

A parallel linear solver exploiting the physical properties of the underlying mechanical problem

F. J. Lingen · P. G. Bonnier · R. B. J. Brinkgreve ·
M. B. van Gijzen · C. Vuik

Received: 23 December 2013 / Accepted: 24 June 2014 / Published online: 27 August 2014
© Springer International Publishing Switzerland 2014

Abstract The iterative solution of large systems of equations may benefit from parallel processing. However, using a straight-forward domain decomposition in “layered” geomechanical finite element models with significantly different stiffnesses may lead to slow or non-converging solutions. Physics-based domain decomposition is the answer to such problems, as explained in this paper and demonstrated on the basis of a few examples. Together with a two-level preconditioner comprising an additive Schwarz preconditioner that operates on the sub-domain level, an algebraic coarse grid preconditioner that operates on the global level, and additional load balancing measures, the described solver provides an efficient and robust solution of large systems of equations. Although the solver has been developed primarily for geomechanical problems, the ideas are applicable to the solution of other physical problems involving large differences in material properties.

Keywords Parallel · Preconditioner · Domain-decomposition · Deflation · Iterative method · Mechanical

F. J. Lingen
Dynaflow Research Group, Houtsingel 95, 2719 EB Zoetermeer,
The Netherlands

P. G. Bonnier · R. B. J. Brinkgreve
Plaxis bv, P.O. Box 572, 2600 AN Delft, The Netherlands

M. B. van Gijzen · C. Vuik (✉)
Faculty of Electrical Engineering, Mathematics and Computer
Science, Delft Institute of Applied Mathematics, Delft University
of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands
e-mail: c.vuik@tudelft.nl

1 Introduction

Finite element models of 3-D geomechanical problems result in linear systems

$$Ka = f \quad (1)$$

in which K is a sparse matrix that is often symmetric positive definite. In realistic models, the system of equations can involve more degrees of freedom than can be handled at the present time by direct solver implementations on current workstation-class computers. Switching to an iterative solver, such as the Conjugate Gradient method, is one possible way to overcome this problem, provided that the solver is combined with a good preconditioner [12]. The solver will not converge without such a preconditioner because the large variation in material parameters typically results in a very ill-conditioned system. For instance, the stiffness of a loose soil layer can be four to five orders of magnitude smaller than the stiffness of a concrete structure within the same model. This results in a comparable difference in the magnitude of the (diagonal) coefficients of the matrix K .

An effective preconditioner is based on an incomplete Cholesky decomposition of the matrix K combined with an adaptive drop tolerance [11, 14]. Indeed, this preconditioner typically limits the number of solver iterations to 100 and often much less. It has been used successfully in the commercial computer program PLAXIS [1] for solving a wide range of 3-D geomechanical problems. Unfortunately, this preconditioner cannot be implemented in a way that efficiently exploits parallel computers, including multi-core processors [3, 18]. This is a significant drawback as the sequential processing speed more or less stagnates while the number of processing cores increases. Another drawback of the preconditioner is that it operates only at the matrix level;

it does not directly use any information from the underlying physical model.

This paper presents an iterative solver that can be implemented efficiently on multi-core processors and parallel computers, and that exploits the physical properties of the problem to be solved. Although it has been developed primarily for solving problems in geomechanics, it can also be used to solve other types of problems in solid mechanics involving large jumps in material properties. The solver makes use of domain decomposition to obtain parallelism and to separate regions with large differences in material properties. It features a two-level preconditioner comprising an additive Schwarz preconditioner that operates on the sub-domain level, and an algebraic coarse grid preconditioner that operates on the global level [16, 17, 19]. The former is based on the incomplete Cholesky decomposition—similar to the one used in the original PLAXIS solver—of the sub-domain matrices. The latter uses the rigid body modes of the sub-domains to construct the algebraic coarse grid. Its task is to ensure that the convergence rate of the solver does not deteriorate as the number of sub-domains increases.

Previous research [7] also involved a second-level preconditioner for mechanical problems. There are some important differences, however, between that research and the current research:

- The sub-domains in the previous research are formed before the sub-domain matrices are assembled; the sub-domain matrices are obtained by assembling the element matrices associated with each sub-domain. In the current research, the domain is decomposed after the global matrix has been assembled; the sub-domain matrices are obtained by extracting sections from the global matrix.
- The target machines in the previous research are clusters of PCs, whereas the current research targets multi-core shared-memory machines.
- The primary preconditioner used in the previous research is based on the additive Schwarz method. The current research, on the other hand, uses the restricted additive Schwarz method.
- The implementation of the solver developed in the previous research is aimed at a specific class of problems, whereas the implementation in this research is more of a black-box nature.

The current implementation of the solver targets multi-core shared-memory computers because this architecture represents almost all modern (desktop) computers. There is no other reason why this architecture has been selected; the algorithms on which the solver is based can be implemented equally well on parallel computers with a distributed memory architecture.

