
Subdomain Deflation Combined with Local AMG: a Case Study Using AMGCL Library

Denis Demidov^{1,*} and Riccardo Rossi^{2,3}

(Submitted by A. V. Lapin)

¹ Kazan Branch of Joint Supercomputer Center, Russian Academy of Sciences, Lobachevskii st. 2/31, 420111 Kazan, Russian Federation

² Department of Civil and Environmental Engineering, Universitat Politècnica de Catalunya (UPC), Barcelona 08034, Spain

³ CIMNE – Centre Internacional de Mètodes Numèrics en Enginyeria, CIMNE, Barcelona 08034, Spain

Received October 10, 2019, revised December 15, 2019; accepted December 23, 2019

Abstract—The paper proposes a combination of the subdomain deflation method and local algebraic multigrid as a scalable distributed memory preconditioner that is able to solve large linear systems of equations. The implementation of the algorithm is made available for the community as part of an open source AMGCL library. The solution targets both homogeneous (CPU-only) and heterogeneous (CPU/GPU) systems, employing hybrid MPI/OpenMP approach in the former and a combination of MPI, OpenMP, and CUDA in the latter cases. The use of OpenMP minimizes the number of MPI processes, thus reducing the communication overhead of the deflation method and improving both weak and strong scalability of the preconditioner. The examples of scalar (single degree of freedom per grid node), Poisson-like, systems as well as non-scalar problems, stemming out of the discretization of the Navier-Stokes equations, are considered in order to estimate performance of the implemented algorithm. A comparison with a traditional global AMG preconditioner based on a well-established Trilinos ML package is provided.

2010 Mathematical Subject Classification: 35-04, 65-04, 65Y05, 65Y10, 65Y15, 97N80

Keywords and phrases: *Subdomain Deflation, Algebraic Multigrid, Scalability, MPI, OpenMP, CUDA*

1. INTRODUCTION

The need to solve large and sparse linear systems of equations is ubiquitous in engineering and physics. Sparse linear solvers thus represent one of the cornerstones of modern numerical methods, justifying the variety of alternative techniques available for the purpose. Direct solvers, based on some form of matrix factorization are simply not viable beyond a certain size (typically of the order of a few millions of unknowns [16, 17]), due to their intrinsic memory requirements and sheer computational cost. Iterative methods thus become the only feasible alternative in addressing such large scale problems. These techniques typically fall in the class of Krylov-subspace solvers [22] or of Algebraic Multigrid algorithms (AMG) [3, 21, 26]. Among the common problems faced by Krylov approaches is their lack of robustness (compared to direct alternatives) as well as their limited algorithmic scalability as the problem size grows. Indeed a common feature of Krylov techniques is that the number of iterations needed for convergence increases with the problem size, thus making the methods less attractive for the solution of extremely large problems. On the other hand, advantages of Krylov methods are their simplicity and computational efficiency. These methods also typically show very good strong scalability for a given problem size (that is, the solution time decreases as more processes are used).

Conversely, multigrid methods show excellent weak scalability, exhibiting a convergence largely independent on the problem scale [5]. In turn, their limitation is related to their higher complexity and to the need for an expensive hierarchy, which may reduce the overall scalability.

* E-mail: dennis.demidov@gmail.com

In practice it is customary to blend the capabilities of multigrid and Krylov techniques so as to preserve some of the advantages of each [29]. A number of successful libraries exist implementing different flavors of such blending for distributed memory machines [14, 32]. Even though such implementations were proven to be weakly scalable, their strong scalability remains questionable. This fact motivated the renowned interest in Domain Decomposition (DD) methods [25] as well as the rise of a new class of approaches named *deflation techniques* [31]. Deflation techniques are closely related to Balancing Domain Decomposition with Constraints (BDDC) methods. However, it has been shown that deflation approaches are slightly advantageous [20].

Essentially all of the recent efforts in this area [12] take on the idea of constructing a multi level decomposition of the solution space, understood as a very aggressive coarsening of the original problem. The use of such coarser basis is proved sufficient to guarantee good weak scaling properties, while implying a minimal additional computational complexity and thus good strong scaling properties. An additional advantage of such approaches is the ease of combining them *in a modular way* with serial, SMP, or GPGPU accelerated preconditioners, which allows to make use of the vast number of existing local preconditioners in order to provide an effective MPI-based solver.

The current paper focuses specifically on the combination of a deflation approach with a local OpenMP-parallel AMG preconditioner. We explore the use of a constant and linearly-varying subspace defined on each of the MPI subdomains, verifying their weak and strong scalability properties.

Although the methods employed are presented in some detail, the focus of the contribution is on the software implementation within the open-source C++ library AMGCL¹. AMGCL provides several parallel backends, such as OpenMP or CUDA, which allows us to consider advantages of hybrid MPI/OpenMP and MPI/CUDA parallelism.

The paper begins by briefly describing deflation-based techniques. Then we introduce the cases of constant and linear deflation. Next, we detail the characteristics of the local AMG preconditioner used on the subdomains. Finally, in the benchmarks section, we prove the scalability of the developed solver by performing both strong and weak scalability tests on both CPU-only and GPU-accelerated systems. The well known Trilinos ML package [14] is used in order to provide a baseline comparison with the traditional approach based on a global AMG preconditioner. The package was chosen because it is currently one of the fastest publicly available implementations of AMG. For example, MueLU, which is a newer implementation of AMG available in Trilinos is still slower than ML by a non-negligible factor. See e.g. Figure 6 in [18]. Even though the publication is of the year 2014, and MueLU surely progressed since then, to the best of our knowledge (and after speaking to MueLU authors at conferences) ML remains a reference in terms of speed.

2. SUBDOMAIN DEFLATION

In order to provide a short description of the deflation approach, let us consider a linear system of equations in the following form:

$$\mathbf{A}\mathbf{x} = \mathbf{b}. \quad (1)$$

Here the matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is assumed to be invertible.

The idea of deflation is to decompose the unknown \mathbf{x} into a fine scale solution \mathbf{y} and a coarse scale solution $\bar{\mathbf{x}} = \mathbf{Z}\lambda$, so that

$$\mathbf{x} = \mathbf{y} + \mathbf{Z}\lambda. \quad (2)$$

Here $\mathbf{Z} \in \mathbb{R}^{n \times m}$ is a rectangular matrix whose columns contain the basis vectors of a subspace of size m .

By premultiplying (1) with \mathbf{Z}^T and using (2) we may write

$$\mathbf{Z}^T \mathbf{A} (\mathbf{y} + \mathbf{Z}\lambda) = \mathbf{Z}^T \mathbf{b}. \quad (3)$$

¹ <https://github.com/ddemidov/amgcl>

Substituting (2) into (1), and using (3) the original problem may be expressed in matrix form as

$$\begin{pmatrix} \mathbf{A} & \mathbf{AZ} \\ \mathbf{Z}^T \mathbf{A} & \mathbf{Z}^T \mathbf{AZ} \end{pmatrix} \begin{pmatrix} \mathbf{y} \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathbf{b} \\ \mathbf{Z}^T \mathbf{b} \end{pmatrix} \quad (4)$$

At this point it is customary to define the deflation matrix $\mathbf{E} = \mathbf{Z}^T \mathbf{AZ}$. The matrix is invertible if \mathbf{A} is invertible and the columns of \mathbf{Z} are linearly independent [12]. Under this assumption we can eliminate λ from the equations and obtain a modified system for \mathbf{y} equivalent to the original one. Namely, this corresponds to expressing λ as

$$\lambda = -\mathbf{E}^{-1} \mathbf{AZy} + \mathbf{E}^{-1} \mathbf{Z}^T \mathbf{b}. \quad (5)$$

Substituting (5) into the first row of (4) yields the following (singular) system:

$$(\mathbf{I} - \mathbf{AZE}^{-1} \mathbf{Z}^T) \mathbf{Ay} = (\mathbf{I} - \mathbf{AZE}^{-1} \mathbf{Z}^T) \mathbf{b}, \quad (6)$$

which may be understood as an operator P applied to both sides of the following equality:

$$P(\mathbf{Ay}) = P(\mathbf{b}). \quad (7)$$

Using the definition of \mathbf{E} , it may be easily shown that for any vector \mathbf{v}

$$P(P(\mathbf{v})) = P(\mathbf{v}), \quad (8)$$

thus proving that the operator P is mathematically a projection. The overall point is that instead of solving the original problem for the complete solution, we can now solve (7) for \mathbf{y} and recover the original solution from (2) and (5).

Depending on the choice of the basis in \mathbf{Z} the system in (7) may become easier to solve than the original. In order to get some intuition for why this may be true, think of constructing \mathbf{Z} by using the subspace spanned by the eigenvectors of \mathbf{A} with lowest associated eigenvalue. Such technique is called “eigenvalue deflation” and its properties are well known in the mathematical community. Physically such eigenvectors describe global features of the solution. In this case λ would describe the solution of the original problem when projected onto such space, making the vector $\bar{\mathbf{x}}$ a global approximation to the solution. The correction \mathbf{y} would then be associated only to the remaining eigenvectors which are relatively “rapidly varying”. Hence, the effect of deflation would be to remove the modes with low eigenvectors from the system, thus improving the condition number and in turn improving the convergence of Krylov solvers.

We can make a number of remarks at this point:

- The solution \mathbf{y} resides in a space orthogonal to the space described by \mathbf{Z} . Since the size of the problem $P(\mathbf{A})$ is the same as the original problem (1), this implies that the problem (7) is singular. Krylov methods allow a solution of such system since the right-hand side \mathbf{b} is also subjected to the same projection.
- If columns of \mathbf{Z} are eigenvectors of \mathbf{A} , a choice known as “eigenvalue deflation”, then it is easy to show that such eigenvectors and the respective eigenvalues are removed from the projected system. In other words, $P(\mathbf{A})$ is “deflated” of such eigenvectors. In particular, if the eigenvectors corresponding to the smallest eigenvalues were chosen for construction of \mathbf{Z} , then the smallest eigenvalues in the system would be removed, implying that the spectral radius of the projected system improves.
- If a slowly varying space is used in constructing \mathbf{Z} then the same argument holds *heuristically*, meaning that the spectral radius is improved as long as the deflation space used for construction of \mathbf{Z} allows approximating the eigenvectors with lower associated eigenvalues.
- The implementation of the projector P depends on the inverse of matrix \mathbf{E} . If this inverse is only approximated then the application of the projector *can only be understood as a left preconditioner applied to the original system*.

