

A staggered-grid multilevel ILU for steady incompressible flows

Saturday 19th October, 2019, 12:33

Sven Baars^a, Mark van der Kloek^a, Jonas Thies^b, Fred W. Wubs^a

^aUniversity of Groningen, Groningen, the Netherlands

^bGerman Aerospace Center (DLR), Simulation and Software Technology, Linder Höhe, 51147 Cologne, Germany

Abstract

We present a parallel fully coupled multi-level incomplete factorization preconditioner for the 3D stationary incompressible Navier-Stokes equations on a structured grid. The algorithm and software are based on the robust two-level method developed in [1]. In this paper, we identify some of the weak spots of the two-level scheme and propose remedies such as a different domain partitioning and recursive application of the method. We apply the method to the well-known 3D lid-driven cavity benchmark problem, and demonstrate its superior robustness by comparing with a segregated SIMPLE-type preconditioner.

Keywords: Parallel, incomplete factorization, multi-level, incompressible Navier-Stokes, \mathcal{F} -matrix, lid-driven cavity

1. Introduction

To study the stability of steady states in fluid flow problems as a function of a parameter, e.g. the Reynolds Number (see [2] for an introduction to such problems), a continuation method [3] can be used. To apply such a method to high-dimensional systems, we require an efficient and robust method for solving linear systems associated with the discretized incompressible Navier-Stokes equations. The most elegant way to do this is to replace the complete LU factorization by an accurate incomplete one, which can be used as a preconditioner in an iterative procedure. Moreover, this preconditioner can be used to find approximate eigenvalues and eigenvectors of the Jacobian matrix during the continuation process. By an appropriate ordering technique and dropping procedure, one can construct an incomplete LU (ILU) factorization that yields grid independent convergence behavior and approaches an exact factorization as the amount of allowed fill is increased.

The incompressible Navier-Stokes equations can be written as

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{1}{\text{Re}} \Delta \mathbf{u}, \\ \nabla \cdot \mathbf{u} &= 0, \end{aligned} \tag{1}$$

where $\text{Re} = \frac{\rho U D}{\mu}$ is the Reynolds number, ρ is the density and μ is the dynamic viscosity, and D and U are characteristic length and velocity scales of the flow.

These equations are discretized using a second-order symmetry-preserving finite volume method on an Arakawa C-grid; see Figure 1. The discretization leads to a system of ordinary differential equations (ODEs)

$$M \frac{d\mathbf{u}}{dt} + N(\mathbf{u}, \mathbf{u}) = -G\mathbf{p} + \frac{1}{\text{Re}}L\mathbf{u} + \mathbf{f}_u,$$

$$-G^T\mathbf{u} = \mathbf{f}_p,$$

where \mathbf{u} and \mathbf{p} now represent the velocity and pressure in each grid point, $N(\cdot, \cdot)$ is the bilinear form arising from the convective terms, L is the discretized Laplace operator, G is the discretized gradient operator, M is the mass matrix, which contains the volumes of the grid cells on its diagonal and \mathbf{f} contains the known part of the boundary conditions. Note that minus the transpose of the gradient operator gives us the divergence operator, which is a property that is exploited by our method.

The aforementioned $N(\mathbf{u}, \mathbf{v})$ is the discretized form of $\mathbf{u} \cdot \nabla \mathbf{v}$. We indicate this expression by $N_1(\mathbf{u})\mathbf{v}$. In the Jacobian we also have an additional contribution resulting from $N(\mathbf{u}, \mathbf{v})$, i.e. the discretization of $\mathbf{v}(\nabla \cdot \mathbf{u})$, which we indicate by $N_2(\mathbf{u})\mathbf{v}$. It holds that

$$N(\mathbf{u}, \mathbf{v}) = N_1(\mathbf{u})\mathbf{v} = N_2(\mathbf{v})\mathbf{u}.$$

Using this notation, the linear system of saddle point type [4] that has to be solved within each Newton step is given by

$$\begin{pmatrix} N_1(\mathbf{u}) + N_2(\mathbf{u}) - \frac{1}{\text{Re}}L & G \\ G^T & 0 \end{pmatrix} \begin{pmatrix} \Delta \mathbf{u} \\ \Delta \mathbf{p} \end{pmatrix} = - \begin{pmatrix} \mathbf{f}_u \\ \mathbf{f}_p \end{pmatrix}. \quad (2)$$

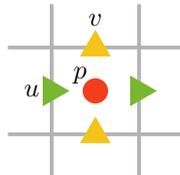


Figure 1: Positioning of unknowns in an Arakawa C-grid

It is known [5], that, on closed domains, dissipation of kinetic energy only occurs by diffusion. Our discretization preserves this property, which has the consequence that for divergence free \mathbf{u} , the operator $N_1(\mathbf{u})$ is skew-symmetric. This means that if we omit $N_2(\mathbf{u})$, the approximate Jacobian will become negative definite on the space of discrete divergence free velocities. Omitting $N_2(\mathbf{u})$ greatly simplifies the solution process since definiteness is a condition under which many standard iterative methods converge. This is a simplification that many authors make, and results in replacing the Newton iteration by a Picard iteration [6]. This works well for reasonably low Reynolds numbers, and far away from bifurcation points where steady solutions become unstable. The Picard iteration, however, seriously impairs the convergence rate of the nonlinear iteration [7].

Similarly, some authors use time dependent approaches to study the stability of steady states [8]. This approach, however, also requires some tricks to obtain the desired information.

Since we want to study precisely the phenomena where the above methods experience difficulty, we would rather use the full Jacobian matrix of the nonlinear equations and apply Newton's method directly.

There are many methods that are based on segregation of physical variables which can solve the linear equations that arise in every Newton iteration. In this approach the system is split into subproblems of an easier form for which standard methods exist. The segregation can already be done at the discretization level, e.g. by doing a time integration and solving a pressure-correction equation independently of the momentum equations [5, 9]. Another class of methods performs the segregation during the linear system solve, often in a preconditioning step. Important are physics based preconditioners [10, 11, 12, 13], which try to split the problem in subsystems which capture the bulk of the physics. The subsystems are again solved by iterative procedures, e.g. algebraic multigrid (AMG) for Poisson-like equations. These methods consist of nested loops for: the nonlinear iteration, iterations over the coupled system, and iterations over the subsystems. The stopping criteria of all these different iterations have to be tuned to make the solver efficient. Furthermore, the total number of iterations in the innermost loop is given by the product of the number of iterations performed on all three levels of iteration and thus easily becomes excessive. This is a major problem when trying to achieve extreme parallelism, because each innermost iteration typically requires global communication in inner products. The number of levels of nested iteration may increase even more if additional physical phenomena are added [10, 14]. Our ultimate aim is to get rid of the inner iterations altogether and to solve the nonlinear equations using a subspace accelerated inexact Newton method. In [15] we did this for simple eigenvalue problems using the Jacobi–Davidson method, which is in fact an accelerated Newton method. The method we present here will make this feasible for the 3D Navier–Stokes equations.

To achieve this, fully aggregated approaches are important. In this category, multigrid and multilevel ILU methods for systems of PDEs exist (see [16, 17] and references therein). The former is attractive but for those methods smoothers may fail due to a loss of diagonal dominance for higher Reynolds Numbers, for which a common solution is to use time integration [18]. The latter comprise AMG algorithms and multilevel methods like MRILU [19] and ILUPACK [20]. ILUPACK has been successful in many fields since it uses a bound to preclude very unstable factorizations. However, this method does not show grid independence for Navier–Stokes problems and is difficult to parallelize [21]. It works well for large problems, but not yet beyond a single shared memory system.

In this paper, we present a novel multilevel preconditioning method which is specially designed for the 3D Navier–Stokes equations. In Section 2, we first describe the two-level ILU preconditioner as introduced in [1] and [22]. After this, we generalize the two-level method to a multilevel method in Section 3. To make this method work for the 3D Navier–Stokes equations, we introduce a skew partitioning method in Section 4. In Section 5 we then discuss the scalability and general performance of the method, and compare it to a popular physics based method, after which we finalize by providing a summary and discussion in Section 6.

2. The two-level ILU preconditioner

In [1] a robust parallel two-level method was developed for solving

$$A\mathbf{x} = \mathbf{b},$$

with $A \in \mathbb{R}^{n \times n}$ for a class of matrices known as \mathcal{F} -matrices. An \mathcal{F} -matrix is a matrix of the form

$$A = \begin{pmatrix} K & B \\ B^T & 0 \end{pmatrix},$$

with K symmetric positive definite and B such that it has at most two entries per row and the entries in each row sum up to 0, as is the case for our gradient matrix G [23, 24]. It was shown that the two-level preconditioner leads to grid-independent convergence for the Stokes equations on an Arakawa C-grid, and that the method is robust even for the Navier–Stokes equations, which strictly speaking do not yield \mathcal{F} -matrices because K becomes nonsymmetric and possibly indefinite.

Rather than reviewing the method and theory in detail, we will only briefly present it here. For details, see [1] and [22].

