it can leverage linear solvers made available by the toolkit.

### 2.1. An example of use

To demonstrate the applicability of the library and its possibility of use on different machines, we consider one of the test examples distributed with AMG4PSBLAS using different combinations of architectures and development environments; see Table 1. The test relies on the classic Poisson problem with homogeneous Dirichlet boundary

conditions:

$$\begin{cases} \nabla^2 u(\mathbf{x}) = f(\mathbf{x}), & \mathbf{x} \in \Omega = [0, 1]^3, \\ u(\mathbf{x}) = 0, & \mathbf{x} \in \partial\Omega, \end{cases}$$

discretized by the 7-points finite-differences method on the homogeneous tensor mesh. This discretization generates a symmetric positive definite and ill-conditioned linear system. For its solution, we choose to use the flexible conjugate gradient (FCG) method to achieve a residual on the error of  $10^{-6}$  preconditioned by one of the algebraic multigrid preconditioners available in AMG4PSBLAS. Specifically, we consider a single sweep of V-cycle leveraging an AMG hierarchy with a smoothed prolongator obtained by the Coarsening based on Compatible Weighted Matching from [16] with aggregates of size at most 8. We use 4 sweeps of  $\ell_1$ -Jacobi [36] both as pre- and post-smoother. On the coarsest grid, we use again a Krylov method preconditioned by a block-Jacobi method with incomplete LU factorization on the local blocks. In Fig. 2(a) we report the total solve time for a *weak-scaling* test using the MPI functionalities of the library, i.e., we show how the solution time varies with the number of computing units for a fixed problem size per unit — here fixed to approximately 256000 unknowns.

In the second test, we look again at the same problem and consider the same algorithmic configuration as above for the AMG preconditioner, except for the coarsest solver, where we use 30 iterations of the  $\ell_1$ -Jacobi method. In this case, we compare a sequential run for the whole reachable size in terms of unknowns, versus a pure MPI run

**Table 1**

List of machines on which the default tests have been executed. The table shows the type of machine, the hardware used and the software environment in which the tests were performed. The *Large Cluster* is the Marconi-100 system by CINECA, the 24th machine in the TOP500 November 2022 list: <https://www.top500.org/lists/top500/list/2022/11/>.

Machine	Hardware	Environment
Laptop	CPU: Intel <sup>(R)</sup> Core <sup>(TM)</sup> i7-8750H CPU @ 2.20 GHz; <b>Memory:</b> 16 Gb.	<b>Compiler:</b> GNU Suite 11.3.0; <b>MPI:</b> OpenMPI 4.1.2; <b>BLAS:</b> OpenBLAS 0.3.20.
Workstation 1	CPU: Intel <sup>(R)</sup> Xeon <sup>(R)</sup> Silver 4210 CPU @ 2.20 GHz; <b>Memory:</b> 64 Gb; <b>GPU:</b> NVIDIA Quadro RTX 5000.	<b>Compiler:</b> GNU Suite 11.3.0; <b>MPI:</b> mpich 3.4.3; <b>CUDA:</b> Cuda compilation tools, release 11.3, V11.3.109; <b>BLAS:</b> ATLAS 3.10.
Workstation 2	CPU: Intel <sup>(R)</sup> Xeon <sup>(R)</sup> Gold 6238R CPU @ 2.20 GHz; <b>Memory:</b> 1.48T.	<b>Compiler:</b> GNU Suite 6.1.0; <b>MPI:</b> OpenMPI 1.10.7; <b>BLAS:</b> OpenBLAS 0.3.3.
Small Cluster	CPU: 1 Node (cl1) with Intel <sup>(R)</sup> Xeon <sup>(R)</sup> CPU E5-2643 v4 @ 3.40 GHz, 4 nodes (cl2) with Intel <sup>(R)</sup> Xeon <sup>(R)</sup> CPU E5-2650 v4 @ 2.20 GHz; <b>Memory:</b> (cl1) 126 Gb, (cl2) 252 Gb; <b>Network:</b> Intel 10-Gigabit X540-AT2.	<b>Compiler:</b> GNU Suite 12.2.0; <b>MPI:</b> OpenMPI 4.1.4; <b>BLAS:</b> OpenBLAS 0.3.20.
Large Cluster	CPU 980 nodes with 2 × 16 cores IBM POWER9 AC922 at 2.6(3.1) GHz; <b>Memory:</b> 256 GB/node; <b>Network:</b> Mellanox IB EDR DragonFly++ 100 Gb/s; <b>GPU:</b> 4 × NVIDIA Volta V100 GPUs/node, Nvlink 2.0, 16 GB.	<b>Compiler:</b> GNU Suite 11.2.0; <b>MPI:</b> OpenMPI 4.1.2; <b>BLAS:</b> BLAS 3.10.0; <b>CUDA:</b> Cuda compilation tools, release 11.6, V11.6.124.