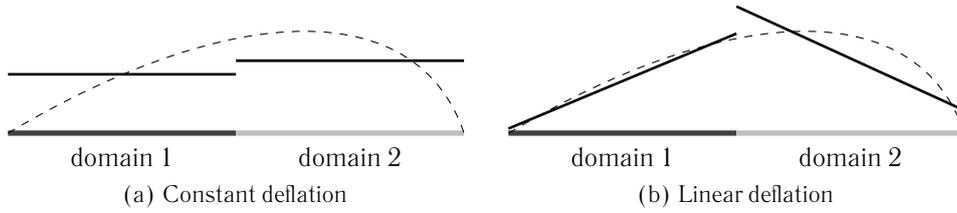


Figure 1. True (dashed line) and deflated (solid line) solution in case of constant and linear deflation.

- Even after the application of the projector, the system in (7) is still large and sparse. Deflation does not change the subspace in which \mathbf{y} resides, hence a preconditioner is still needed to allow the effective solution in terms of \mathbf{y} .

In practice the eigenvalues of \mathbf{A} are generally not known. Nevertheless it is possible to construct a global subspace by grouping the unknowns which are “close” in the graph of \mathbf{A} , in a way similar to what is done in algebraic multigrid. Such choice is valid when the solution is geometrically smooth, which we assume to be the case for all of the benchmark examples considered in the current work. When such smoothness cannot be assumed, as happens for example in the case of multiple materials or bubbly flows, the deflation idea can still be exploited effectively, but some problem-specific knowledge is needed to group the unknowns based on common physical characteristics [28]. This is possible through the AMGCL user interface, however we do not target this case in our benchmarks. In practice, for smooth solutions one may construct a basis by assigning a value of 1 to one of these groups at a time while zeroing out other groups, so that each deflation vector acts as an indicator for its respective subdomain. This approach is known as “subdomain deflation” [12]. Figure 1(a) provides a visual demonstration of the approach: the ideal solution is approximated by a constant solution over each MPI domain (two in the case shown) which represents the effect of the term $\mathbf{Z}\lambda$. Clearly the piecewise constant solution only represents a rough approximation of the real solution. The term \mathbf{y} is thus interpreted as the correction needed to complete the deflated solution and to arrive to the real one. The whole point here is that the correction \mathbf{y} is smaller than the total solution \mathbf{x} , it is thus intuitive to understand how finding \mathbf{y} requires a lesser effort (read, less iterations) than finding the whole solution. Figure 1(b) shows how the idea can be further improved by allowing the deflated solution to vary linearly within the domain. This choice provides better approximation for the real solution, hence reducing the residual and requiring less iterations to convergence. The drawback of this approach is the need to store the additional deflation vectors and the increased size of the coarse system, the latter being more important with the increasing number of MPI processes.

In the context of distributed linear algebra, it is customary that matrix \mathbf{A} is already mapped to MPI domains so as to preserve the locality of entries in the graph. Basing on this observation one can assign one deflation subdomain to each MPI process.

2.1. On precision of the coarse system solution

The approach just shown consists in identifying a modified “deflated” system, written in terms of an auxiliary variable \mathbf{y} which is easier to solve by iterative techniques. The core idea of the approach is that the real solution \mathbf{x} is easily recovered from \mathbf{y} by performing a simple correction step. As observed in the remarks, such approach relies heavily on the correct definition of the deflated system, that is, on the exact application of the deflation projector. In practice this implies that if we were to compute an inexact approximation $\hat{\mathbf{E}} \approx \mathbf{E}^{-1}$, then we would be solving the following system:

$$\left(\mathbf{I} - \mathbf{AZ}\hat{\mathbf{E}}\mathbf{Z}^T\right) \mathbf{A}\hat{\mathbf{y}} = \left(\mathbf{I} - \mathbf{AZ}\hat{\mathbf{E}}\mathbf{Z}^T\right) \mathbf{b} \quad (9)$$

thus yielding a solution $\hat{\mathbf{y}}$ which only approximately complies with the requirements. The immediate impact of this is that even after the final correction term, the solution will depend on the local approximation $\hat{\mathbf{E}}$ in a way that is difficult to quantify. This is not acceptable in a black box solver, since

it does not guarantee the final level of precision. Hence, the application of the inverse of \mathbf{E} must be computed very accurately if the projector is intended to be precise. Such requirement practically rules out the choice of Krylov or multigrid methods in the solution of the coarse problem, since the cost of enforcing a very high precision typically exceeds that of a direct approach for the problem sizes that are typical in the coarse space.

Even though the current paper focuses mostly on the application of the exact “projector” the possibility exists to consider the deflation as “inexact” [27] thus allowing the use of a more relaxed tolerance in the solution of the coarse problem. This implies considering the projector as a left preconditioner for the original linear problem, to be combined multiplicatively with other preconditioners. Such preconditioner should be considered as “variable” during the iterations (due to the inexact inverse of \mathbf{E}) so its application might require the use of a flexible Krylov solver, for example the flexible GMRES variant [22].

2.2. Domain decomposition approaches

The practical importance of deflation is in the context of domain decomposition (DD) methods. The idea is that, in the context of distributed linear algebra, the problem to be solved is divided into subdomains normally obtained by grouping the unknowns so that they represent physically contiguous portions of the domain to be solved. Under such assumption, an appealing possibility is to construct the deflation space in \mathbf{Z} by assuming that each of the basis vectors is purely local to one of the subdomains [12]. At this point one may construct a preconditioner following a classic domain decomposition approach, that is, discarding the columns of \mathbf{A} corresponding to nonlocal entries. If the deflation space is properly defined (for example taking it as constant over each subdomain) then the solution \mathbf{y} is discontinuous across the subdomain borders, but such discontinuity may easily be fixed with local corrections needed to retrofit the coarse solution to the real one. Such corrections are of local nature, and tend to show “high frequency” behavior. As such they are quickly addressed by local preconditioners. It should be intuitive how an increase in the number of subdomains improves the quality of the deflation space and thus reduces the correction needed in computing \mathbf{y} . This reflects in a reduction in the number of iterations needed to have the overall problem to converge, thus combining optimally with local preconditioners.

2.3. Constant deflation

The easiest option in constructing the deflation is named “constant” or “subdomain deflation” [12]. In this case the solution on each of the subdomains is approximated with a constant. If the size of the system is n and the number of subdomains considered is m , then $\mathbf{Z} \in \mathbb{R}^{n \times m}$ can be constructed as

$$\mathbf{z}_{ij} = \begin{cases} 1, & i \in S_j; \\ 0, & \text{otherwise;} \end{cases} \quad (10)$$

where S_j is the set of unknowns belonging to the j -th subdomain.

Note that \mathbf{Z} defined this way is completely local to each of the subdomains. This means that in practice it is more convenient to define and store just the local subblock \mathbf{Z}^l of the deflation space. In case of the constant deflation \mathbf{Z}^l consist of a single column filled with ones.

2.4. Linear Deflation

In some cases it is possible to associate physical coordinates with each unknown i in the system. In a three dimensional space we could define the coordinates as $(\phi_x^i, \phi_y^i, \phi_z^i)$. Then we could approximate the solution on each of the subdomains with a plane and interpret λ as the plane equation coefficients. This means we can define the local subblock \mathbf{Z}^l as a four-column matrix with

$$\mathbf{z}_{ij}^l = \begin{cases} 1, & j = 1; \\ \phi_x^i - c_x, & j = 2; \\ \phi_y^i - c_y, & j = 3; \\ \phi_z^i - c_z, & j = 4. \end{cases} \quad (11)$$

Here (c_x, c_y, c_z) are coordinates of the barycenter of the current subdomain.

3. IMPLEMENTATION

The deflation approach we discussed in the previous sections was implemented with AMGCL — an open source C++ library for solution of large sparse linear systems with algebraic multigrid (AMG) method [8, 29]. The library is developed in the Supercomputer center of the Russian academy of sciences and is published under liberal MIT license.

AMGCL has minimal set of dependencies, and is designed in a generic and extensible way that allows one to easily construct an AMG variant from available algorithmic components of the method, such as coarsening strategy or relaxation technique. The created hierarchy is transferred to one of the provided backends for the solution phase. Available backends support various parallelization technologies such as OpenMP (for use with multicore CPUs), CUDA or OpenCL (for use with modern GPUs). To illustrate this point, figure 2 shows OpenMP scalability of AMGCL for solution of a 3D Poisson problem on a single cluster node. The problem is solved on a uniform 150^3 mesh, the system matrix contains 3 375 000 unknowns and 23 490 000 nonzero values. The OpenMP-based AMGCL solution is compared with the MPI-based Trilinos ML solution, where all MPI processes are allocated within the same compute node, and with CUDA-based AMGCL solution, using NVIDIA Tesla K80 GPU available on the same node. The node is a dual socket system with two Intel Xeon E5-2640 v3 CPUs. The OpenMP-based AMGCL solution shows performance comparable with Trilinos ML. The AMGCL version uses one sweep of SPAI-0 [4] relaxation as smoother, and the Trilinos version uses two sweeps of Gauss-Seidel relaxation. This explains the fact that AMGCL needs about twice more iterations than Trilinos ML to converge, but the overall solution time is comparable for both versions. The CUDA-based solution is able to outperform both OpenMP-based AMGCL and Trilinos ML by about 2.5 times during solution phase, and is overall faster by about 30%.