The performance of the solver has been examined on a standard eight-core workstation. The original solver in PLAXIS has been used to obtain a reference point with which the performance of the new solver has been compared. The number of cores has been varied from one to eight to determine the parallel scalability of the solver.

The remainder of this paper is structured as follows. Section 2 starts with a description of the geomechanical finite element models that are targeted by the solver. This section introduces the terminology that is used throughout the subsequent sections. Section 3 continues with a (mathematical) description of the solver. After that, Section 4 explains how the solver has been implemented on top of a small thread-based message passing library. Next, Section 5 discusses the performance of the solver for various geomechanical problems. Finally, Section 6 presents the conclusions.

2 Description of the finite element models

The geomechanical problems of interest typically involve a (large) volume of soil (or rock) and various structural objects. The soil is often composed of multiple (semi-) horizontal layers with different mechanical properties. The structural objects are embedded in the soil or are located on the top of the soil. They may include steel or concrete walls, foundations, piles, anchors, and tunnels.

Finite element models are used to compute the deformation of the soil and the structural objects. Such a model is composed of volume elements that model the soil and thick structural objects; shell elements that model plates and thin-walled structural objects; line elements that model piles and anchors; and interface elements that model the interaction between the soil and structural objects. The displacement field u in an element e is approximated by

$$u(x) \approx N_e(x)a_e \quad (2)$$

in which N_e is a matrix containing the element shape functions and a_e is a vector containing the *degrees of freedom* in the nodes of the element. The latter are the displacements along the three coordinate axes, and the rotations about those axes if the element is a shell element. Note that adjacent elements share the degrees of freedom in their common nodes.

Substitution of the approximate displacement field into the constitutive equations yields a nonlinear system of equations that is solved with a (quasi) Newton method. Each Newton iteration requires the solution of a linear system of equations of the form

$$Ka = f \quad (3)$$

in which K is the tangent stiffness matrix (or an approximation thereof), a is a vector containing the incremental degrees of freedom in all the nodes, and f is a vector containing the unbalanced forces and moments in the nodes. The solution of this linear system can be approximated with moderate accuracy because it is only an intermediate of the nonlinear system to be solved. This is an advantage when using an iterative solver.

3 Description of the solver

The solver is made up of the following building blocks:

- a physics-based domain decomposition method;
- a standard iterative solver;
- a two-level preconditioner.

Although these building blocks are not new, they have been adapted in various ways to obtain an efficient solution algorithm for geomechanical problems.

3.1 Physics-based domain decomposition

The domain decomposition method can be viewed as a variation of the age-old divide-and-conquer strategy. The idea is to divide the problem to be solved into multiple sub-problems which are more or less solved independently. As the sub-problems are not completely independent, information needs to be exchanged between the sub-problems to obtain the solution of the global problem.

In the context of this paper, the domain decomposition method is applied to the degrees of freedom. The degrees of freedom are first partitioned into n_s non-overlapping groups, each one making up one sub-domain. The sub-domains are then extended by adding degrees of freedom from neighboring sub-domains to obtain n_s slightly overlapping sub-domains. The original degrees of freedom in a sub-domain are called the *internal* degrees of freedom, and the additional degrees of freedom are called the *overlap* degrees of freedom. Note that each degree of freedom is an internal degree of freedom in exactly one sub-domain. Also note that an overlap degree of freedom in one sub-domain is also an internal degree of freedom in exactly one other sub-domain.

The sub-domains are assigned to different threads of execution that are responsible for executing all computations related to the degrees of freedom in their sub-domain. The threads are also responsible for storing all data related to the degrees of freedom in their sub-domain. No data are shared between the threads, except for the data structures that are needed to exchange data between the threads.

By defining *restriction operators* for each sub-domain, the global stiffness matrix K can be expressed

in terms of slightly overlapping sub-domain matrices K_i as follows:

$$K = \sum_{i=1}^{n_s} L_i^T K_i R_i$$

where L_i and R_i are the left and right restriction operators, respectively, associated with the i -th sub-domain. Both operators are non-square boolean matrices containing ones and zeroes; each row contains at most one non-zero entry. Note that the domain decomposition method typically involves only the right restriction operators. However, the sub-domain matrices in this context are not obtained by assembling the element matrices but by extracting them from the assembled stiffness matrix.

The left restriction operators determine how the rows in the sub-domain matrices are mapped to the rows in the global matrix, while the right restriction operators determine how the columns in the sub-domain matrices are mapped to the columns in the global matrix.