using the full CPU capabilities of the node, and a run using the available GPU accelerators on the given node; Fig. 2(b). On the *Workstation 2* (Table 1) we obtain for the GPU run a speedup of 21× with respect to the sequential solve time, and of 5× with respect to the solve time when the full CPU is used (40 MPI tasks). On the *Large Cluster*, which has 4 GPUs per node, we obtain a speedup of 68× with respect to the sequential solve time, and – consistently – of 5× with respect to the solve time when the full CPU is used (32 MPI tasks). Finally, in Fig. 2(c) we consider again a weak scaling on the *Large Cluster* fully using the computing nodes and around 8192000 unknowns per node. That is, we compare the scaling using either 32 MPI tasks per computing node or 4 MPI tasks and 4 GPUs per node while using up to 16 nodes. We observe behavior in line with the expected one given the performances on a single node and the performances using pure MPI. Small fluctuations are observed in the solution times which seem mostly attributable to the machine; in fact, the number of iterations for the use case with pure MPI is contained between 8 and 15, while in the hybrid case between 8 and 12.

For results in extreme scalability and comparison tests between the various preconditioners and some competing libraries, we refer to [16].

### 3. Ongoing projects

The PSCToolkit is being used and going under active development in the Euro-HPC project “TEXTAROSSA: Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale”, and in the National Center on HPC, Big Data and Quantum Computing of the Italian PNRR.

### 4. Limitations and further developments

Currently, support for OpenMP is limited to some internal functions. One of the development directions we are following is to extend it to heavier routines, such as the setup phase on the CPU of multi-grid preconditioners to then exploit the application on the GPU. The other two lines of development underway are on one hand the use of communication-avoiding Krylov methods to reduce the weight of the scalar products and (re)orthogonalization steps on the total cost of the method. On the other, we plan to investigate the usage of hierarchical and other data-sparse matrix structure for reducing the memory footprint of the various methods.

### 5. Scholarly publications enabled by PSCToolkit

PSCToolkit has been used in several contexts on a range of different architectures and software stacks. Applications are the construction of relaxation method for high-performance image segmentation on

GPUs [37]; the solution of the pressure equation in the Large-Eddy-Simulation of the Navier–Stokes equations [34,38–40]; the numerical solution of discretized elliptic differential equations [16,41]; the numerical solution of the linear systems arising from the application of a Quasi-Newton method to a discrete version of the Richards equation for the filtration of fluids in the vadose zone [42]. Some of the tools are discussed in the book [43] where parallel AMG methods for solving linear systems are illustrated.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgments

This work was partially supported by GNCS-INdAM, Italy, by the Euro-HPC project *Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale (TEXTAROSSA)*, Horizon 2020 Program for Research and Innovation, Project ID: 956831, and by the Horizon 2020 Project *Energy oriented Centre of Excellence: toward exascale for energy (EoCoE-II)*, Project ID: 824158.