To simplify the discussion, we focus on the special case of the 3D incompressible Navier–Stokes equations in a cube-shaped domain, discretized on an Arakawa C-grid (see Figure 1). We refer to the velocity variables, which are located on the cell faces as V -nodes, and to the the pressure, which is located in the cell center, as P -node. The variables are numbered cell-by-cell, i.e. the first three variables are the $u/v/w$ -velocity at the north/east/top face of the cell in the bottom south west corner of the domain, and variable four is the P -node in its center. Appropriate boundary conditions (e.g. Dirichlet conditions) are easily implemented in this situation. We remark that the algorithm (and software) can handle more general situations like rectangular domains, interior boundary cells, etc., and could be generalized to arbitrary domain shapes and partitionings.

First we describe the initialization phase of the preconditioner, which is the necessary preprocessing that has to be done only once for a series of linear systems with matrices with the same sparsity pattern. Then we describe the factorization phase, which has to be done separately for every matrix. Finally we describe the solution phase, which has to be performed for every application of the preconditioner.

2.1. Initialization phase

First the system is partitioned into N non-overlapping subdomains Ω_i , with $i = 1, \dots, N$. These subdomains may be distributed over different processes, which allows for parallelization of the computation. The partitioning may be done based on the graph of the matrix, as implemented for instance in METIS [25], or by a geometric approach, e.g. by dividing the domain into rectangular subdomains. An example of a Cartesian partitioning of a square domain is shown in Figure 2. The non-overlapping subdomains are denoted by the black lines.

In the discretization, variables in a subdomain may be coupled to variables in other subdomains. This means that the linear systems associated with the

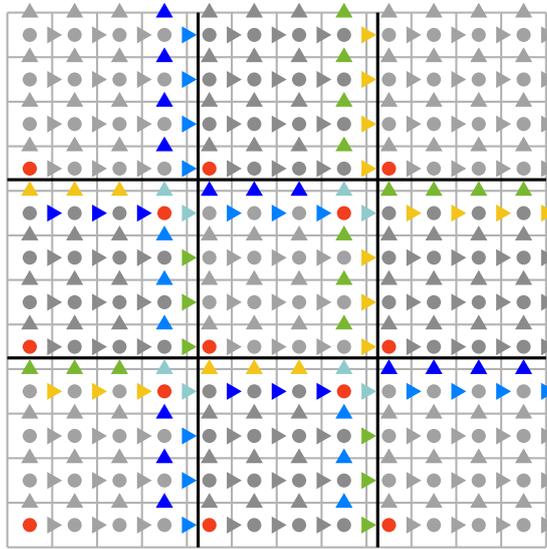


Figure 2: Cartesian partitioning of a 12×12 C-grid discretization of the Navier–Stokes equations into 9 subdomains of size $s_x \times s_y$ with $s_x = s_y = 4$. The interiors are shown in gray. Velocities of the same (non-gray) color that are on the same face of neighboring cells form a separator group. The red circles are pressure nodes that are retained.

subdomains may not be solved independently. For this reason we define overlapping subdomains $\bar{\Omega}_i$, which have *interior nodes* that are in the non-overlapping subdomain Ω_i and are coupled only to variables in the overlapping subdomain $\bar{\Omega}_i$. The interiors of all subdomains may be solved independently. The nodes in the overlapping subdomains that are not only coupled to variables in the same overlapping subdomain constitute the *separator nodes*. One *separator* is defined as a set of separator nodes that are coupled with the same set of subdomains. One *separator group* is defined as the variables on the same separator that have the same variable type (u, v, w or p). In Figure 2 the interior nodes are shown in gray and the different separator groups are denoted by different colors.

We can now reorder the matrix A such that the interiors (I) and separators (S) are grouped. This gives us the system

$$\begin{pmatrix} A_{II} & A_{IS} \\ A_{SI} & A_{SS} \end{pmatrix} \begin{pmatrix} \mathbf{x}_I \\ \mathbf{x}_S \end{pmatrix} = \begin{pmatrix} \mathbf{b}_I \\ \mathbf{b}_S \end{pmatrix},$$

where A_{II} consists of independent diagonal blocks. Since A_{II} consists of independent diagonal blocks, applying A_{II}^{-1} is easy and trivial to parallelize. For this reason, we aim to solve

$$\begin{aligned} S\mathbf{x}_S &= \mathbf{b}_S - A_{SI}A_{II}^{-1}\mathbf{b}_I, \\ \mathbf{x}_I &= A_{II}^{-1}\mathbf{b}_I - A_{II}^{-1}A_{IS}\mathbf{x}_S, \end{aligned}$$

where S is the Schur complement given by $S = A_{SS} - A_{SI}A_{II}^{-1}A_{IS}$.

Whether a variable is coupled to a different subdomain can be determined from the graph of the matrix. It may again also be determined geometrically by additionally defining the overlapping subdomains during the partitioning

step, and checking what overlapping subdomains a node of the non-overlapping subdomain belongs to. The separators can be determined by, for every node, making a set of subdomains it is coupled to or overlapping subdomains it belongs to, and then grouping the nodes that have the same set.

There are three types of separators: faces (in 3D), edges and corners. For the Navier–Stokes problem on a C-grid, these separators only consist of V -nodes. The P -nodes are only connected to V -nodes that belong to the same overlapping subdomain, so these should never lie on a separator. We arbitrarily choose one P -node in every interior which we also define to be its own separator group to make sure the Schur complement stays nonsingular.

We want to eliminate velocities on a separator in such a way that the velocities that remain on a separator face provide an approximation of the collective flux through that face. It is therefore important that the variables are correctly scaled before the factorization in a way that they represent physical fluxes. In the matrix in (2) this gives a G -part with entries of constant magnitude. In this case we can define a Householder transformation which exactly decouples all but one V -node from the P -nodes [1]. This transformation is of the form

$$H_j = I - 2 \frac{\mathbf{v}_j \mathbf{v}_j^T}{\mathbf{v}_j^T \mathbf{v}_j}, \quad (3)$$

for some separator group g_j with

$$\mathbf{v}_j^{(k)} = \begin{cases} \mathbf{e}_j^{(k)} + \|\mathbf{e}_j\|_2 & \text{if node } k \text{ is the first node of separator group } g_j \\ \mathbf{e}_j^{(k)} & \text{otherwise} \end{cases} \quad (4)$$

and

$$\mathbf{e}_j^{(k)} = \begin{cases} 1 & \text{if node } k \text{ is in separator group } g_j \\ 0 & \text{if node } k \text{ is not in separator group } g_j \end{cases}$$

for all $k = 1, \dots, n$. We call the one V -node that is still coupled to the P -nodes a V_Σ node. Since the Householder transformation can be applied independently for every separator, we may collect all these transformations in one single transformation matrix H .

2.2. Factorization phase

For every overlapping subdomain $\bar{\Omega}_i$, $i = 1, \dots, N$, where N is the total number of overlapping subdomains, we can define a local matrix

$$A^{(i)} = \begin{pmatrix} A_{II}^{(i)} & A_{IS}^{(i)} \\ A_{SI}^{(i)} & A_{SS}^{(i)} \end{pmatrix}.$$

After computing the factorization $A_{II}^{(i)} = L_{II}^{(i)} U_{II}^{(i)}$, the local contribution to the Schur complement is given by

$$S_i = A_{SS}^{(i)} - A_{SI}^{(i)} (L_{II}^{(i)} U_{II}^{(i)})^{-1} A_{IS}^{(i)},$$

and the global Schur complement is given by

$$S = \sum_{\Omega_i} A_{SS}^{(i)} - \sum_{\bar{\Omega}_i} A_{SI}^{(i)} (L_{II}^{(i)} U_{II}^{(i)})^{-1} A_{IS}^{(i)}.$$

Here we take the sum of the $A_{SS}^{(i)}$ contributions over the non-overlapping subdomains to make sure contributions from separators are not counted multiple times.

We now apply the Householder transformation

$$\begin{aligned} S_T &= HSH^T = H \left(\sum_{\Omega_i} A_{SS}^{(i)} - \sum_{\bar{\Omega}_i} A_{SI}^{(i)} (L_{II}^{(i)} U_{II}^{(i)})^{-1} A_{IS}^{(i)} \right) H^T \\ &= \sum_{\Omega_i} H_i A_{SS}^{(i)} H_i^T - \sum_{\bar{\Omega}_i} H_i A_{SI}^{(i)} (L_{II}^{(i)} U_{II}^{(i)})^{-1} A_{IS}^{(i)} H_i^T, \end{aligned}$$

which can be done separately for every separator or subdomain, or on the entire Schur complement. We choose to do this separately for every subdomain, since H is very sparse, and sparse matrix-matrix products are very expensive and memory inefficient.

We now drop all connections between V_Σ and non- V_Σ nodes and between non- V_Σ nodes and non- V_Σ nodes in different separator groups. The dropping that is applied here is what makes our factorization inexact. Not applying the dropping gives us a factorization that can still be partially parallelized, but is also exact.

Our transformed Schur complement is now reduced to a block-diagonal matrix with blocks for the non- V_Σ nodes for every separator and one block for all V_Σ nodes, which we will call $S_{\Sigma\Sigma}$. Separate factorizations can again be made for all these diagonal blocks, which can again be done in parallel. For the non- V_Σ blocks, sequential LAPACK solves can be used, and for $S_{\Sigma\Sigma}$ we can employ a (distributed) sparse direct solver. We denote the factorization that is computed by these solvers by $L_S U_S$.