The users of the library can easily implement their own backends in order for the library to work transparently with their custom data structures and algorithms. The main focus of AMGCL, as opposed to many available non-commercial packages, is to provide a non-MPI, “single-node” implementation of the AMG. The proposed approach allows to extend the library functionality onto MPI clusters in a modular way.

Listing 1 shows a basic example of using the library. Line 7 selects the backend to use. Here we use the builtin backend with `double` as a value type. Lines 10 to 17 define the solver type that we will use. The `make_solver` class binds together two concepts: a preconditioner (in this case an `amg` class), and an iterative solver (`cg`). Lines 20 to 22 show how to set solver parameters. Here we set the desired relative residual norm and a damping parameter for the damped Jacobi relaxation algorithm chosen in line 14. Most of the parameters have reasonable default values, so we only need to change what is necessary. An instance of the solver is constructed in line 25 for the assembled sparse matrix in the Compressed Row Storage (CRS) format. The instance is then used to solve the system for the given right-hand side in line 33.

Now, if we decided to use another backend, for example, the NVIDIA CUDA one, we would only need to change the definition of the `Backend` type in line 7 from `amgcl::backend::builtin<double>` to `amgcl::backend::cuda<double>` and replace the type of the vectors `rhs` and `x` with `thrust::device_vector<double>` in line 28. As a result, the constructed solver would be transferred to the supported GPU, and the solution would be performed on the GPU. In a similar way, to change any of the algorithmic components of the solver, we just need to adjust its type definition in lines 10 to 17.

This simple example already exposes some of the library design choices. AMGCL uses policy-based design [1, 13] for the classes like `make_solver` or `amg`, so that individual algorithmic components of the classes are selected at compile time by the corresponding template parameters. The most important concepts of the library are backends, iterative solvers, and preconditioners. A backend is a unified set of data types and operations on them packaged into a class. Iterative solvers and preconditioners are implemented in terms of backend operations, so that backends are easily switched by changing a template parameter. The algebraic multigrid class `amgcl::amg` is an implementation of the preconditioner concept and in turn depends on concepts of coarsening and relaxation schemes. In listing 1 we use smoothed aggregation method as coarsening and damped Jacobi method as relaxation (lines 13 and 14). Each component in the library defines a parameter structure `params` complete with reasonable default values. Whenever a parent component depends on a child, the child’s parameters are included into the parent’s ones. For example, parameters for `amgcl::make_solver` are defined as shown in listing 2.

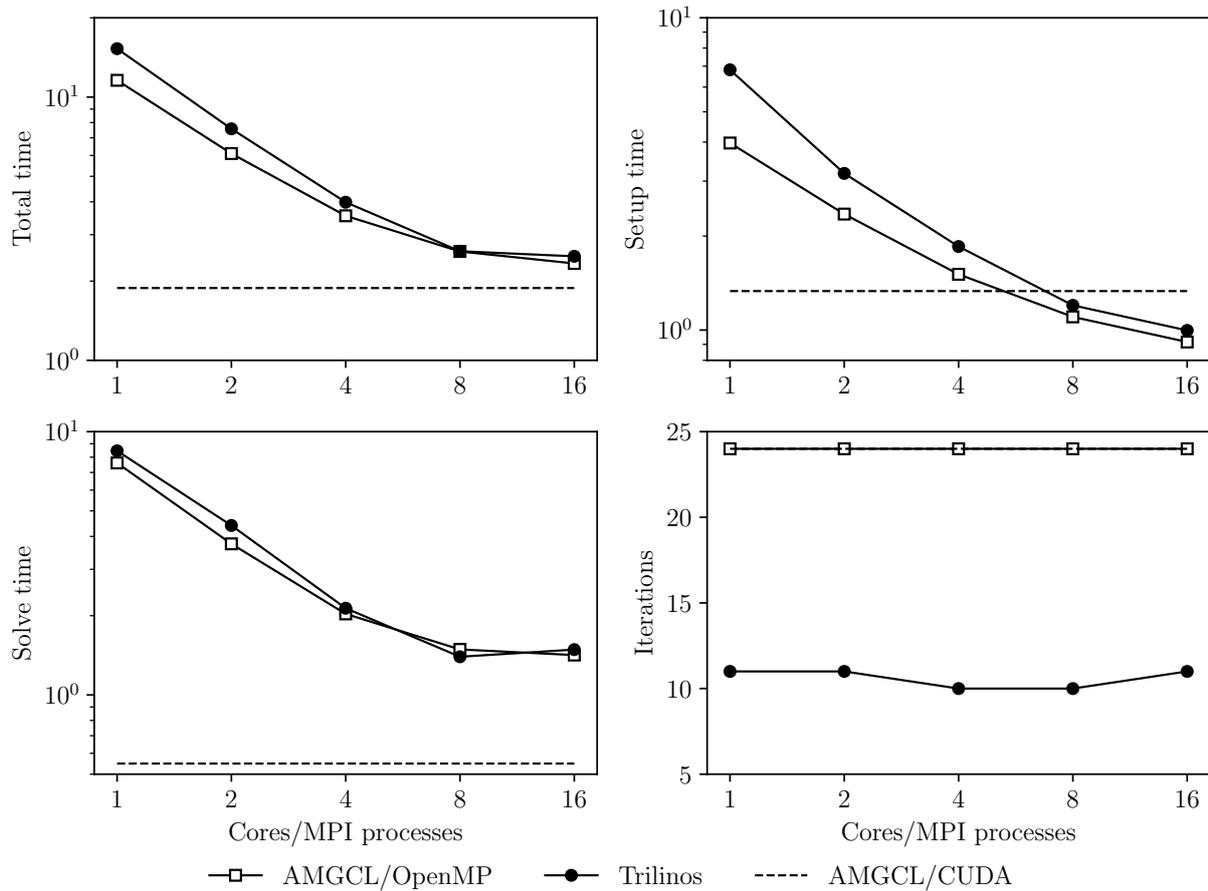


Figure 2. Performance of OpenMP and CUDA backends of AMGCL vs Trilinos ML on a single compute node for solution of a 3D Poisson problem.

This allows to have generic constructor interface for all classes in AMGCL and to seamlessly integrate various concepts together.

The compile-time definition of the solver components allows the compiler to apply efficient optimization techniques, but in practice the complete specification of the algorithm may be problem-dependent and impossible to provide in advance. Thus AMGCL provides a runtime interface that allows the user to select the algorithm components with runtime parameters. The runtime solver identical to the one used in listing 1 is shown in listing 3. The only component that still has to be defined at compile time is the selected backend. An instance of `boost::property_tree::ptree` class is now used for input parameters, but the parameters have the same generic structure as before. An attractive advantage of the Boost.PropertyTree library [23] is that it supports import of the parameters from a number of file formats, such as XML or JSON.

The subdomain deflation approach in AMGCL is implemented using the concepts described above. The components it depends on are a backend, an iterative solver, a local preconditioner, and a direct solver (used to solve the coarse deflated system). All of these may be readily reused from the core library, and we only need to define the global projection method in order to finish the implementation of the algorithm.

An important feature of the deflation approach is that each of the local problems is well posed. Indeed, ignoring non-local matrix coefficients is largely equivalent to applying Dirichlet conditions to all of the non-local unknowns. This allows the use of local preconditioners within every computational domain, and also constitutes an important advantage with respect to classical DD techniques in which the local problems often lack sufficient boundary conditions.

Listing 1. Solving linear system with AMGCL.

```

1 // Assemble system matrix in CRS format:
2 std::vector<int> ptr, col;
3 std::vector<double> val;
4 int n = assemble(ptr, col, val);
5
6 // Select the backend to use:
7 typedef amgcl::backend::builtin<double> Backend;
8
9 // Define solver components:
10 typedef amgcl::make_solver<
11     amgcl::amg<                               // preconditioner:
12         Backend,
13         amgcl::coarsening::smoothed_aggregation,
14         amgcl::relaxation::damped_jacobi
15     >,
16     amgcl::solver::cg<Backend> // iterative solver:
17 > Solver;
18
19 // Set solver parameters:
20 Solver::params prm;
21 prm.solver.tol = 1e-6;
22 prm.precond.relax.damping = 0.8;
23
24 // Setup solver:
25 Solver solve(boost::tie(n, ptr, col, val), prm);
26
27 // The RHS and the solution vectors:
28 std::vector<double> rhs(n, 1.0), x(n, 0.0);
29
30 // Solve the system for the given RHS and initial solution:
31 int iters;
32 double error;
33 boost::tie(iters, error) = solve(rhs, x);

```

Listing 2. Nested definition of solver parameters.

```

1 template <class PrecondType, class SolverType>
2 class make_solver {
3     ...
4     struct params {
5         typename PrecondType::params precond;
6         typename SolverType::params solver;
7     };
8     ...
9 };

```

The global projector, or the deflation operator, is implemented in a matrix free fashion. That is,

$$\mathbf{r}^* := (\mathbf{I} - \mathbf{A}\mathbf{Z}\mathbf{E}^{-1}\mathbf{Z}^T) \mathbf{r} \quad (12)$$

is computed as

$$\mathbf{t}_1 = \mathbf{Z}^T \mathbf{r}, \quad \mathbf{t}_2 = \mathbf{E}^{-1} \mathbf{t}_1, \quad \mathbf{r}^* = \mathbf{r} - (\mathbf{A}\mathbf{Z}) \mathbf{t}_2, \quad (13)$$

where \mathbf{t}_* are temporary vectors used during the computations. It should be noted that in our implementation the $\mathbf{A}\mathbf{Z}$ matrix product and matrix \mathbf{E} are computed explicitly for the sake of performance and to be able to apply state-of-the art solvers to the solution of the coarse matrix \mathbf{E} . Avoiding the explicit computation of \mathbf{E} would be feasible if one goes for inexact solution strategies as discussed in section 2.2.1. This might be mandatory if extreme scalability is to be achieved.