### References

- [1] G. Karypis, V. Kumar, A fast and high quality multilevel scheme for partitioning irregular graphs, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 20 (1) (1998) 359–392, <http://dx.doi.org/10.1137/S1064827595287997>.
- [2] P.R. Amestoy, T.A. Davis, I.S. Duff, Algorithm 837: AMD, an approximate minimum degree ordering algorithm, *ACM Trans. Math. Softw.* (ISSN: 0098-3500) 30 (3) (2004) 381–388, <http://dx.doi.org/10.1145/1024074.1024081>.
- [3] T.A. Davis, Algorithm 832: UMFPACK V4.3—an unsymmetric-pattern multifrontal method, *ACM Trans. Math. Softw.* (ISSN: 0098-3500) 30 (2) (2004) 196–199, <http://dx.doi.org/10.1145/992200.992206>.
- [4] P. Amestoy, A. Buttari, I. Duff, A. Guermouche, J.-Y. L'Excellent, B. Uçar, Mumps, in: D. Padua (Ed.), *Encyclopedia of Parallel Computing*, Springer US, Boston, MA, ISBN: 978-0-387-09766-4, 2011, pp. 1232–1238, [http://dx.doi.org/10.1007/978-0-387-09766-4\\_204](http://dx.doi.org/10.1007/978-0-387-09766-4_204).
- [5] X.S. Li, An overview of SuperLU: Algorithms, implementation, and user interface, *ACM Trans. Math. Softw.* 31 (3) (2005) 302–325.
- [6] X.S. Li, J.W. Demmel, SuperLU-DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems, *ACM Trans. Math. Softw.* 29 (2) (2003) 110–140.
- [7] S. Filippone, M. Colajanni, PSBLAS: A library for parallel linear algebra computation on sparse matrices, *ACM Trans. Math. Software* (ISSN: 0098-3500) 26 (4) (2000) 527–550, <http://dx.doi.org/10.1145/365723.365732>.
- [8] S. Filippone, A. Buttari, Object-oriented techniques for sparse matrix computations in Fortran 2003, *ACM Trans. Math. Software* (ISSN: 0098-3500) 38 (4) (2012) <http://dx.doi.org/10.1145/2331130.2331131>.