The full factorization obtained by the method is given by

$$A = \begin{pmatrix} L_{II} & 0 \\ A_{SI} U_{II}^{-1} & H^T L_S \end{pmatrix} \begin{pmatrix} U_{II} & L_{II}^{-1} A_{IS} \\ 0 & U_S H \end{pmatrix}.$$

2.3. Solution phase

After the preconditioner has been computed, it can be applied to a vector in each step of a preconditioned Krylov subspace method, for which we use the well-known GMRES method [26]. Communication has to occur in the application of A_{IS} and A_{SI} , and when solving the distributed V_Σ system. The latter was addressed by using a parallel sparse direct solver in [22], but in the next section we propose a different road that does not rely on the availability of such a solver.

3. The multilevel ILU preconditioner

The main issue with the above two-level ILU factorization that prevents scaling to very large problem sizes in three space dimensions is that as the problem size increases and the subdomain size remains constant, the size of $S_{\Sigma\Sigma}$ will increase steadily and any (serial or parallel) sparse direct solver will eventually limit the feasible problem sizes. Increasing the subdomain size, on the other hand, leads to more iterations and therefore more synchronization points.

One of the strong points, on the other hand, is the fact that it preserves the structure of the original problem in the sense that, when applied to an \mathcal{F} -matrix, it produces a strongly reduced matrix $S_{\Sigma\Sigma}$ which is again an \mathcal{F} -matrix. It is therefore intuitive to try applying the scheme recursively to avoid the problem of the growing sparse matrix that has to be factorized. From the structure preserving properties of the algorithm, it is immediately clear that such a recursive scheme will again lead to grid-independent convergence if the number of recursions is kept constant as the grid size is increased.

From now on we will refer to the number of recursions, or the number of times a reduced matrix $S_{\Sigma\Sigma}$ is computed, as the number of levels. Note that this means that the method, which was previously referred to the two-level method is in fact the first level of the multilevel method. Applying a direct solver to S_T from (??) would be level zero, since in this case the preconditioner itself is in fact just a direct solver.

In order to apply the method to the reduced matrix $S_{\Sigma\Sigma}$, we need a coarser partitioning for which it is optimal in terms of communication to make sure subdomains are not split up in the new partitioning. In case we have a regular partitioning like a rectangular partitioning this may be done by multiplying the *separator length* by a certain *coarsening factor*. Having a coarsening factor of 2 would mean that in 3D the separator length is increased by a factor 2, and the number of subdomains is reduced by a factor 8.

As stated in the previous section, we require the velocity variables to be correctly scaled to be able to apply the Householder transformation. However, the V_Σ -variables from the previous level that lie on one separator had a different number of variables in their separator groups. In case of a regular partitioning, an edge separator, for instance, consists of V_Σ -nodes from two edges and one corner from the previous level. This leads to a different scaling of the V_Σ -nodes and thus to non-constant entries in the G -part of $S_{\Sigma\Sigma}$. Instead of re-scaling the $S_{\Sigma\Sigma}$ matrix on every level, we instead use a test vector \mathbf{t} . The test vector is defined initially as a constant vector of ones, and is multiplied by the Householder transformation H at each consecutive level. The Householder transformation is as defined in (3) and (4), but now with

$$\mathbf{e}_j^{(k)} = \begin{cases} \mathbf{t}^{(k)} & \text{if node } k \text{ is in separator group } g_j \\ 0 & \text{if node } k \text{ is not in separator group } g_j \end{cases}$$

for all $k = 1, \dots, n$. After applying the Householder transformation, we can again apply dropping to remove connections between V_Σ and non- V_Σ nodes and between non- V_Σ nodes and non- V_Σ nodes in different separator groups. When the matrix $S_{\Sigma\Sigma}$ is sufficiently small, a direct solver is applied to factorize it.

Instead of just having one separator group per variable per separator, we may also choose to have multiple separator groups, meaning that instead of retaining only one V_Σ node per variable per separator we retain multiple V_Σ nodes. This in turn means that less dropping is applied, and therefore the factorization is more accurate. Retaining all nodes in this way, possibly only after reaching a certain level, gives us an exact factorization, which, in terms of iterations for the outer iterative solver, should give the same results as using any other direct solver at that level. A visual representation of this process is given in Figure 3.

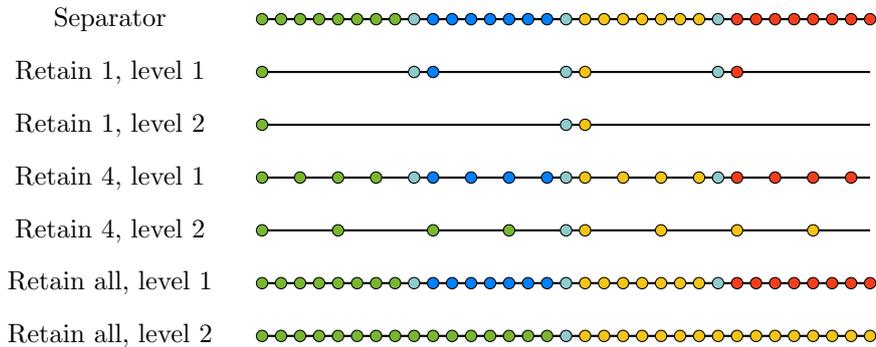


Figure 3: One-dimensional representation of the process of retaining multiple nodes per separator.

4. Skew partitioning in 2D and 3D

Looking at Figure 2, we see that there are pressures that are located at the corners of the subdomains that are surrounded by velocity separators. This means that if we add these pressures to the interior, as was suggested above, we get a singular interior block. We call these pressure nodes that are surrounded by velocity separators *isolated pressure nodes*. For the two-level preconditioner in 2D, it was possible to solve this problem by turning these pressures into their own separator groups. This unfortunately does not work for the multilevel method, since in this case velocity nodes around the isolated pressure nodes get eliminated. It also does not work in 3D because there the isolated pressure nodes also exist inside of the edge separators, where they form tubes of isolated nodes.

A possible way to solve this for the multilevel case and in 3D is to also turn all velocity nodes around the isolated pressure nodes into separate separator groups. This means that they will all be treated as V_Σ nodes and will never be eliminated until they are in the interior of the domain at a later level. This, however, has a downside that a lot more nodes have to be retained at every level, resulting in much larger $S_{\Sigma\Sigma}$ matrices at every level. Another downside is that a lot of bookkeeping is required to properly keep track of all the extra separator groups.

In 2D, we can resolve these problems by rotating the Cartesian partitioning by 45 degrees. The result is shown in Figure 4a. It is easy to see that in this case no pressure nodes are surrounded by only velocity separators. We call this partitioning method *skew partitioning*. In Figure 4, we also show the workings of the multilevel method, with all the steps being displayed in the different subfigures.

For the skew partitioning to work with our multilevel method, we have two requirements on the shapes of the subdomains: (1) it should be *space-filling*, meaning that it can be used to fill the entire domain and (2) it should be *scalable*, meaning that we can construct a larger subdomain of the same shape from multiple smaller subdomains. It is easy to see that these two properties hold for the 2D skew partitioning.

The most basic idea for generalizing the rotated square shape to a 3D setting was to use octahedral subdomains. Partitioning with this design turned out to

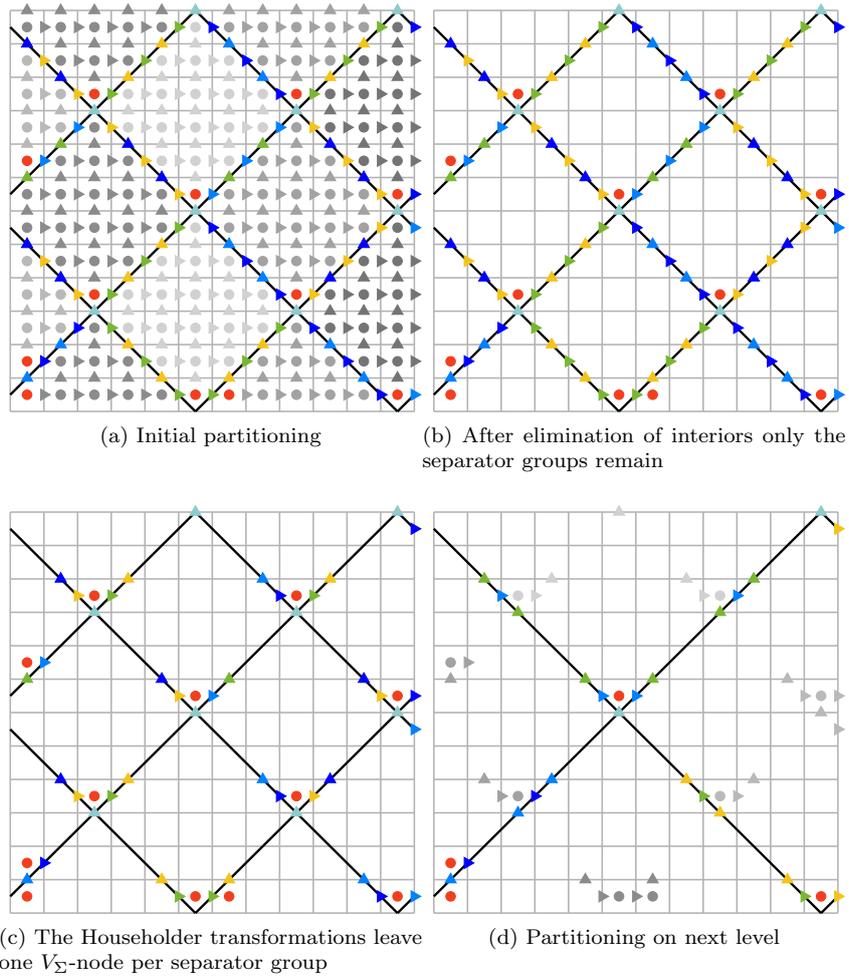


Figure 4: Skew partitioning of a 12×12 C-grid discretization of the Navier–Stokes equations into 12 subdomains. The interiors are shown in gray. Velocities of the same (non-gray) color that are on the same face of neighboring cells form a separator group. The red circles are pressure nodes that are retained. This example uses a coarsening factor of 2, i.e. the separators on the next level have twice the length of those on the previous level.