The fact that the deflation basis stored in \mathbf{Z} is by construction purely local to each of the subdomains allows us to optimize the operations in the above matrix-free scheme and reduce the communication cost

Listing 3. Runtime definition of the AMGCL solver.

```

1  typedef amgcl::make_solver<
2      amgcl::runtime::amg<Backend>,
3      amgcl::runtime::iterative_solver<Backend>
4      > Solver;
5
6  boost::property_tree::ptree prm;
7  prm.put("solver.type", "cg");
8  prm.put("solver.tol", 1e-6);
9  prm.put("precond.coarsening.type", "smoothed_aggregation");
10 prm.put("precond.relax.type", "damped_jacobi");
11 prm.put("precond.relax.damping", 0.8);
12
13 Solver solve(boost::tie(n, ptr, col, val), prm);

```

of the algorithm. The same optimization can also be applied to the computation of \mathbf{E} thus improving the overall efficiency of the process.

To reiterate, the algorithm considered in the current paper is equivalent to solution of the system (7) preconditioned with block AMG. Each block is constructed locally on each MPI process, and corresponds to a square diagonal subblock of the system matrix A . There is no overlap between the blocks, and the only communication is happening during the computation of the global projector (12).

4. BENCHMARKS

In this section we demonstrate performance and scalability of the proposed approach on the example of a Poisson problem and a Navier-Stokes problem in a three dimensional space. To provide a reference, we compare performance of the AMGCL library with that of the well-known Trilinos ML package [14]. The source code of the benchmarks is available at a GitHub repository².

The benchmarks were run on MareNostrum 4³, and PizDaint⁴ clusters which we gained access to via PRACE program. The MareNostrum 4 cluster has 3 456 compute nodes, each equipped with two 24 core Intel Xeon Platinum 8160 CPUs, and 96 GB of RAM. The peak performance of the cluster is 6.2 Petaflops. The PizDaint cluster has 5 320 hybrid compute nodes, where each node has one 12 core Intel Xeon E5-2690 v3 CPU with 64 GB RAM and one NVIDIA Tesla P100 GPU with 16 GB RAM. The peak performance of the PizDaint cluster is 25.3 Petaflops.

4.1. 3D Poisson problem

The first set of our experiment uses the classical 3D Poisson problem [19]. Namely, we look for the solution of the problem

$$-\Delta u = 1, \quad (14)$$

in the unit cube $\Omega = [0, 1]^3$ with homogeneous Dirichlet boundary conditions.

The AMGCL implementation uses a BiCGStab(2) [24] iterative solver preconditioned with subdomain deflation, as it showed the best behaviour in our tests. Smoothed aggregation AMG is used as the local preconditioner. The Trilinos implementation uses a CG solver [2] preconditioned with smoothed aggregation AMG with default “SA” settings, or domain decomposition method with default “DD-ML” settings [14]. All implementations targeted relative residual norm $\varepsilon = \|b - Ax\|/\|b - Ax_0\|$ of 10^{-8} .

Figure 3 shows weak scaling of the solution on the MareNostrum 4 cluster. Here the problem size is chosen to be proportional to the number of CPU cores with about 100^3 unknowns per core. The rows in the figure from top to bottom show total computation time, time spent on constructing the preconditioner, solution time, and the number of iterations. The AMGCL library results are labelled “OMP= n ”, where $n = 1, 4, 12, 24$ corresponds to the number of OpenMP threads controlled by each

² https://github.com/ddemidov/amgcl_benchmarks

³ Barcelona, Spain, <https://www.bsc.es/marenostrum/>

⁴ Lugano, Switzerland, http://www.cscs.ch/computers/piz_daint/

MPI process. The Trilinos library uses single-threaded MPI processes. The Trilinos data is only available for up to 1536 MPI processes, which is due to the fact that only 32-bit version of the library was available on the cluster. The AMGCL data points for 19 200 cores with “OMP=1” are missing because factorization of the deflated matrix becomes too expensive for this configuration. AMGCL plots in the left and the right columns correspond to the linear deflation and the constant deflation correspondingly. The Trilinos and Trilinos/DD-ML lines correspond to the smoothed AMG and domain decomposition variants accordingly and are depicted both in the left and the right columns for convenience.

In the case of ideal scaling the timing plots on this figure would be strictly horizontal. This is not the case here: instead, we see that both AMGCL and Trilinos loose about 6-8% efficiency whenever the number of cores doubles. The proposed approach performs about three times worse than the AMG-based Trilinos version, and about 2.5 times better than the domain decomposition based Trilinos version. This is mostly governed by the number of iterations each version needs to converge.

We observe that AMGCL scalability becomes worse at the higher number of cores. We refer to the table 1 for the explanation. The table presents the profiling data for the solution of the Poisson problem on the MareNostrum 4 cluster. The first two columns show time spent on the setup of the preconditioner and the solution of the problem; the third column shows the number of iterations required for convergence. The “Setup” column is further split into subcolumns detailing the total setup time and the time required for factorization of the coarse system. It is apparent from the table that factorization of the coarse (deflated) matrix starts to dominate the setup phase as the number of subdomains (or MPI processes) grows, since we use a sparse direct solver for the coarse problem. This explains the fact that the constant deflation scales better, since the deflation matrix is four times smaller than for a corresponding linear deflation case.

The advantage of the linear deflation is that it results in a better approximation of the problem on a coarse scale and hence needs less iterations for convergence and performs slightly better within its scalability limits, but the constant deflation eventually outperforms linear deflation as the scale grows.

Figure 4 shows weak scaling of the Poisson problem on the PizDaint cluster. The problem size here is chosen so that each node owns about 200^3 unknowns. On this cluster we are able to compare performance of the OpenMP and CUDA backends of the AMGCL library. Intel Xeon E5-2690 v3 CPU is used with the OpenMP backend, and NVIDIA Tesla P100 GPU is used with the CUDA backend on each compute node. The scaling behavior is similar to the MareNostrum 4 cluster. We can see that the CUDA backend is about 9 times faster than OpenMP during solution phase and 4 times faster overall. The discrepancy is explained by the fact that the setup phase in AMGCL is always performed on the CPU, and in the case of CUDA backend it has the additional overhead of moving the generated hierarchy into the GPU memory. It should be noted that this additional cost of setup on a GPU (and the cost of setup in general) often can be amortized by reusing the preconditioner for different right-hand sides. This is often possible for non-linear or time dependent problems. The performance of the solution step of the AMGCL version with the CUDA backend here is on par with the Trilinos ML package. Of course, this comparison is not entirely fair to Trilinos, but it shows the advantages of using CUDA technology.

Figure 5 shows strong scaling results for the MareNostrum 4 cluster. The problem size is fixed to 512^3 unknowns and ideally the compute time should decrease as we increase the number of CPU cores. The case of ideal scaling is depicted for reference on the plots with thin gray dotted lines.

Here, AMGCL demonstrates scalability slightly better than that of the Trilinos ML package. At 384 cores the AMGCL solution for OMP=1 is about 2.5 times slower than Trilinos/AMG, and 2 times faster than Trilinos/DD-ML. As is expected for a strong scalability benchmark, the drop in scalability at higher number of cores for all versions of the tests is explained by the fact that work size per each subdomain becomes too small to cover both setup and communication costs.

The profiling data for the strong scaling case is shown in the table 2, and it is apparent that, as in the weak scaling scenario, the deflated matrix factorization becomes the bottleneck for the setup phase performance.

An interesting observation is that convergence of the method improves with growing number of MPI processes. In other words, the number of iterations required to reach the desired tolerance decreases with as the number of subdomains grows, since the deflated system is able to describe the main problem better and better. This is especially apparent from the strong scalability results (figure 5), where the problem size remains fixed, but is also observable in the weak scaling case for “OMP=1” (figure 3).