- [9] A. Buttari, P. D'Ambra, D. di Serafino, S. Filippone, Extending PSBLAS to build parallel Schwarz preconditioners, in: J. Dongarra, K. Madsen, J. Waśniewski (Eds.), *Applied Parallel Computing. State of the Art in Scientific Computing*, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN: 978-3-540-33498-9, 2006, pp. 593–602.
- [10] A. Buttari, P. D'Ambra, D. di Serafino, S. Filippone, 2LEV-D2P4: a package of high-performance preconditioners for scientific and engineering applications, *Appl. Algebra Eng. Comm. Comput.* (ISSN: 0938-1279) 18 (3) (2007) 223–239, <http://dx.doi.org/10.1007/s00200-007-0035-z>.
- [11] P. D'Ambra, D. di Serafino, S. Filippone, MLD2P4: a package of parallel algebraic multilevel domain decomposition preconditioners in Fortran 95, *ACM Trans. Math. Softw.* (ISSN: 0098-3500) 37 (3) (2010) <http://dx.doi.org/10.1145/1824801.1824808>, Art. 30, 23.
- [12] M. Brezina, P. Vaněk, A black-box iterative solver based on a two-level Schwarz method, *Computing* (ISSN: 0010-485X) 63 (3) (1999) 233–263, <http://dx.doi.org/10.1007/s006070050033>.
- [13] P. Vaněk, J. Mandel, M. Brezina, Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems, in: *International GAMM-Workshop on Multi-Level Methods* (Meisdorf, 1994), vol. 56, (3) 1996, pp. 179–196, <http://dx.doi.org/10.1007/BF02238511>.
- [14] P. D'Ambra, P.S. Vassilevski, Adaptive AMG with coarsening based on compatible weighted matching, *Comput. Vis. Sci.* (ISSN: 1432-9360) 16 (2) (2013) 59–76, <http://dx.doi.org/10.1007/s00791-014-0224-9>.
- [15] P. D'Ambra, S. Filippone, P.S. Vassilevski, BootCMatch: a software package for bootstrap AMG based on graph weighted matching, *ACM Trans. Math. Softw.* (ISSN: 0098-3500) 44 (4) (2018) <http://dx.doi.org/10.1145/3190647>, Art. 39, 25.
- [16] P. D'Ambra, F. Durastante, S. Filippone, AMG preconditioners for linear solvers towards extreme scale, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 43 (5) (2021) S679–S703, <http://dx.doi.org/10.1137/20M134914X>.
- [17] V. Cardellini, S. Filippone, D.W. Rouson, Design patterns for sparse-matrix computations on hybrid CPU/GPU platforms, *Sci. Program.* 22 (2014) 1–19, <http://dx.doi.org/10.3233/SPR-130363>, 1.
- [18] M. Martone, S. Filippone, S. Tucci, M. Paprzycki, M. Ganzha, Utilizing recursive storage in sparse matrix-vector multiplication - preliminary considerations, in: T. Philips (Ed.), *Proceedings of the ISCA 25th International Conference on Computers and their Applications, CATA 2010, March 24–26, 2010, Sheraton Waikiki Hotel, Honolulu, Hawaii, USA, ISCA, 2010*, pp. 300–305.
- [19] M. Metcalf, J.K. Reid, M. Cohen, *Fortran 95/2003 Explained*, vol. 3, Oxford University Press Oxford, 2004.
- [20] Y. Saad, *Iterative Methods for Sparse Linear Systems*, second ed., Society for Industrial and Applied Mathematics, Philadelphia, PA, ISBN: 0-89871-534-2, 2003, p. xviii+528, <http://dx.doi.org/10.1137/1.9780898718003>.
- [21] M.R. Hestenes, E. Stiefel, *Methods of conjugate gradients for solving linear systems*, *J. Res. Natl. Bur. Stand.* (ISSN: 0160-1741) 49 (1952) 409–436 (1953).
- [22] Y. Notay, Flexible conjugate gradients, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 22 (4) (2000) 1444–1460, <http://dx.doi.org/10.1137/S1064827599362314>.
- [23] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* (ISSN: 0196-5204) 7 (3) (1986) 856–869, <http://dx.doi.org/10.1137/0907058>.
- [24] H.A. van der Vorst, Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* (ISSN: 0196-5204) 13 (2) (1992) 631–644, <http://dx.doi.org/10.1137/0913035>.
- [25] G.L.G. Sleijpen, H.A. van der Vorst, D.R. Fokkema, BiCGstab(l) and other hybrid Bi-CG methods, *Numer. Algorithms* (ISSN: 1017-1398) 7 (1) (1994) 75–109, <http://dx.doi.org/10.1007/BF02141261>.
- [26] D. Bertaccini, S. Filippone, Sparse approximate inverse preconditioners on high performance GPU platforms, *Comput. Math. Appl.* (ISSN: 0898-1221) 71 (3) (2016) 693–711, <http://dx.doi.org/10.1016/j.camwa.2015.12.008>.