be unsuccessful, but it is still briefly discussed here since it led to some new insights. Since regular octahedra (the dual to cubes, having their vertices at the centers of the cube faces) by themselves are not space-filling, the octahedra can be slightly distorted to make them fit within a single cubic repeat unit. The resulting subdomains are space-filling, but only by using three mutually orthogonal subdomain types. The fact that it requires the use of three types of templates increases the programming efforts significantly since it introduces a lot of exceptional cases that should be considered, e.g. for subdomains located at the boundary of the domain.

A problem with the octahedral subdomains is that they are not scalable, meaning that we can not construct a larger octahedral subdomain from multiple smaller octahedral subdomains. This is of course required for the multilevel method. However, scalability can be achieved by splitting the octahedra into four smaller tetrahedra, of which six different types are required to fill 3D space. This would introduce additional separator planes that are similar to the 2D skew case and hence it increases the risk of isolating a pressure node when such planes intersect. Especially planar intersections which are parallel to either of the Cartesian axes have a high risk of producing isolating pressure nodes.

We did indeed not manage to find any scalable tiling using tetrahedra that would not give rise to any isolated pressure nodes. Moreover, we would like to have a single subdomain shape that we can use instead of six, since this would greatly simplify the implementation. A lesson we learned is that isolated pressure nodes always seem to arise when having faces that are aligned with the grid. Therefore, we looked into a rotated parallelepiped shape that does not have any faces that are aligned with the grid [27]. This shape is shown in Figure 5a, where the cubes represent a set of $s_x \times s_x \times s_x$ grid cells. A welcome property of this domain is that its central cross section resembles the proposed skew 2D partitioning method.

A schematic view of the position of the separator nodes is shown in Figure 5b. One important detail to note is that on the side that is facing towards us, only half of the w -nodes are displayed. This is because the w -nodes have to be positioned in an alternating fashion on the inside and outside of the non-overlapping subdomain to prevent isolated pressure nodes from appearing. A consequence of this alternating property is that the w -planes have to be divided into two separate separator groups; one for the w -nodes that are located inside the non-overlapping subdomain, and one for w -nodes that are only present in the overlapping subdomain.

Another advantage of using a skew domain partitioning is that the amount of communication that is required is reduced compared to a square partitioning. In [28], it is estimated that for the Laplace problem, the communication is asymptotically reduced by a factor of $\sqrt{2}$ for the 2D diamond shape. If we instead compare the diamond shape to a rectangular domain with the same number of nodes (having the same number of nodes with a square domain is impossible), we find that communication is reduced by a factor of $\frac{3}{2}$. In the same manner, we can compare a 3D domain of size $s_x \times s_x \times s_x/2$ to the rotated parallelepiped, and find a factor of $\frac{4}{3}$. We remark that the truncated octahedron that is used in [28] for the 3D domain and has a factor of 1.68 can not be used for our multilevel method, since truncated octahedra are not scalable.

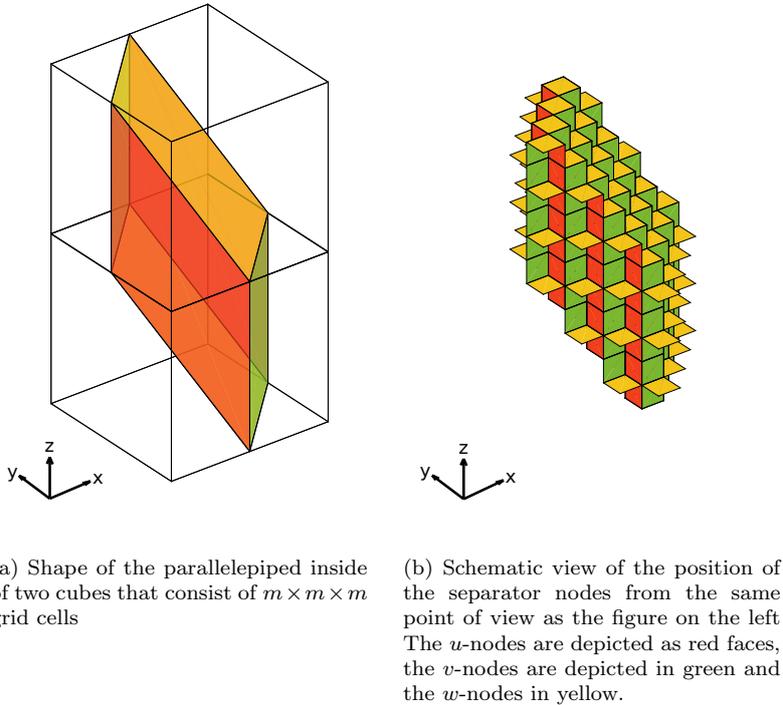


Figure 5: Parallelepiped shape for partitioning the domain.

5. Numerical results

A common benchmark problem in fluid dynamics is the lid-driven cavity problem. We consider an incompressible Newtonian fluid in a square three-dimensional domain of unit length, with a lid at the top which moves at a constant speed U . The equations are given by (1). No-slip boundary conditions are applied at the walls, meaning that they are Dirichlet boundary conditions, and the equations are discretized on an Arakawa C-grid as described before. We take $n_x = n_y = n_z$ grid cells in every direction.

At first, however, we will only look at the Stokes equations of the form

$$\begin{aligned}\Delta \mathbf{u} - \nabla p &= 0, \\ \nabla \cdot \mathbf{u} &= 0.\end{aligned}$$

This is because our preconditioner is constructed in such a way that memory usage and time cost for both computation and application of the preconditioner should not be influenced by inclusion of the convective term. After this, we will further investigate the behavior of the method on the lid-driven cavity problem for increasing Reynolds numbers, which constitute harder problems. Therefore we expect an increase in iterations of the iterative solver, but otherwise the same behavior.

For obtaining the exact memory usage, we developed a custom library which overrides all memory allocation routines when linking against it. The library

contains a hash table in which the amount of memory that is allocated is stored by its memory address. We keep track of the total amount of memory that is allocated, which is increased on memory allocation, and reduced by the amount that is stored in the hash table when memory is freed. The reason we did this is that existing methods rely on rough estimates of the memory usage of the used data structures, use the data that is available from `/proc/meminfo` which is inaccurate, or actually count memory usage in a similar way as we do (i.e. `valgrind`), but have a large performance impact.

We perform every experiment twice: once to determine the memory usage, and once to determine the run time, without linking to the memory usage library. This means that when reporting timing results, we are not impacted by the performance impact of tools to determine memory usage. The reason that we are still concerned about their performance impact, even though we developed our own library, is that it adds roughly a constant amount of time per process, which impacts scalability results. The memory usage that we report is the exact difference in memory usage before and after a certain action is performed, e.g. before and after the construction of the preconditioner.

For the timing results, we do not include the time it takes to compute the partitioning, because we did not optimize this step. The partitioning is computed by first constructing a template sequentially, which contains all possible nodes of exactly one overlapping subdomain, which may even partially be outside of the domain. This template is then mapped to the correct position for every overlapping subdomain to determine the interiors and separator groups. However, since the template contains all possible nodes, every time we increase the number of levels by 1, the amount of time to compute this template increases by a factor 8, which is the worst possible scenario. The reason we do not include this in the timing results is that this may be resolved by only computing the partitioning for nodes that are still present in the Schur complement at that level. This would, however, require a full rewrite of the existing partitioning code, while it would not have any impact on the timing results of the remainder of the implementation, since this is completely decoupled. This means that even though the partitioning method does not scale at all, the preconditioning method itself can be studied reliably without including the timing of the partitioning.

For the implementation of the preconditioner and solver, we use libraries from the Trilinos project [29]. The libraries we use are Epetra (vector and matrix data structures), IFPACK (direct solver and preconditioning interfaces) [30], Amesos (direct solvers) [31] and Belos (iterative solvers) [32]. As iterative solver we use GMRES(250) [26], as parallel sparse direct solver on the coarsest level we use SuperLU_DIST 6.1 [33], and as direct solver for the interior blocks we use KLU [34] with the fill-reducing ordering from [24]. The implementation of our preconditioner can be found on GitHub ¹.