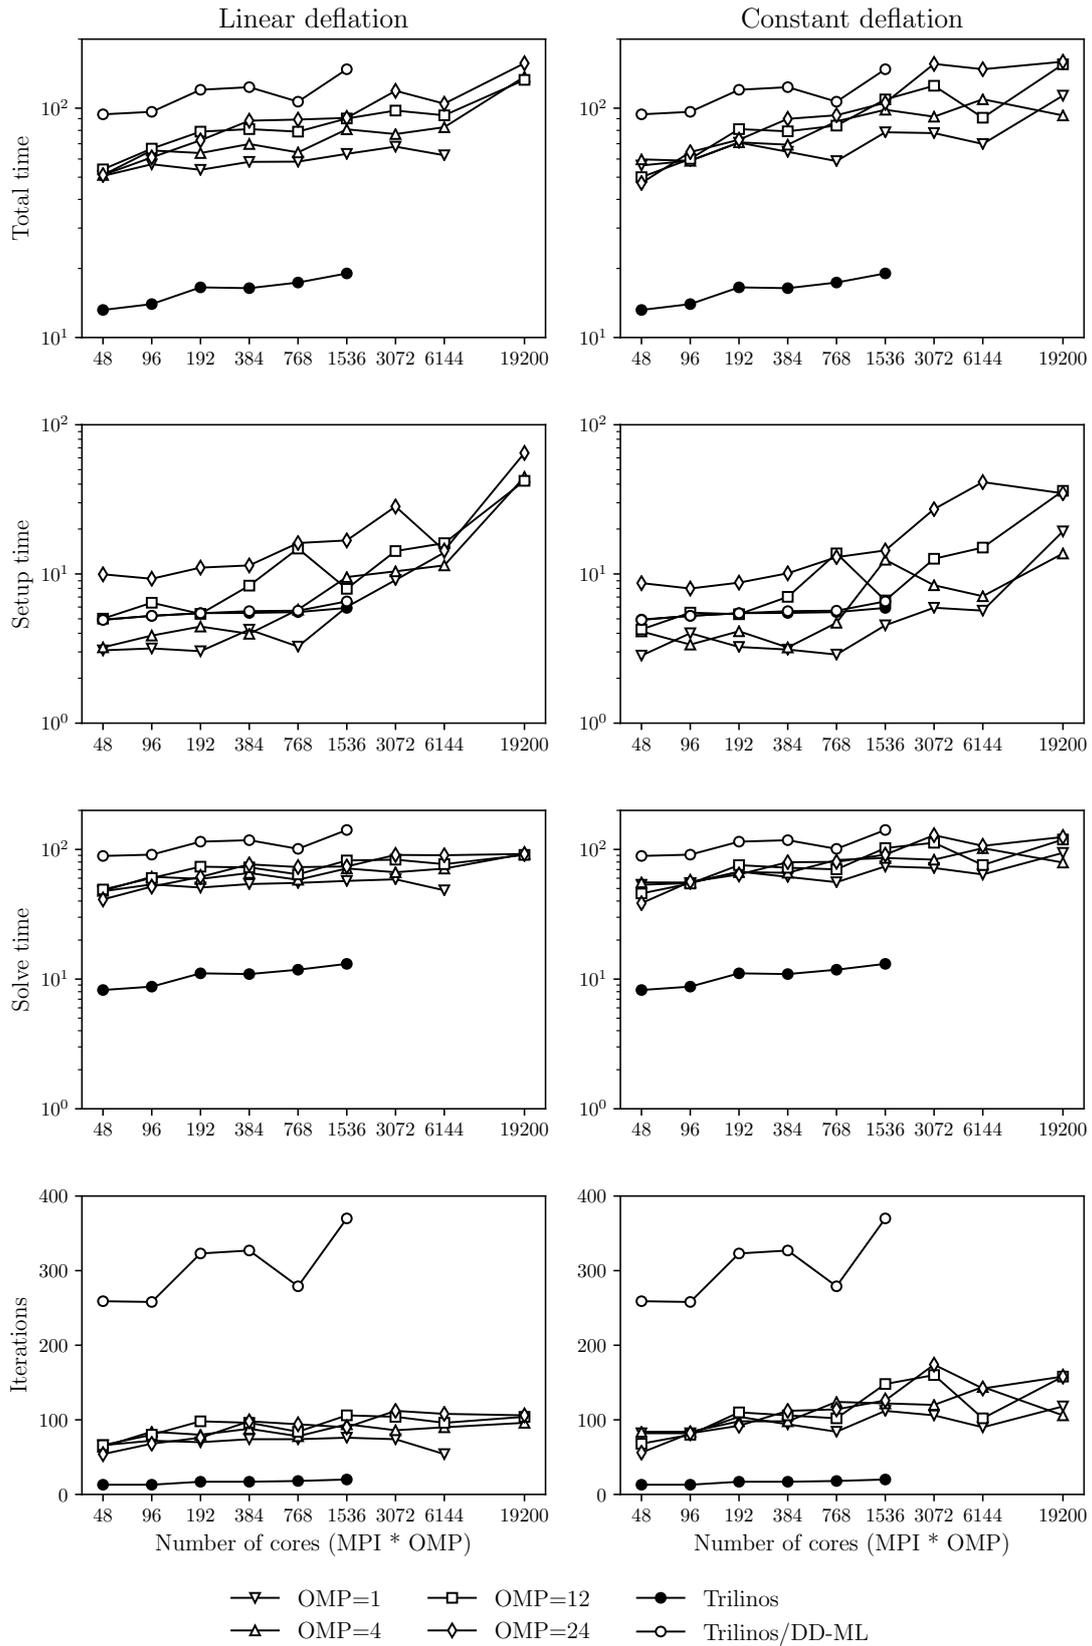


Figure 3. Weak scalability of Poisson problem on the MareNostrum 4 cluster.

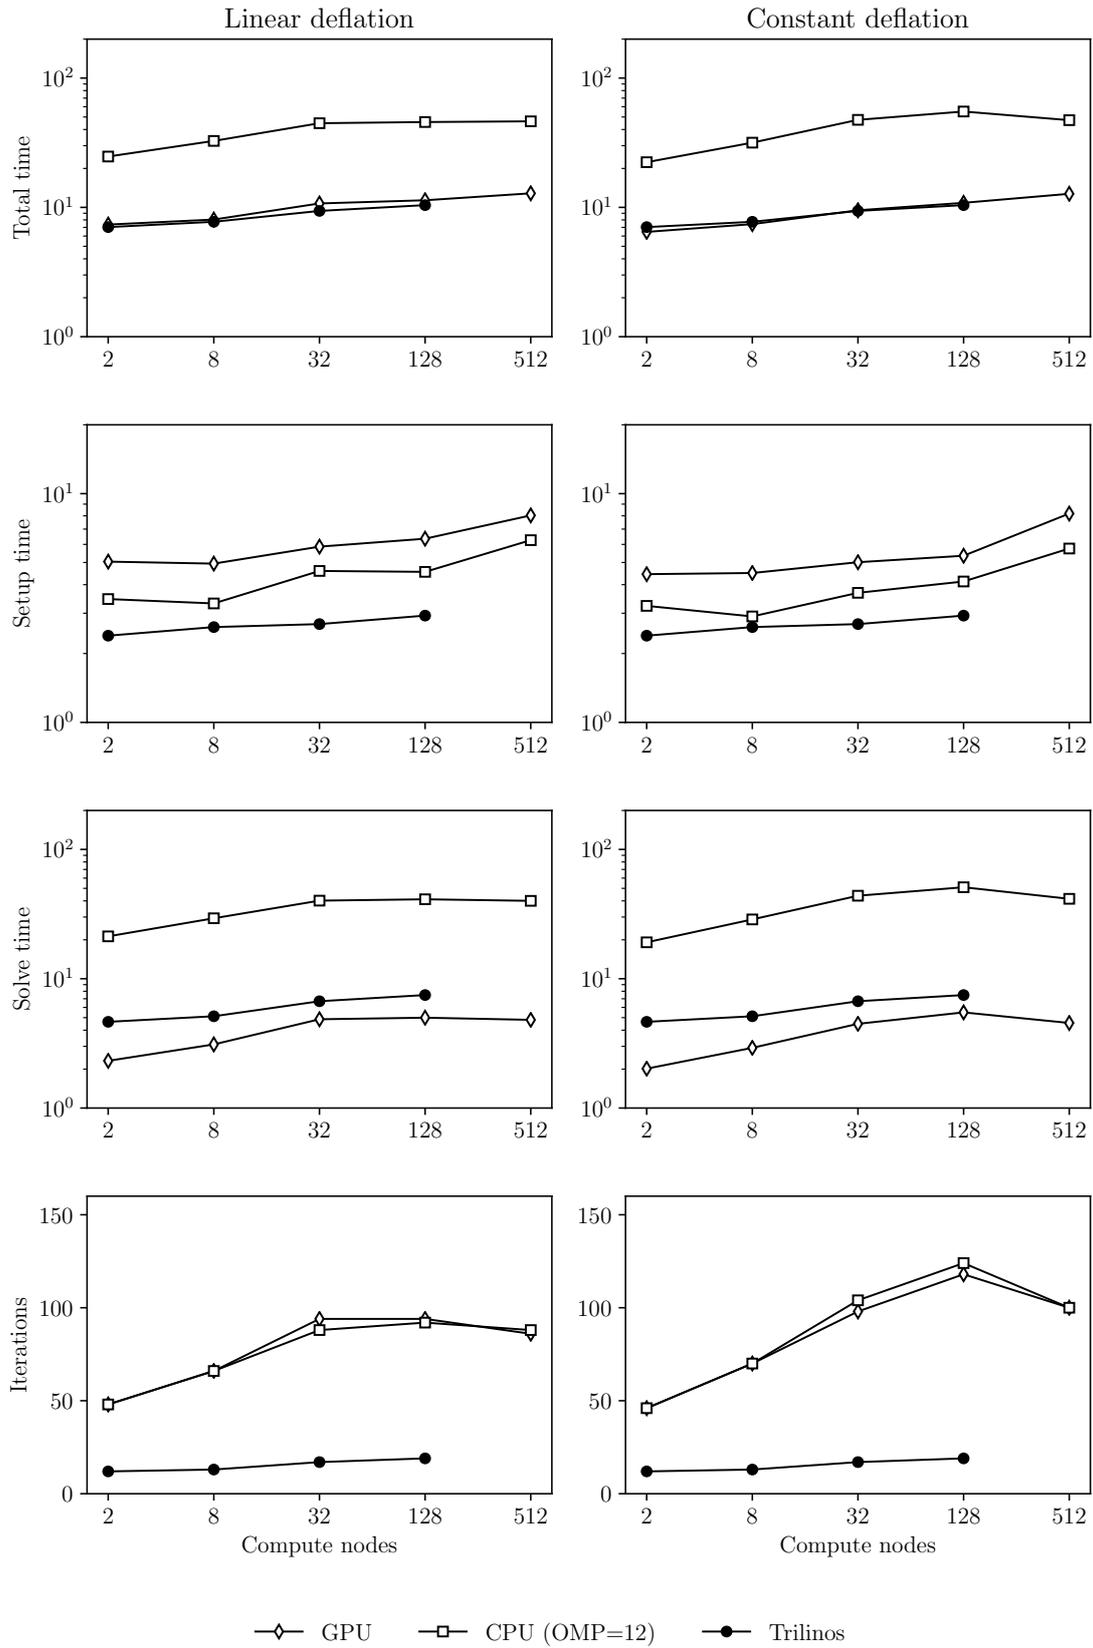


Figure 4. Weak scalability of Poisson problem on the PizDaint cluster.