- [27] D. Bertaccini, F. Durastante, Interpolating preconditioners for the solution of sequence of linear systems, *Comput. Math. Appl.* (ISSN: 0898-1221) 72 (4) (2016) 1118–1130, <http://dx.doi.org/10.1016/j.camwa.2016.06.023>.
- [28] Y. Saad, ILUT: a dual threshold incomplete LU factorization, *Numer. Linear Algebra Appl.* (ISSN: 1070-5325) 1 (4) (1994) 387–402, <http://dx.doi.org/10.1002/nla.1680010405>.
- [29] M. Benzi, J.K. Cullum, M. Tüma, Robust approximate inverse preconditioning for the conjugate gradient method, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 22 (4) (2000) 1318–1332, <http://dx.doi.org/10.1137/S1064827599356900>.
- [30] M. Benzi, M. Tüma, A sparse approximate inverse preconditioner for nonsymmetric linear systems, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 19 (3) (1998) 968–994, <http://dx.doi.org/10.1137/S1064827595294691>.
- [31] M. Benzi, C.D. Meyer, M. Tüma, A sparse approximate inverse preconditioner for the conjugate gradient method, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 17 (5) (1996) 1135–1149, <http://dx.doi.org/10.1137/S1064827594271421>.
- [32] A.C.N. van Duin, Scalable parallel preconditioning with the sparse approximate inverse of triangular matrices, vol. 20, (4) (ISSN: 0895-4798) 1999, pp. 987–1006, <http://dx.doi.org/10.1137/S0895479897317788>, Sparse and structured matrices and their applications (Coeur d'Alene, ID, 1996),
- [33] P. Vassilevski, *Multilevel Block Factorization Preconditioners: Matrix-Based Analysis and Algorithms for Solving Finite Element Equations*, Springer, New York, USA, 2008.
- [34] H. Owen, O. Lehmkuhl, P. D'Ambra, F. Durastante, S. Filippone, Alya towards exascale: Algorithmic scalability using psctoolkit, 2022, <http://dx.doi.org/10.48550/ARXIV.2210.16660>.
- [35] A.C. Hindmarsh, P.N. Brown, K.E. Grant, S.L. Lee, R. Serban, D.E. Shumaker, C.S. Woodward, SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers, *ACM Trans. Math. Software* (ISSN: 0098-3500) 31 (3) (2005) 363–396, <http://dx.doi.org/10.1145/1089014.1089020>.
- [36] A.H. Baker, R.D. Falgout, T.V. Kolev, U.M. Yang, Multigrid smoothers for ultraparallel computing, *SIAM J. Sci. Comput.* (ISSN: 1064-8275) 33 (5) (2011) 2864–2887, <http://dx.doi.org/10.1137/100798806>.
- [37] P. D'Ambra, S. Filippone, A parallel generalized relaxation method for high-performance image segmentation on GPUs, *J. Comput. Appl. Math.* (ISSN: 0377-0427) 293 (2016) 35–44, <http://dx.doi.org/10.1016/j.cam.2015.04.035>.
- [38] A. Arovitola, P. D'Ambra, F.M. Denaro, D. di Serafino, S. Filippone, SPARLES: enabling large eddy simulations with parallel sparse matrix computation tools, *Comput. Math. Appl.* (ISSN: 0898-1221) 70 (11) (2015) 2688–2700, <http://dx.doi.org/10.1016/j.camwa.2015.06.028>.
- [39] P. D'Ambra, D. di Serafino, S. Filippone, Performance analysis of parallel Schwarz preconditioners in the LES of turbulent channel flows, *Comput. Math. Appl.* (ISSN: 0898-1221) 65 (3) (2013) 352–361, <http://dx.doi.org/10.1016/j.camwa.2012.06.023>.
- [40] A. Arovitola, P. D'Ambra, F. Denaro, D. di Serafino, S. Filippone, Scalable algebraic multilevel preconditioners with application to CFD, in: *Parallel Computational Fluid Dynamics 2008*, in: *Lect. Notes Comput. Sci. Eng.*, vol. 74, Springer, Heidelberg, 2010, pp. 15–27, [http://dx.doi.org/10.1007/978-3-642-14438-7\\_2](http://dx.doi.org/10.1007/978-3-642-14438-7_2).
- [41] B. Alfio, V. De Simone, D. di Serafino, Parallel algebraic multilevel Schwarz preconditioners for a class of elliptic PDE systems, *Comput. Vis. Sci.* (ISSN: 1432-9360) 16 (1) (2013) 1–14, <http://dx.doi.org/10.1007/s00791-014-0220-0>.
- [42] D. Bertaccini, P. D'Ambra, F. Durastante, S. Filippone, Why diffusion-based preconditioning of richards equation works: spectral analysis and computational experiments at very large scale, 2021, <http://dx.doi.org/10.48550/ARXIV.2112.05051>, URL <https://arxiv.org/abs/2112.05051>.
- [43] D. Bertaccini, F. Durastante, *Iterative Methods and Preconditioning for Large and Sparse Linear Systems with Applications*, in: *Monographs and Research Notes in Mathematics*, CRC Press, Boca Raton, FL, ISBN: 978-1-4987-6416-2, 2018, p. xviii+353.