The benchmark is performed on the SuperMUC-NG cluster at LRZ, Munich ². The nodes contain two Intel Xeon “Skylake” Platinum 8174 processors with 24 cores and have 96 GB of memory. For all experiments, we disabled multithreading inside the MPI processes, since the implementation in Epetra is

¹<https://github.com/nlesc-smcm/hymls>

²<https://www.top500.org/system/179566/>

not compatible with our need for very small objects (e.g. sparse matrices living on a very small subdomain).

5.1. Weak scalability

First, we look at results obtained when increasing the grid size n_x at the same rate as the number of used cores n_p , i.e. the problem size on each core is kept constant. The exact configurations that we use are 1 core for $n_x = 16$, 1 core for $n_x = 32$, 8 cores for $n_x = 64$, 64 cores for $n_x = 128$, 512 cores for $n_x = 256$ and 4096 cores for $n_x = 512$. The size of the subdomains (the size of the cubes in Figure 5a) at the first level is $s_x = 8$, while we choose the coarsening factor to be 2, meaning we increase s_x by a factor of 2 at each level. We perform experiments when keeping the number of levels constant at $L = 2$, and when increasing the number of levels by 1 when doubling the grid size. For the latter, we look at three cases where we retain a different number of separator nodes starting at level 2: 1, 4, and all. Since we start retaining more nodes after two levels, results with only two levels ($n_x = 16$) will be the same for all configurations. When a result is missing, this means the method ran out of memory.

For the fixed number of levels, we expect the number of iterations of the iterative solver to converge to a constant number as the grid is refined. For the case where we increase the number of levels as the domain size increases, we expect the number of iterations to only increase mildly, and we expect that retaining more separator nodes starting at level 2 decreases the number of iterations until it again becomes constant as we retain more and more nodes per separator.

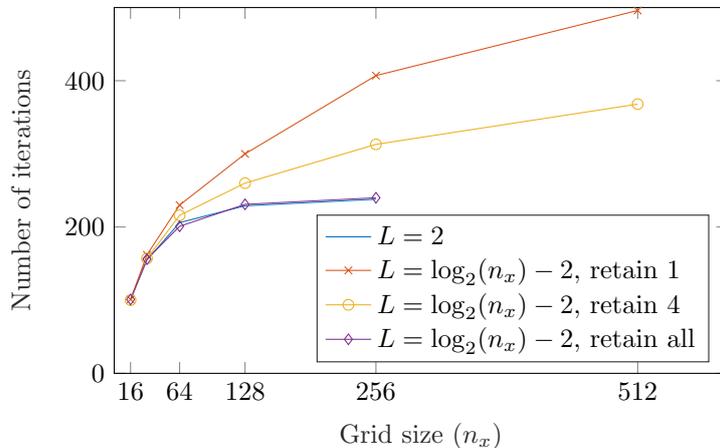


Figure 6: Number of iterations of GMRES for the Stokes problem on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled. A relative residual of 10^{-8} was used as convergence tolerance.

The results are shown in Figure 6. We see that indeed the number of iterations becomes constant when fixing the number of levels or when retaining all separator nodes. When increasing the number of levels with the grid size, we see that the number of iterations keeps increasing gradually. Interesting is that when retaining 4 instead of 1 separator node per separator group after level 2,

the number of iterations that is required decreases drastically, even though this does not have a significant impact on the memory usage as we will see later.

The computational time of both computing the preconditioner, as well as the application of the preconditioner would ideally become constant when keeping the problem size at each core constant while increasing the problem size. However, in practice this is not possible since increasing the number of cores also increases the required amount of communication. The results are shown in Figure 7 and Figure 8.

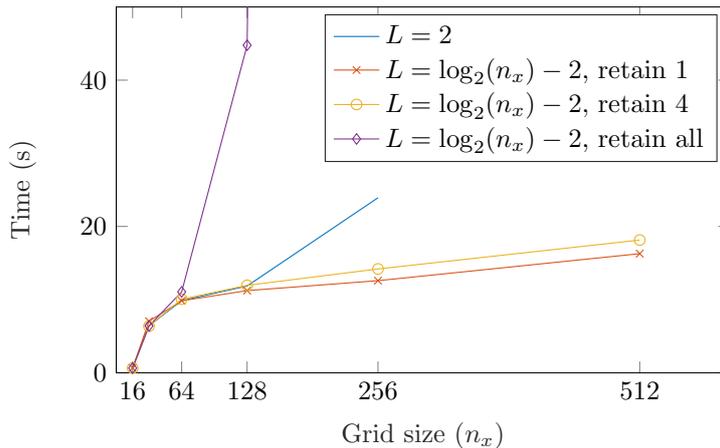


Figure 7: Time to compute the preconditioner for the Stokes problem on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled.

When computing the preconditioner, we see that if we increase the number of levels with the grid size, the computational time only slightly increases. Ideally, there would be no increase at all, but since increasing the number of computational nodes also means increased communication, there will always be an increase in practice. In our case the communication mainly happens at the point where contributions of neighboring subdomains have to be added up in the Schur complement. Since retaining more nodes per level means an increase in the amount of communication, we also see that retaining 4 or all nodes takes more time than retaining only 1 node per separator at every level. It must be noted that due to system variability, we also found results that were about 10% better or worse than the results displayed here. This could be improved by disabling frequency adaptation in the CPUs, but we are aware that at very large core counts the synchronous nature of our algorithm may become a scalability problem.

If we keep the number of levels constant, the computational time required to compute the factorization at the last level increases rapidly, and for larger grid sizes the amount of memory that is required for the factorization is too large for the computation to complete. We also notice that retaining all nodes after level 2 is much more inefficient than just using SuperLU_DIST at level 2, meaning that our preconditioner performs very poorly as a direct solver. This is mainly due to the fact that the Schur complement at the last level is quite large and full. The last level Schur complement for a grid of size 256^3 has size 20961×20961

and its factorization is filled with 72% nonzeros. Computing the factorization of this matrix and applying it is therefore really expensive. Using a parallel dense solver instead of SuperLU_DIST might help to make it more efficient. Moreover, since we chose $n \sim n_p$, the cost of computing the factorization in a sparse direct solver goes at best with $\mathcal{O}(n^2/n_p) = \mathcal{O}(n) = \mathcal{O}(n_x^3)$ [33], and expect that the cost of computing the preconditioner when keeping the number of levels constant, or when retaining all separator nodes, increases with n_x^3 .

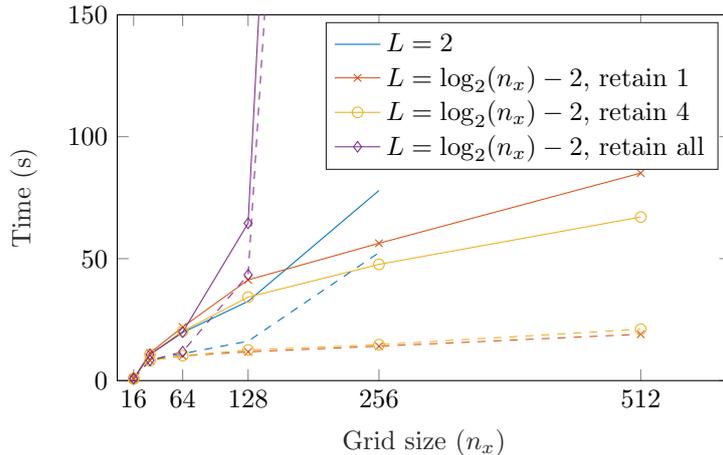


Figure 8: Time to solve the linear system with GMRES (lines), and time of 200 applications of the preconditioner (dashed lines). This is for the Stokes problem on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled. A relative residual of 10^{-8} was used as convergence tolerance.

In Figure 8, we show both the time it takes to solve the linear system after computation of the preconditioner, and the time of 200 applications of the preconditioner, which is not influenced by the number of iterations of the iterative solver and does not include time spent on for instance matrix-vector products and orthogonalization. First of all, we again observe that retaining all separator nodes is a bad idea, since the computational time goes off the chart. For the case where we only use 2 levels, we also see the unwanted behavior of a time that keeps increasing linearly for the total solution time, and superlinearly for the application of the preconditioner, where we would actually expect $\mathcal{O}(n^{4/3}/n_p) = \mathcal{O}(n_x)$ behavior from the triangular solve [33]. This may be caused by disabling multithreading support, which results in an increased amount of communication inside of SuperLU_DIST.

For both cases where we retain 1 and 4 separator nodes after 2 levels, we see that the computational time only slightly increases for larger grid sizes. We also note that even though the case where we retain 4 nodes is slower per application, it is faster in total solution time due to the lower number of iterations that is required, as can also be seen in Figure 6, and the fact that applying the preconditioner is so cheap.