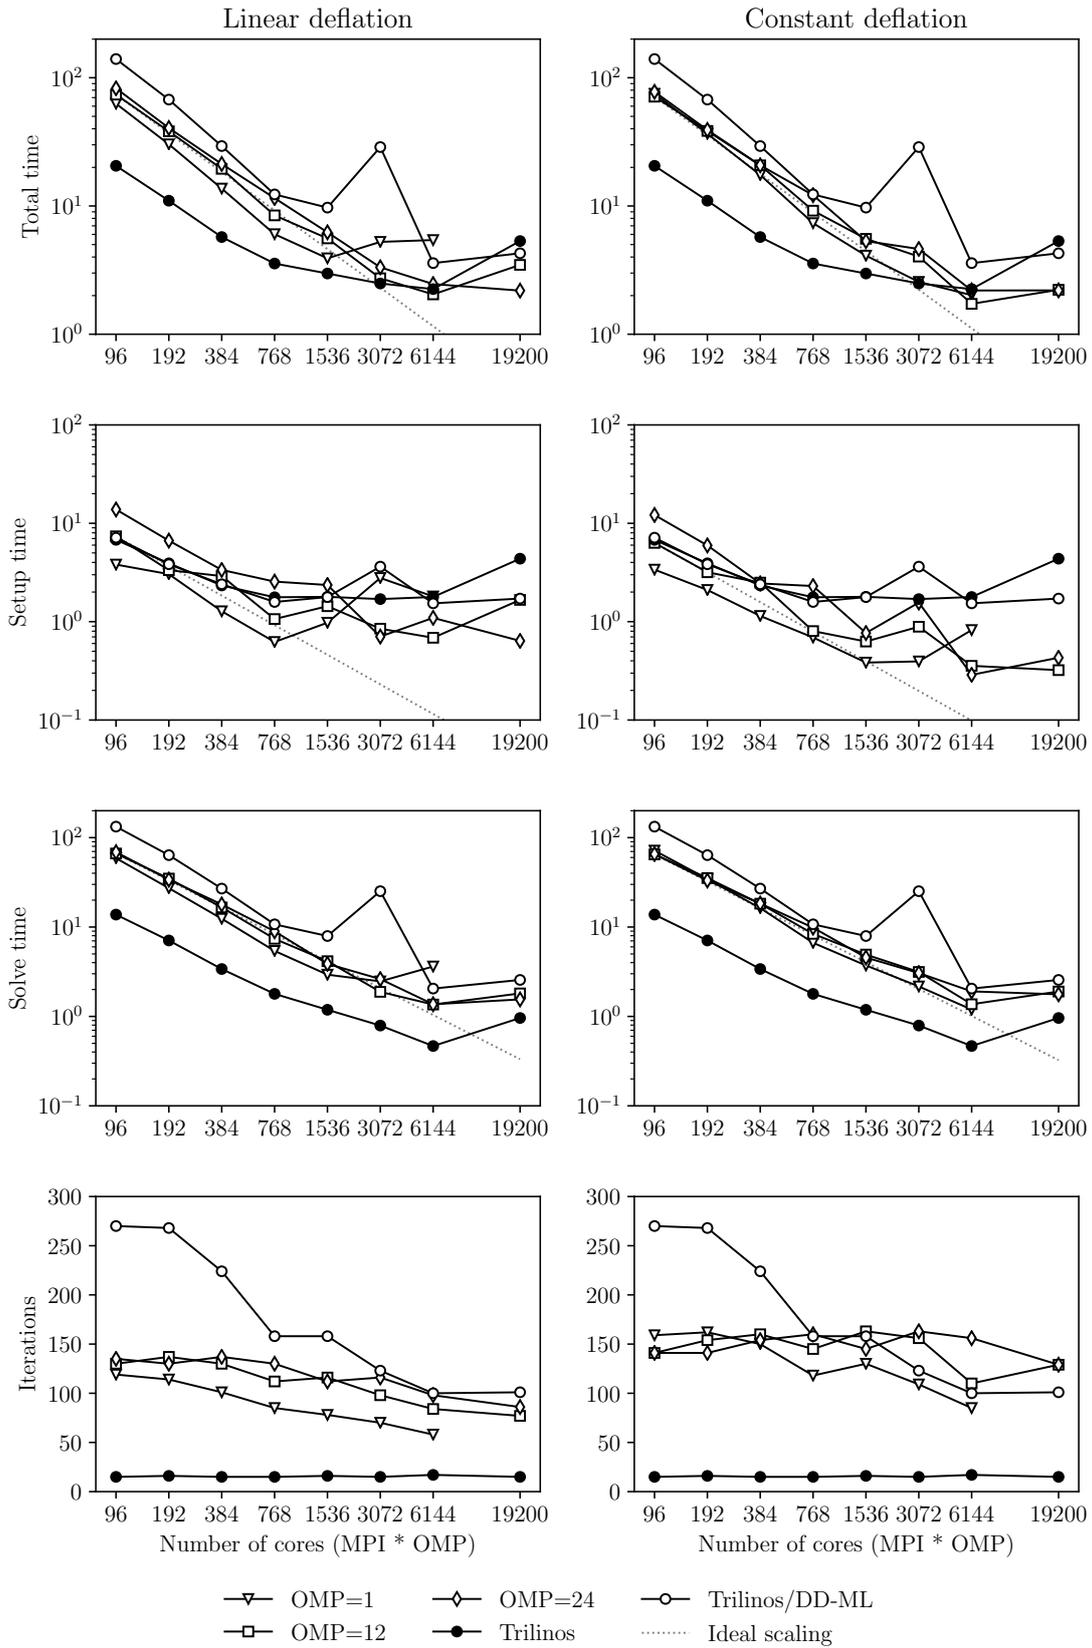


Figure 5. Strong scalability of Poisson problem on the MareNostrum 4 cluster.

Table 1. Profiling data for the weak scaling of the Poisson problem on MareNostrum 4.

Cores	Setup		Solve	Iterations
	Total	Factorize E		
OMP=1		Linear deflation		
384	4.23	0.02	54.08	74
1 536	6.01	0.64	57.19	76
6 144	13.92	8.41	48.40	54
OMP=1		Constant deflation		
384	3.11	0.00	61.41	94
1 536	4.52	0.01	73.98	112
6 144	5.67	0.16	64.13	90
OMP=12		Linear deflation		
384	8.35	0.00	72.68	96
1 536	7.95	0.00	82.22	106
6 144	16.08	0.03	77.00	96
19 200	42.09	1.76	90.74	104
OMP=12		Constant deflation		
384	7.02	0.00	72.25	106
1 536	6.64	0.00	102.53	148
6 144	15.02	0.00	75.82	102
19 200	36.08	0.03	119.25	158

To conclude the benchmark, we compare the results of the proposed deflation solver, which as we recall combines deflation and AMG as local preconditioner (employing the AMGCL OpenMP version), with the result that can be obtained by omitting the use of deflation and using just the local AMG. As shown in figure 6 the use of the local preconditioner alone is not weakly scalable, and leads to a number of iterations that grows, approximately linearly with the problem size. On the contrary, the use of deflation results in a flat number of iterations. Clearly such improvement in the iteration count comes at the cost of solving a global coarse problem, thus implying that each deflation iteration is slightly more costly than applying the local AMG preconditioner. The timings however suggest that the price is reasonably low, and that the use of deflation pays off.

4.2. Navier–Stokes problem

The second test problem is an incompressible Navier–Stokes problem discretized on a non uniform 3D mesh with a finite element method:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p = \mathbf{b}, \quad (15a)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (15b)$$

The problem is discretized using an equal-order tetrahedral Finite Elements, stabilized employing an ASGS-type (algebraic subgrid-scale) approach [10]. This results in a discretized linear system of

Table 2. Profiling data for the strong scaling of the Poisson problem on MareNostrum 4.

Cores	Setup		Solve	Iterations
	Total	Factorize E		
OMP=1		Linear deflation		
384	1.27	0.02	12.39	101
1 536	0.97	0.45	2.93	78
6 144	9.09	8.44	3.61	58
OMP=1		Constant deflation		
384	1.14	0.00	16.30	150
1 536	0.38	0.01	3.71	130
6 144	0.82	0.16	1.19	85
OMP=12		Linear deflation		
384	2.90	0.00	16.57	130
1 536	1.43	0.00	4.15	116
6 144	0.68	0.03	1.35	84
19 200	1.66	1.29	1.80	77
OMP=12		Constant deflation		
384	2.49	0.00	18.25	160
1 536	0.62	0.00	4.91	163
6 144	0.35	0.00	1.37	110
19 200	0.32	0.02	1.89	129

equations with a block structure of the type

$$\begin{pmatrix} \mathbf{K} & \mathbf{G} \\ \mathbf{D} & \mathbf{S} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{b}_u \\ \mathbf{b}_p \end{pmatrix} \quad (16)$$

where each of the matrix subblocks is a large sparse matrix, and the blocks \mathbf{G} and \mathbf{D} are non-square. The overall system matrix for the problem was assembled in the Kratos⁵ multi-physics package [6, 7] developed in CIMNE, Barcelona. The matrix contains 4 773 588 unknowns and 281 089 456 nonzeros. Such problem is routinely solved within Kratos by employing the ML solver with the default “NSSA” settings. From the developer experience, this constitutes the fastest “out-of-the-box” solver option, which is used here as the reference line. Here we explored the use of the Schur complement solver capabilities available within AMGCL. The pressure variables are identified by a “pressure-mask” array, and the AMGCL solver constructs an inexact pressure-Schur preconditioner [11, 15, 30]. The logical matrix blocks in (16) are explicitly assembled for the sake of performance. Each application of the preconditioner corresponds to the following steps:

1. Inexact solve of $\mathbf{K}\hat{\mathbf{u}} = \mathbf{b}_u - \mathbf{G}\mathbf{p}$.

⁵ <http://www.cimne.com/kratos/>

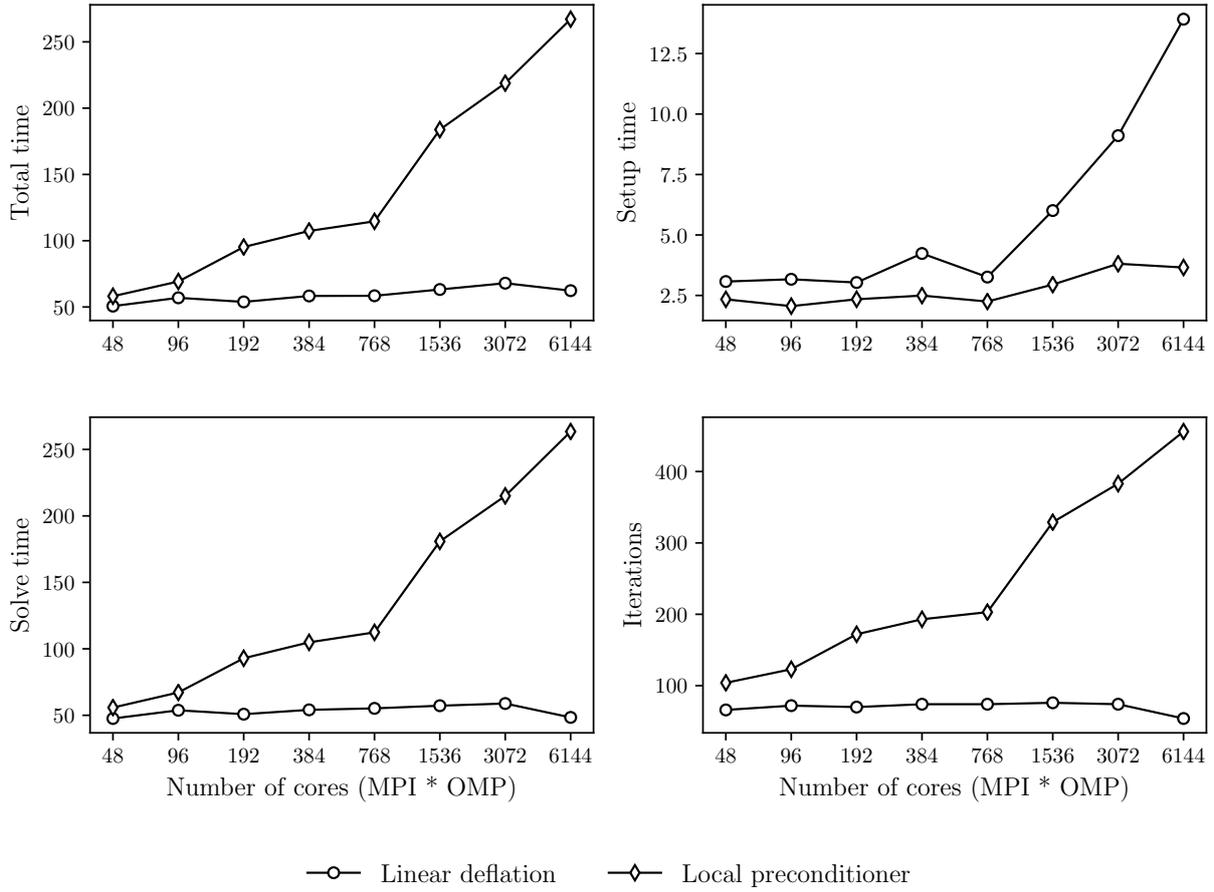


Figure 6. Weak scaling of the Poisson problem. The proposed approach is compared with the case when AMG is used as a purely local preconditioner. OMP=1 in both cases.

2. Inexact matrix-free solve of $(\mathbf{S} - \mathbf{D} \text{diag}(\mathbf{K})^{-1} \mathbf{G}) \mathbf{p} = \mathbf{b}_p - \mathbf{D}\hat{\mathbf{u}}$.
3. Inexact solve of $\mathbf{K}\mathbf{u} = \mathbf{b}_u - \mathbf{G}\mathbf{p}$.

A purely local SPAI-0 (sparse approximate inverse) preconditioner [4] is employed in the solution of the first and third step. Deflation in combination with local AMG is employed as a preconditioner in the second step (the solution of the pressure Schur complement). Here we choose $\mathbf{M} \approx \mathbf{S}^{-1}$ so that each application of the preconditioner corresponds to an approximate solution of a linear problem having \mathbf{S} as a system matrix. Since a low tolerance is targeted in the application of such preconditioner, a flexible GMRES (FGMRES) solver is employed in the solution of the matrix-free problem, to ensure that a preconditioner delivering a varying level of precision can be correctly handled. FGMRES is also used in the top level in targeting the overall solution procedure. The complete preconditioner is composed from the reusable components provided by the AMGCL library, and listing 4 shows the definition of the resulting solver class as used in our benchmarks. The class is using runtime interface, and listing 5 shows JSON file containing AMGCL options used for the tests. In each variant of the solution we used relative residual norm of 10^{-4} as exit condition.

Figure 7 shows scalability results for solving problem (15) on the MareNostrum 4 cluster. Since we are solving a fixed-size problem, this is essentially a strong scalability test. It should be noted that Trilinos ML [14] does not provide field-split type preconditioners, which is why the Trilinos benchmark was performed using the non-symmetric smoothed aggregation variant (NSSA) applied to the monolithic problem. Default NSSA parameters were employed in the tests.

Listing 4. AMGCL class definition for the Navier-Stoker solver.

```

1  typedef
2      amgcl::mpi::make_solver<
3      amgcl::mpi::schur_pressure_correction<    // top level preconditioner
4          amgcl::mpi::make_solver<            // flow block solver
5              amgcl::mpi::block_preconditioner<
6                  amgcl::runtime::relaxation::as_preconditioner<Backend>
7              >,
8              amgcl::runtime::iterative_solver
9          >,
10         amgcl::mpi::subdomain_deflation<    // pressure block solver
11             amgcl::runtime::amg<Backend>,
12             amgcl::runtime::iterative_solver,
13             amgcl::runtime::mpi::direct_solver<double>
14         >
15     >,
16     amgcl::runtime::iterative_solver        // top level iterative solver
17 >
18 Solver;

```

Listing 5. JSON file with AMGCL options used to solve the Navier-Stokes test problem.

```

1  {
2      "solver": {
3          "type" : "fgmres",
4          "M" : 50,
5          "tol" : 1e-4
6      },
7      "precond": {
8          "usolver": {
9              "solver": {
10                 "type" : "gmres",
11                 "tol" : 1e-3,
12                 "maxiter" : 5
13             }
14         },
15         "psolver": {
16             "isolver": {
17                 "type" : "fgmres",
18                 "tol" : 1e-2,
19                 "maxiter" : 20
20             },
21             "local" : {
22                 "coarse_enough" : 500
23             }
24         }
25     }
26 }

```

Both AMGCL and ML preconditioners deliver a very flat number of iterations with growing number of MPI processes. As expected, the pressure-Schur approach pays off and the tested approach performs better than the monolithic approach in the solution of the problem. Overall the AMGCL implementation shows a decent, although less than optimal parallel scalability. This is not unexpected since the problem size quickly becomes too little to justify the use of more parallel resources (note that at 192 processes, less than 25 000 unknowns are assigned to each MPI subdomain). Unsurprisingly, in this context the use of OpenMP within each domain pays off and allows delivering a greater level of scalability.

5. CONCLUSION

The work presents a scalable preconditioner for sparse distributed linear systems in the form of a combination of the subdomain deflation method and local algebraic multigrid. The implementation of the algorithm is included as a reusable component into the open source AMGCL library. The