In Figure 9, we see the average memory usage of the preconditioner per core. Here we again observe that the 2 level case, and the case where we retain all separator nodes after level 3 performs poorly. The case where we retain 1 or 4

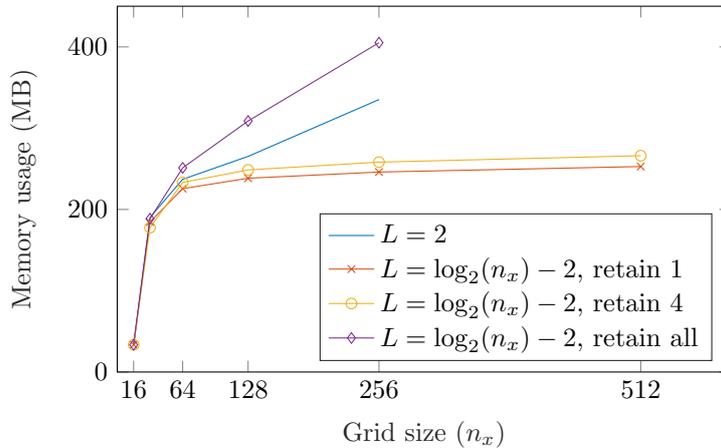


Figure 9: Memory usage of the preconditioner per core for the Stokes problem on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled.

separator nodes per separator group after level 2 show similar behavior in terms of memory usage, and becomes constant as expected.

5.2. Strong scalability

In this section, we will look at a problem with size $n_x = 128$, with 6 levels and retaining only one node per separator group. We use 1 to 128 cores with steps of a factor of 2 for all cases except the application time, where we use 2 to 128 cores. The reason we do not look at 1 core for this timing is that this configuration caused a memory allocation error in the iterative solver (Belos).

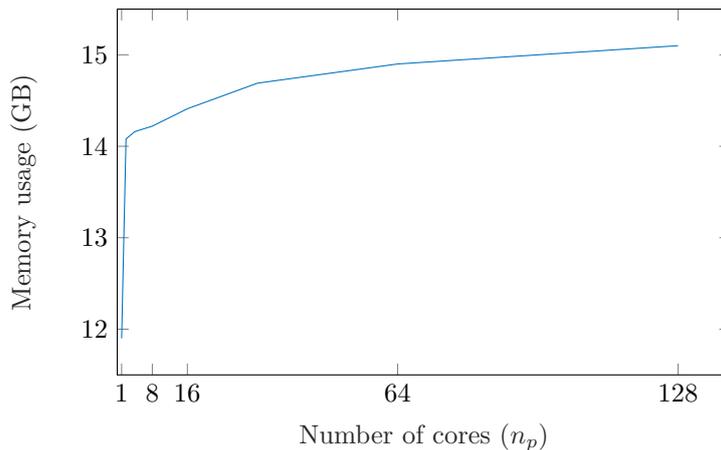


Figure 10: Total memory usage of the preconditioner for the Stokes problem on a grid of size 128^3 , with 1 to 128 cores.

We first look at the total memory usage in Figure 10. We observe that there is a large difference between the memory usage on one core, and the memory

usage on two cores. This is due to the fact that mainly Epetra uses different implementations of its data structures for serial and parallel computations.

We also observe that the total memory usage increases slightly when using more cores. This can be explained by the overlap that exists between different processes. Since the difference in communication between a cubical domain and a parallelepiped shaped domain is a constant factor of $\frac{3}{4}$, we may instead look at a cubical domain to explain this behavior. Say we have a subdomain of size $s_x \times s_y \times s_z$, then communication is required for $2s_x s_y + 2s_x s_z + 2s_y s_z$ grid cells. If we now split this subdomain in the x -direction, then we require communication for $2s_x s_y + 2s_x s_z + 4s_y s_z$ grid cells. If we now assume that $s_x = s_y = s_z$, we see that communication increases with a factor $\frac{4}{3}$ when doubling the number of cores. We verified that this indeed explains the slight increase in memory usage.

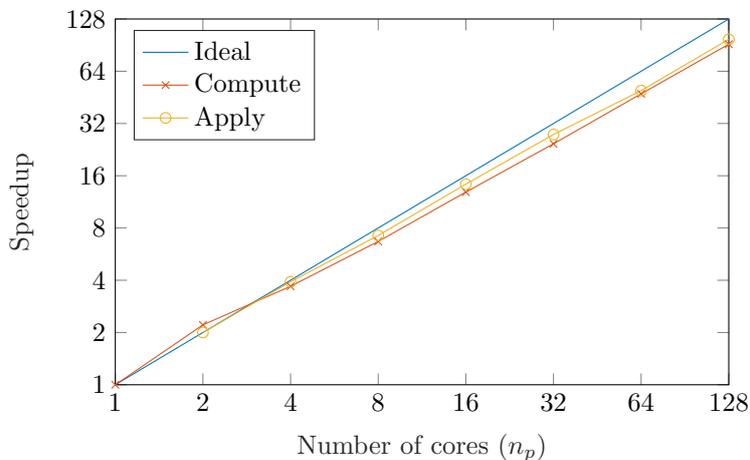


Figure 11: Speedup of computation and application of the preconditioner for the Stokes problem on a grid of size 128^3 , with 2 to 128 cores.

In Figure 11 we look at the speedup when using more cores. Ideally, we would see that using twice the number of cores means half of the computational time. We plotted this ideal line for reference. We see that the speedup of both the computation and application of the preconditioner is very close to this ideal line. There is one outlier at 2 cores which is above the ideal line, which may be explained by system variability. Other than that the behavior is as expected.

5.3. Lid-driven cavity

In the previous sections we determined the weak and strong scalability properties of the preconditioner, which means that we can now continue with the robustness of the solver on the lid-driven cavity problem with increasing Reynolds numbers. We perform a continuation with steps of $Re = 100$ starting from the solution of the Stokes problem. We show the results of the first iteration of Newton at $Re = 500$ and $Re = 2000$. Reason we go up to $Re = 2000$ is that a Hopf bifurcation is located between $Re = 1900$ and $Re = 2000$, which is of interest [18, 35].

In Figure 12 and Figure 13, we show the results at Reynolds number 500, which show good correspondence with the results on the Stokes problem. The

main difference is that more iterations are needed, since higher Reynolds numbers constitute harder problems. We do see that the number of iterations that is required when keeping the number of levels constant actually decreases, which is due to the grid refinement having a positive effect on the mesh Péclet number, i.e. the coefficients of the diffusion increase with respect to those of the convection.

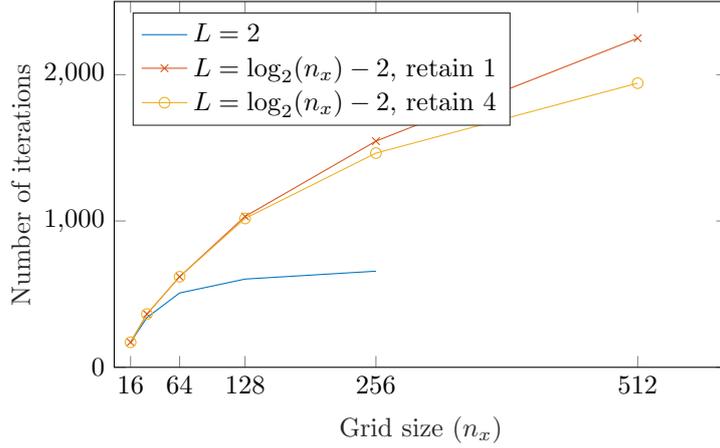


Figure 12: Number of iterations of GMRES for the lid-driven cavity problem at $\text{Re} = 500$ on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled. A relative residual of 10^{-8} was used as convergence tolerance. As initial guess we used the solution at $\text{Re} = 400$.

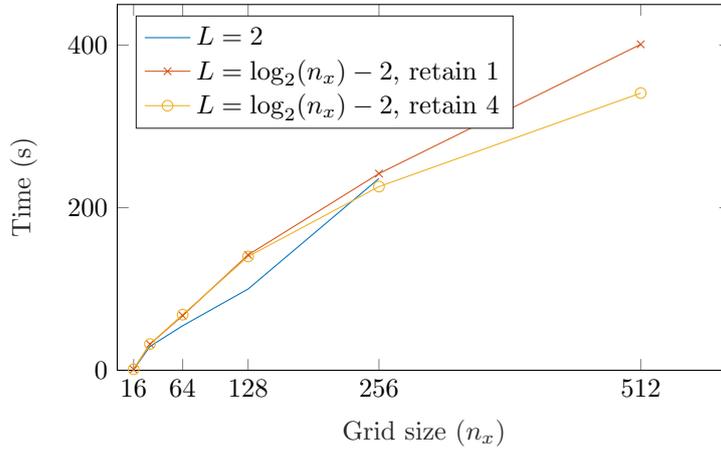


Figure 13: Time for GMRES to converge for the lid-driven cavity problem at $\text{Re} = 500$ on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled. A relative residual of 10^{-8} was used as convergence tolerance. As initial guess we used the solution at $\text{Re} = 400$.

The results at Reynolds number 2000 are shown in Figure 14 and Figure 15. We again observe that the number of iterations is much larger. What is odd,

however, is that retaining more nodes now actually gives worse convergence. This may be because we use GMRES(250) instead of GMRES due to memory limitations, and therefore do not preserve the convergence properties of GMRES. This effect is more prevalent for this problem because of the large number of iterations that is required. We do observe that for the first 250 iterations, retaining 4 nodes after level 2 gives rise to better convergence, as we expected. For Reynolds number 2000, we do not show results with a grid of size 512^3 , because the method did not converge within 10 000 GMRES iterations, which we set as the maximum. It did converge, however, and extrapolating the results, we would expect that around 15000 iterations are needed to meet the tolerance. Alternatively, instead of increasing the number of levels with the grid size, we could also have used 6 levels as we did for a grid of size 256^3 , in which case we would have expected it to actually use less iterations due to the effect of grid refinement on the mesh Péclet number as explained earlier.