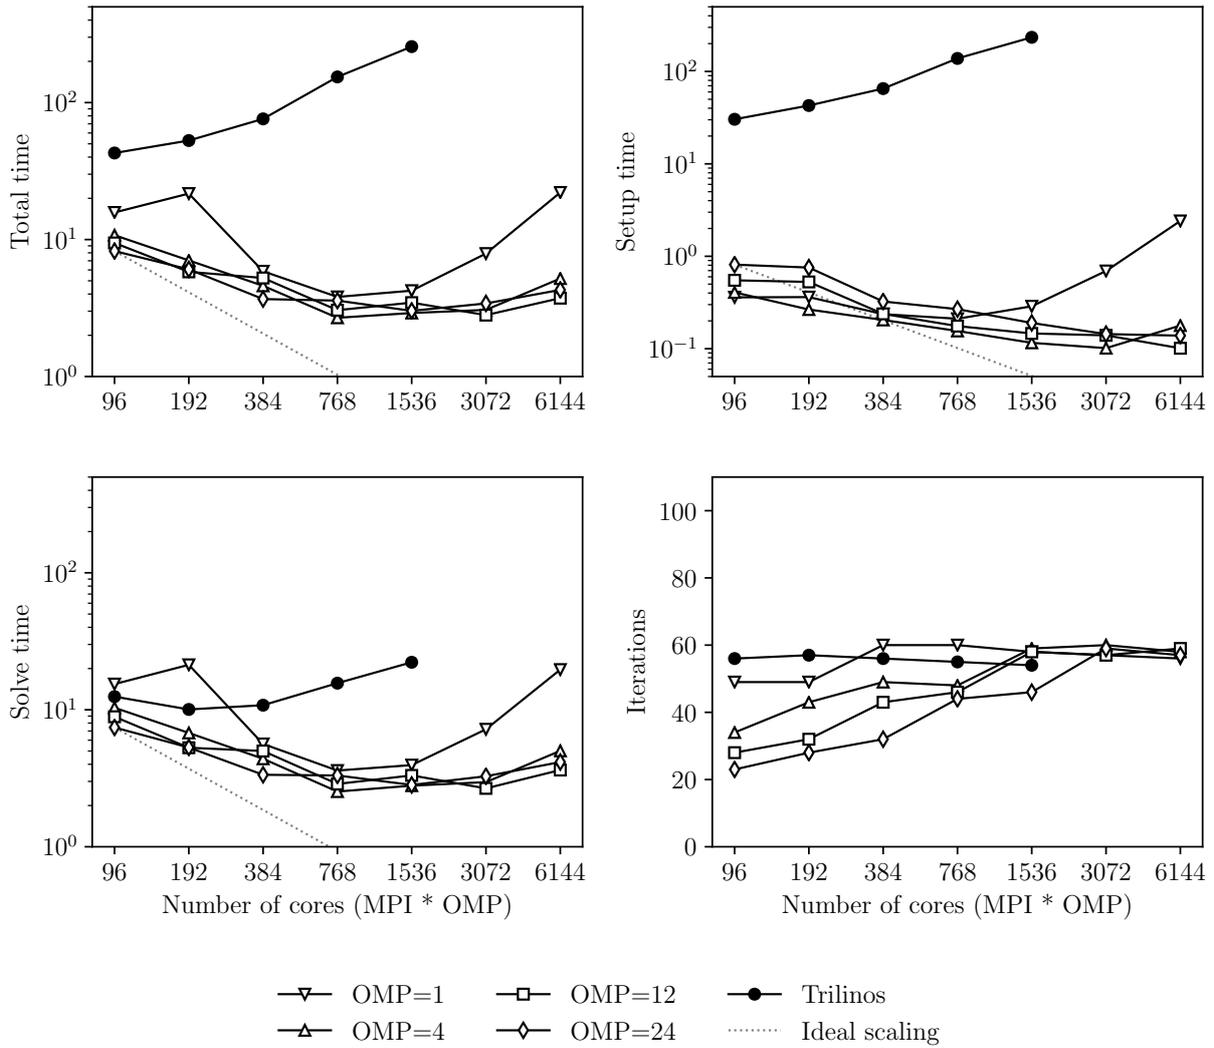


Figure 7. Strong scalability of Navier-Stokes problem on the MareNostrum 4 cluster.

results presented prove that the solver enjoys good weak scalability up to 19 200 cores, the maximum core count available during the testing. Strong scalability tests show the solver efficiency for fixed size problems. The solution targets both homogeneous (CPU-only) and heterogeneous (CPU/GPU) systems, employing hybrid MPI/OpenMP approach in the former and MPI/CUDA in the latter cases. The use of a heterogeneous cluster with GPUs installed on each node allows to gain an overall 4x speedup for the complete solution of the benchmark problem under consideration. The speedup factor results from a much faster solution phase (around 10x speedup) and of a slightly slower setup phase.

To the best of our knowledge these tests represent some of the largest systematic benchmarks of deflation approaches available in the literature and allow to evaluate the efficiency of the proposed approach against reference libraries. The outcome of the investigation is that the proposed deflation approach, although faster than alternative implementation of a similar idea in form of DD-ML preconditioner in the Trilinos library, is still inferior to AMG when applied to the solution of Poisson-like problems. The limitation here is related to a (much) slower performance with respect to AMG at small core counts which is maintained as the number of cores grows. This result ultimately pushed the authors to implement a fully MPI-parallel AMG within the AMGCL library. The results of this effort fall outside of the scope of the current paper and can be found in [9].

From the scalability analysis it can be observed that the use of OpenMP or CUDA in combination with MPI proves important in achieving scalability at high core counts. The reasons for such behavior are easily identified in the smaller size of the deflation matrix \mathbf{E} to be inverted, and in the minor MPI communication volume. Taking advantage of such a feature, we believe that larger core counts are possible. For example, by using 48 cores per MPI on MareNostrum 4 one may target using 100K processor cores with a deflation matrix of size 2000, easily within the scalable range. However, such extreme scalability scenarios are still untested and will require some additional effort in ironing out performance bottlenecks. Targeting such scenarios will be the focus of future works.

6. ACKNOWLEDGMENTS

Contribution of Dr. Demidov was funded by the state assignment to the Joint Supercomputer Center of the Russian Academy of Sciences for scientific research and RFBR grant No 18-07-00964.

Dr. Rossi acknowledges the financial support to CIMNE via the CERCA Programme/Generalitat de Catalunya and the support of the ExaQUte FetHPC project GA 800898.

The authors thankfully acknowledge the support of the PRACE program (project 2010PA4058), in providing access to the MareNostrum 4 and PizDaint clusters. Without such resources the testing would not have been possible. The help of Prof. Labarta of the POP Center of Excellence in improving the NUMA scalability of the solver is also gratefully acknowledged.

REFERENCES

1. Andrei Alexandrescu. *Modern C++ design: generic programming and design patterns applied*. Addison-Wesley, 2001.
2. Richard Barrett, Michael Berry, Tony F Chan, James Demmel, June Donato, Jack Dongarra, Victor Eijkhout, Roldan Pozo, Charles Romine, and Henk Van der Vorst. *Templates for the solution of linear systems: building blocks for iterative methods*. SIAM, 1994.
3. A Brandt, S McCoruick, and J Hüge. Algebraic multigrid (amg) for sparse matrix equations. *Sparsity and its Applications*, 257, 1985.
4. Oliver Bröker and Marcus J Grote. Sparse approximate inverse smoothers for geometric and algebraic multigrid. *Applied numerical mathematics*, 41(1):61–80, 2002.
5. Andrew J Cleary, Robert D Falgout, Van Emden Henson, Jim E Jones, Thomas A Manteuffel, Stephen F McCormick, Gerald N Miranda, and John W Ruge. Robustness and scalability of algebraic multigrid. *SIAM Journal on Scientific Computing*, 21(5):1886–1908, 2000.
6. P. Dadvand, R. Rossi, M. Gil, X. Martorell, J. Cotela, E. Juanpere, S. R. Idelsohn, and E. Oñate. Migration of a generic multi-physics framework to HPC environments. *Computers and Fluids*, 80(1):301–309, jul 2013.
7. P. Dadvand, R. Rossi, and E. Oñate. An object-oriented environment for developing finite element codes for multi-disciplinary applications. *Archives of Computational Methods in Engineering*, 17(3):253–297, 2010.
8. Denis Demidov. AMGCL: a C++ library for solution of large sparse linear systems with algebraic multigrid method, 2017.
9. Denis Demidov. Amgcl: An efficient, flexible, and extensible algebraic multigrid implementation. *Lobachevskii Journal of Mathematics*, 40(5):535–546, May 2019.
10. J. (Jean) Donéa and Antonio Huerta. *Finite element methods for flow problems*. Wiley, 2003.
11. H C Elman. Preconditioning Strategies for Models of Incompressible Flow. *Journal of Scientific Computing*, 25(9):347–366, 2005.
12. Jason Frank and C Vuik. On the construction of deflation-based preconditioners. *SIAM Journal on Scientific Computing*, 23(2):442–462, 2001.
13. Erich Gamma. *Design patterns: elements of reusable object-oriented software*. Pearson Education India, 1995.
14. M W Gee, C M Siefert, J J Hu, R S Tuminaro, and M G Sala. ML 5.0 Smoothed Aggregation User’s Guide. Technical Report SAND2006-2649, Sandia National Laboratories, 2006.
15. Björn Gmeiner, Markus Huber, Lorenz John, Ulrich Rüde, and Barbara Wohlmuth. A quantitative performance study for stokes solvers at the extreme scale. *Journal of Computational Science*, 17:509–521, 2016.
16. Pascal Hénon, Pierre Ramet, and Jean Roman. PASTIX: a high-performance parallel direct solver for sparse symmetric positive definite systems. *Parallel Computing*, 28(2):301–321, 2002.

17. Jonathan Hogg and Jennifer Scott. New parallel sparse direct solvers for multicore architectures. *Algorithms*, **6**(4):702–725, 2013.
18. Paul Lin, Matthew Bettencourt, Stefan Domino, Travis Fisher, Mark Hoemmen, Jonathan Hu, Eric Phipps, Andrey Prokopenko, Sivasankaran Rajamanickam, Christopher Siefert, et al. Towards extreme-scale simulations for low mach fluids with second-generation trilinos. *Parallel processing letters*, **24**(04):1442005, 2014.
19. Jon Mathews and Robert Lee Walker. *Mathematical methods of physics*, volume 501. WA Benjamin New York, 1970.
20. R Nabben and C Vuik. A comparison of deflation and the balancing preconditioner. *SIAM Journal on Scientific Computing*, **27**:1742–1759, 2006.
21. John W Ruge and Klaus Stüben. Algebraic multigrid. In *Multigrid methods*, pages 73–130. SIAM, 1987.
22. Yousef Saad. *Iterative methods for sparse linear systems*. Siam, 2003.
23. Boris Schäling. *The Boost C++ libraries*. XML Press, Aufl. Laguna Hills, Calif, 2014.
24. Gerard L G Sleijpen and Diederik R Fokkema. BiCGstab (l) for linear equations involving unsymmetric matrices with complex spectrum. *Electronic Transactions on Numerical Analysis*, **1**:11–32, 1993.
25. Barry Smith, Petter Bjorstad, and William Gropp. *Domain decomposition: parallel multilevel methods for elliptic partial differential equations*. Cambridge university press, 2004.
26. K Stuben. Algebraic multigrid (AMG): an introduction with applications. GMD Report 70, GMD, Sankt Augustin, Germany, 1999.
27. J.M. Tang, R Nabben, C Vuik, and Y.A. Erlangga. Comparison of two-level preconditioners derived from deflation, domain decomposition and multigrid methods. *SIAM Journal on Scientific Computing*, **39**:340–370, 2009.
28. JM Tang and Cornelis Vuik. *New variants of deflation techniques for bubbly flow problems*. Delft University of Technology, 2006.
29. U Trottenberg, C Oosterlee, and A Schüller. *Multigrid*. Academic Press, London, 2001.
30. Stefan. Turek. *Efficient solvers for incompressible flow problems : an algorithmic and computational approach*. Springer-Verlag, 1999.
31. C Vuik, A Segal, and J A Meijerink. An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients. *Journal of Computational Physics*, **152**(1):385–403, 1999.
32. Ulrike Meier Yang and Others. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, **41**(1):155–177, 2002.