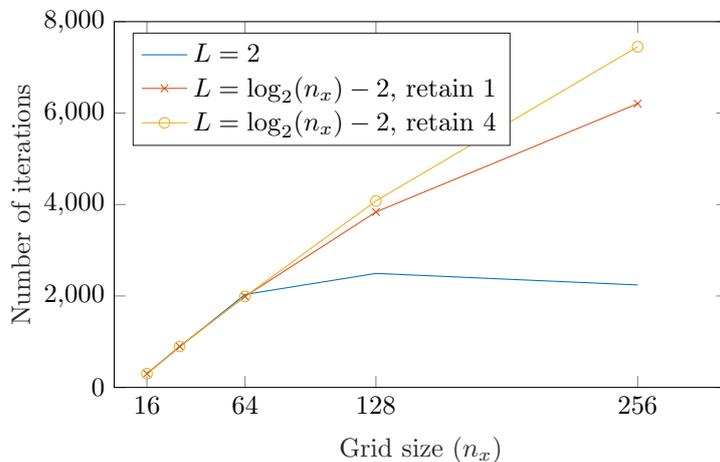


Figure 14: Number of iterations of GMRES for the lid-driven cavity problem at $\text{Re} = 2000$ on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled. A relative residual of 10^{-8} was used as convergence tolerance. As initial guess we used the solution at $\text{Re} = 1900$.

In Table 1 we show results for Reynolds number 500 using the block preconditioner LSC (Least-Squares Commutator) implemented in Teko [12]. We chose Teko for comparison because it is available in Trilinos and can therefore easily be coupled to our code. We also tried other preconditioners, but those did unfortunately not yield convergence. The stopping criterion of the linear solver is 10^{-8} , as before.

Compared to our method, we see that Teko has much more difficulty with the grid refinement, leading to a huge computational cost already at a grid size of 64^3 . A crude computation shows that per grid point the method becomes about 65 times more expensive per iteration. This must be attributed to slow convergence of algebraic multigrid on the subblocks. One of these blocks is the $L + N$ block from (2), with N indefinite, and this is something that is difficult for a standard AMG method. We chose the number of levels to be 2 for grid sizes 16^3 and 32^3 , and 3 levels for 64^3 since these seemed to give the optimal

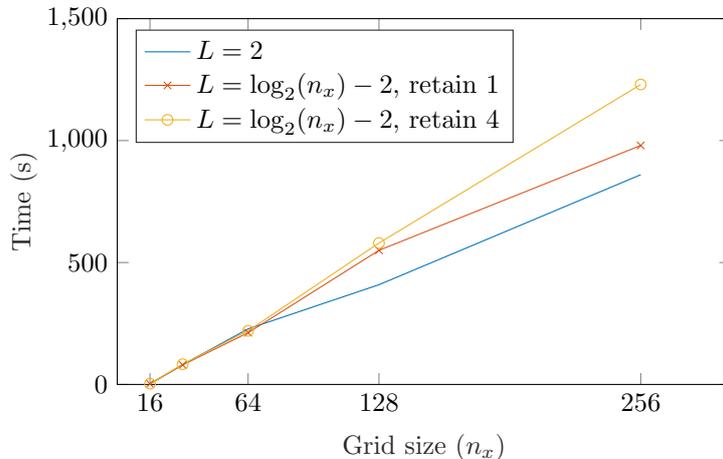


Figure 15: Time for GMRES to converge for the lid-driven cavity problem at $\text{Re} = 2000$ on a grid of size $n_x \times n_x \times n_x$, while keeping the number of levels at $L = 2$, and when increasing the number of levels by 1 when n_x is doubled. A relative residual of 10^{-8} was used as convergence tolerance. As initial guess we used the solution at $\text{Re} = 1900$.

n_p	n_x	L	its	t_c (s)	t_s (s)
1	16	2	142	0.12	0.77
1	32	2	187	9.88	18.80
8	64	3	245	1511.12	313.00

Table 1: Performance of Teko with LSC preconditioner for the lid-driven cavity problem at $\text{Re} = 500$ on a grid of size $n_x \times n_x \times n_x$. Here n_p is the number of cores, L is the number of levels, its is the number of iterations, t_c is the time to compute the preconditioner and t_s is the time for solving the linear system.

results. For Reynolds number 2000, we did not observe convergence past the 16^3 grid.

6. Summary and Discussion

We presented a robust method for solving the steady, incompressible Navier–Stokes equations, which makes use of parallelepiped shaped overlapping subdomains. The interiors of these overlapping subdomains can be eliminated in parallel. On the interfaces of the subdomains, Householder transformations are applied to decouple all but one velocity node from the pressure nodes, after which all decoupled nodes can also be eliminated in parallel. The key to the multilevel approach is the resulting reduced Schur complement matrix, which has the same structure as the original matrix. Therefore we can recursively apply the preconditioner to this matrix. The shape of the subdomains makes sure that pressure nodes are not isolated in the factorization process, which would result in a singular Schur complement matrix.

Our weak scalability experiments show that if the number of levels of the multilevel approach is kept constant while increasing the problem size, grid

independent convergence is obtained. We also show that by increasing the number of levels and processors as the problem size increases, we only require a small amount of additional time and memory for both the factorization and solution process. Moreover, by increasing the number of nodes that is retained per separator after level 2, we can further decrease the time that is required to solve the system.

Our strong scalability experiments show that the time that is required to both compute and apply the preconditioner scales very well. The same holds for the memory usage, which behaves as expected.

We also showed that the preconditioner leads to convergence of GMRES for the lid-driven cavity problem at high Reynolds numbers, where other methods, such as the LSC preconditioner that is implemented in Teko, fail to do so.

This leads us to conclude that we implemented a robust solution method for the Navier–Stokes equations on staggered grids which shows good parallel scalability. In the future we want to apply our preconditioner in ocean-climate models. The fact that it has proven to be robust for the lid-driven cavity problem with higher Reynolds numbers leads us to believe that it will also perform well for large scale ocean models.

Acknowledgments

We would like to thank Eric Cyr for his help in our attempts of getting the solvers in Teko to work as expected. This work is part of the Mathematics of Planet Earth research program with project number 657.014.007, which is financed by the Netherlands Organization for Scientific Research (NWO) (SB and FW). Access to the SuperMUC-NG system during the ‘friendly user phase’ was kindly granted by LRZ Munich via a test account.

- [1] Fred W. Wubs and Jonas Thies. A robust two-level incomplete factorization for (Navier–)Stokes saddle point matrices. *SIAM Journal on Matrix Analysis and Applications*, 32(4):1475–1499, Oct 2011. ISSN 1095-7162. doi: 10.1137/100789439. URL <https://doi.org/10.1137/100789439>.
- [2] Francois Charru and Patricia de Forcrand-Millard. *Hydrodynamic Instabilities*. Cambridge University Press, 2011. ISBN 9780511975172. doi: 10.1017/cbo9780511975172. URL <https://doi.org/10.1017/cbo9780511975172>.
- [3] H. B. Keller. Numerical solution of bifurcation and nonlinear eigenvalue problems. In P. H. Rabinowitz, editor, *Applications of Bifurcation Theory*, pages 359–384. Academic Press, New York, U.S.A., 1977.
- [4] Michele Benzi, Gene H. Golub, and Jörg Liesen. Numerical solution of saddle point problems. *Acta Numerica*, 14:1–137, May 2005. ISSN 1474-0508. doi: 10.1017/s0962492904000212. URL <https://doi.org/10.1017/s0962492904000212>.
- [5] R.W.C.P. Verstappen and A.E.P. Veldman. Symmetry-preserving discretization of turbulent flow. *Journal of Computational Physics*, 187(1): 343–368, May 2003. ISSN 0021-9991. doi: 10.1016/s0021-9991(03)00126-8. URL [https://doi.org/10.1016/s0021-9991\(03\)00126-8](https://doi.org/10.1016/s0021-9991(03)00126-8).

- [6] Howard Elman, David Silvester, and Andy Wathen. *Finite Elements and Fast Iterative Solvers*. Oxford University Press, Jun 2014. ISBN 9780199678792. doi: 10.1093/acprof:oso/9780199678792.001.0001. URL <https://doi.org/10.1093/acprof:oso/9780199678792.001.0001>.
- [7] G.F. Carey and R. Krishnan. Convergence of iterative methods in penalty finite element approximation of the Navier–Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 60(1):1–29, Jan 1987. ISSN 0045-7825. doi: 10.1016/0045-7825(87)90127-7. URL [https://doi.org/10.1016/0045-7825\(87\)90127-7](https://doi.org/10.1016/0045-7825(87)90127-7).
- [8] Henk A. Dijkstra, Fred W. Wubs, Andrew K. Cliffe, Eusebius Doedel, Ioana F. Dragomirescu, Bruno Eckhardt, Alexander Yu. Gelfgat, Andrew L. Hazel, Valerio Lucarini, Andy G. Salinger, Erik T. Phipps, Juan Sanchez-Umbria, Henk Schuttelaars, Laurette S. Tuckerman, and Uwe Thiele. Numerical bifurcation methods and their application to fluid dynamics: Analysis beyond simulation. *Communications in Computational Physics*, 15(01):1–45, Jan 2014. ISSN 1991-7120. doi: 10.4208/cicp.240912.180613a. URL <https://doi.org/10.4208/cicp.240912.180613a>.
- [9] Roel Verstappen and Marc Dröge. A symmetry-preserving cartesian grid method for computing a viscous flow past a circular cylinder. *Comptes Rendus Mécanique*, 333(1):51–57, Jan 2005. ISSN 1631-0721. doi: 10.1016/j.crme.2004.09.021. URL <https://doi.org/10.1016/j.crme.2004.09.021>.
- [10] Arie de Niet, Fred Wubs, Arjen Terwisscha van Scheltinga, and Henk A. Dijkstra. A tailored solver for bifurcation analysis of ocean-climate models. *Journal of Computational Physics*, 227(1):654–679, Nov 2007. ISSN 0021-9991. doi: 10.1016/j.jcp.2007.08.006. URL <https://doi.org/10.1016/j.jcp.2007.08.006>.
- [11] C. M. Klaij and C. Vuik. SIMPLE-type preconditioners for cell-centered, colocated finite volume discretization of incompressible reynolds-averaged Navier–Stokes equations. *International Journal for Numerical Methods in Fluids*, 71(7):830–849, May 2012. ISSN 0271-2091. doi: 10.1002/fld.3686. URL <https://doi.org/10.1002/fld.3686>.
- [12] Eric C. Cyr, John N. Shadid, and Raymond S. Tuminaro. Stabilization and scalable block preconditioning for the Navier–Stokes equations. *Journal of Computational Physics*, 231(2):345–363, Jan 2012. ISSN 0021-9991. doi: 10.1016/j.jcp.2011.09.001. URL <https://doi.org/10.1016/j.jcp.2011.09.001>.
- [13] Xin He, Cornelis Vuik, and Christiaan M. Klaij. Combining the augmented Lagrangian preconditioner with the simple Schur complement approximation. *SIAM Journal on Scientific Computing*, 40(3):A1362–A1385, Jan 2018. ISSN 1095-7197. doi: 10.1137/17m1144775. URL <https://doi.org/10.1137/17m1144775>.
- [14] Jonas Thies, Fred Wubs, and Henk A. Dijkstra. Bifurcation analysis of 3D ocean flows using a parallel fully-implicit ocean model. *Ocean Modelling*,

- 30(4):287–297, Jan 2009. ISSN 1463-5003. doi: 10.1016/j.ocemod.2009.07.005. URL <https://doi.org/10.1016/j.ocemod.2009.07.005>.
- [15] Gerard L. G. Sleijpen and Fred W. Wubs. Exploiting multilevel preconditioning techniques in eigenvalue computations. *SIAM Journal on Scientific Computing*, 25(4):1249–1272, Jan 2004. ISSN 1095-7197. doi: 10.1137/s1064827599361059. URL <https://doi.org/10.1137/s1064827599361059>.
- [16] U Trottenberg, Cornelis W Oosterlee, and Anton Schüller. *Multigrid*. Academic Press, Nov 2000. ISBN 9780127010700.
- [17] A. J. Wathen. Preconditioning. *Acta Numerica*, 24:329–376, Apr 2015. ISSN 1474-0508. doi: 10.1017/s0962492915000021. URL <https://doi.org/10.1017/s0962492915000021>.
- [18] Yuri Feldman and Alexander Yu. Gelfgat. Oscillatory instability of a three-dimensional lid-driven flow in a cube. *Physics of Fluids*, 22(9):093602, Sep 2010. ISSN 1089-7666. doi: 10.1063/1.3487476. URL <https://doi.org/10.1063/1.3487476>.
- [19] E. F. F. Botta and F. W. Wubs. Matrix renumbering ILU: an effective algebraic multilevel ILU preconditioner for sparse matrices. *SIAM Journal on Matrix Analysis and Applications*, 20(4):1007–1026, Jan 1999. ISSN 1095-7162. doi: 10.1137/s0895479897319301. URL <https://doi.org/10.1137/s0895479897319301>.
- [20] Matthias Bollhöfer and Yousef Saad. Multilevel preconditioners constructed from inverse-based ILUs. *SIAM Journal on Scientific Computing*, 27(5):1627–1650, Jan 2006. ISSN 1095-7197. doi: 10.1137/040608374. URL <https://doi.org/10.1137/040608374>.
- [21] José I. Aliaga, Matthias Bollhöfer, Alberto F. Martín, and Enrique S. Quintana-Ortí. Design, tuning and evaluation of parallel multilevel ILU preconditioners. In *High Performance Computing for Computational Science VECPAR 2008*, pages 314–327. Springer Berlin Heidelberg, 2008. doi: 10.1007/978-3-540-92859-1_28. URL https://doi.org/10.1007/978-3-540-92859-1_28.
- [22] Jonas Thies and Fred Wubs. Design of a parallel hybrid direct/iterative solver for CFD problems. *2011 IEEE Seventh International Conference on eScience*, Dec 2011. doi: 10.1109/escience.2011.60. URL <https://doi.org/10.1109/escience.2011.60>.
- [23] Miroslav Tuma. A note on the LDLT decomposition of matrices from saddle-point problems. *SIAM Journal on Matrix Analysis and Applications*, 23(4):903–915, Jan 2002. ISSN 1095-7162. doi: 10.1137/s0895479897321088. URL <https://doi.org/10.1137/s0895479897321088>.
- [24] A. C. de Niet and F. W. Wubs. Numerically stable LDLT-factorization of F-type saddle point matrices. *IMA Journal of Numerical Analysis*, 29(1):208–234, Feb 2008. ISSN 1464-3642. doi: 10.1093/imanum/drn005. URL <https://doi.org/10.1093/imanum/drn005>.

- [25] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, Jan 1998. ISSN 1095-7197. doi: 10.1137/s1064827595287997. URL <https://doi.org/10.1137/s1064827595287997>.
- [26] Youcef Saad and Martin H. Schultz. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, Jul 1986. ISSN 2168-3417. doi: 10.1137/0907058. URL <https://doi.org/10.1137/0907058>.
- [27] Mark van der Kloek. *Skew Partitioning for the Hybrid Multilevel Solver*. Bachelor’s thesis, University of Groningen, Jul 2017.
- [28] Rob H. Bisseling. *Parallel Scientific Computation*. Oxford University Press, Mar 2004. ISBN 9780198529392. doi: 10.1093/acprof:oso/9780198529392.001.0001. URL <https://doi.org/10.1093/acprof:oso/9780198529392.001.0001>.
- [29] Michael A. Heroux, Eric T. Phipps, Andrew G. Salinger, Heidi K. Thornquist, Ray S. Tuminaro, James M. Willenbring, Alan Williams, Kendall S. Stanley, Roscoe A. Bartlett, Vicki E. Howle, Robert J. Hoekstra, Jonathan J. Hu, Tamara G. Kolda, Richard B. Lehoucq, Kevin R. Long, and Roger P. Pawlowski. An overview of the Trilinos project. *ACM Transactions on Mathematical Software*, 31(3):397–423, Sep 2005. ISSN 0098-3500. doi: 10.1145/1089014.1089021. URL <https://doi.org/10.1145/1089014.1089021>.
- [30] Marzio Sala and Michael A. Heroux. Robust algebraic preconditioners using IFPACK 3.0. Technical report, Sandia National Laboratories, Jan 2005. URL <https://doi.org/10.2172/1127118>.
- [31] Marzio Sala, Kendall S. Stanley, and Michael A. Heroux. On the design of interfaces to sparse direct solvers. *ACM Transactions on Mathematical Software*, 34(2):1–22, Mar 2008. ISSN 0098-3500. doi: 10.1145/1326548.1326551. URL <https://doi.org/10.1145/1326548.1326551>.
- [32] Eric Bavier, Mark Hoemmen, Sivasankaran Rajamanickam, and Heidi Thornquist. Amesos2 and Belos: Direct and iterative solvers for large sparse linear systems. *Scientific Programming*, 20(3):241–255, 2012. ISSN 1875-919X. doi: 10.1155/2012/243875. URL <https://doi.org/10.1155/2012/243875>.
- [33] Yang Liu, Mathias Jacquelin, Pieter Ghysels, and Xiaoye S. Li. Highly scalable distributed-memory sparse triangular solution algorithms. *2018 Proceedings of the Seventh SIAM Workshop on Combinatorial Scientific Computing*, pages 87–96, Jan 2018. doi: 10.1137/1.9781611975215.9. URL <https://doi.org/10.1137/1.9781611975215.9>.
- [34] Timothy A. Davis and Ekanathan Palamadai Natarajan. Algorithm 907. *ACM Transactions on Mathematical Software*, 37(3):1–17, Sep 2010. ISSN 0098-3500. doi: 10.1145/1824801.1824814. URL <https://doi.org/10.1145/1824801.1824814>.

- [35] A. Liberzon, Yu. Feldman, and A. Yu. Gelfgat. Experimental observation of the steady-oscillatory transition in a cubic lid-driven cavity. *Physics of Fluids*, 23(8):084106, Aug 2011. ISSN 1089-7666. doi: 10.1063/1.3625412. URL <https://doi.org/10.1063/1.3625412>.