The left restriction operators have the property that:

$$L_i L_j^T = 0 \quad \text{if } i \neq j$$

That is, each left restriction operator L_i has a unique set of non-zero columns corresponding with the internal degrees of freedom in the i -th sub-domain. The right restriction operator is a superset of the left restriction operator; the matrix $(R_i - L_i)$ is a boolean matrix in which the non-zero columns correspond with the overlap degrees of freedom in the i -th sub-domain. The right restriction operator can be used to extract a sub-domain matrix from the global matrix as follows:

$$K_i = R_i K R_i^T$$

The right restriction operator can also be used to extract a sub-domain vector from a global vector as follows:

$$x_i = R_i x$$

in which x_i is a sub-domain vector containing the elements from the global vector x that are associated with the i -th sub-domain. Note that the sub-domain vectors can be assembled into a global vector using the left restriction operators:

$$x = \sum_{i=1}^{n_s} L_i^T x_i$$

The sub-domains are created by a partitioning algorithm—*partitioner* in short—that has three objectives: minimize the number of overlap degrees of freedom; minimize the variation in sub-domain sizes; and separate degrees of freedom associated with different material properties. The first objective is based on the relation between the number of overlap degrees of freedom and the amount of data to be exchanged between the sub-domains. A smaller number

of overlap degrees of freedom results in a smaller amount of data to be exchanged, which results in a lower communication and synchronization overhead, which results in a better parallel performance. The second objective is based on the relation between the size of a sub-domain (the number of degrees of freedom in that sub-domain) and the amount of work associated with that sub-domain. A smaller variation in sub-domain sizes results in a better distribution of the total work load over the sub-domains, which results in a higher parallel efficiency. The third objective is based on the observation that the preconditioner becomes more effective when the large jumps in material properties coincide with the sub-domain boundaries. That is, when degrees of freedom associated with different materials in the underlying finite element model are assigned to different sub-domains.

In general, all three objectives cannot be achieved at the same time. The last objective, in particular, can restrict the way in which the sub-domains are created. The partitioner, therefore, applies a number of heuristics that have been tuned by solving a range of problems, involving different geometries and material compositions, with a varying number of sub-domains.

The partitioner does not actually operate on the level of degrees of freedom but on the level of nodes. That is, it first divides the nodes into overlapping groups and then collects all degrees of freedom that are associated with the nodes in those groups to obtain the sub-domains. This has the advantage that all degrees of freedom associated with the same node are assigned to the same sub-domain. Another advantage is that the number of nodes is less than the number of degrees of freedom so that the partitioner requires less processing time and memory. As there is a one-to-one mapping between the groups of nodes and the sub-domains, the latter term will also be used for the groups of nodes.

The partitioner creates the sub-domain in four stages: form the initial sub-domains by identifying groups of connected nodes associated with the same material type; merge small sub-domains into larger ones; split large sub-domains into smaller ones by applying a conventional graph partitioning algorithm; and extend each sub-domain with one or more levels of adjacent nodes to obtain the final sub-domains. These four steps are described in detail in a separate report [10].

Note that the number of sub-domains varies during the execution of the first three stages. After the third stage, it is equal to the desired number of sub-domains which equals the number of threads. The number of sub-domains remains equal to the number of threads during the execution of the fourth stage.

3.2 The iterative solver

The parallel solver presented here uses a Krylov method to compute the solution of the linear system of equations. It implements both the Conjugate Gradient (CG) and the Generalized Minimum Residual (GMRES) method [13]. Although the latter requires more memory and more work per iteration than CG, it can deal with non-symmetric matrices and preconditioners. This will be shown to be a significant advantage.

A preconditioner is applied (either from the left or from the right) to increase the convergence rate of the Krylov method. The preconditioner is the most important part of the solver as it largely determines the overall performance. It is composed of an additive Schwarz preconditioner that operates on the sub-domain level and an algebraic coarse grid preconditioner that operates on the global level.

The solver executes four types of operations: a vector addition, a scalar vector product, a sparse matrix-vector product, and a preconditioner-vector product. The first operation is trivial to execute in parallel: each thread simply adds its own sub-domain vectors; no data exchange is needed between the threads. The other two types of operations are relatively straightforward to execute and are described below. The last type of operation is explained in the next two sub-sections.

A scalar vector product is computed in parallel as follows:

$$s = x^T y = \sum_{i=1}^{n_s} x_i^T L_i L_i^T y_i$$

in which the matrix $L_i L_i^T$ is a boolean diagonal matrix containing non-zero values for the unique set of degrees of freedom in the i -th sub-domain. The above equation is equivalent with a series of local scalar vector products followed by a global sum that requires data exchanges between threads. The time required to execute these data exchanges is at least proportional to the logarithm of the number of threads, which means that a scalar vector product lowers the parallel efficiency of the solver. The solver, therefore, tries to compute multiple scalar products at the same time with the aim of lowering the communication latency that typically makes up a large fraction of the communication overhead.

A matrix-vector product can be written as:

$$y = Kx = \sum_{i=1}^{n_s} L_i^T K_i R_i x = \sum_{i=1}^{n_s} L_i^T K_i x_i$$

The resulting vector associated with the i -th sub-domain is therefore given by:

$$y_i = R_i y = \sum_{j=1}^{n_s} R_i L_j^T K_j x_j$$

As the matrix product $R_i L_j^T$ is non-zero only if the sub-domains i and j share common degrees of freedom, the matrix-vector product can be computed in two steps. First, each thread independently computes its local matrix-vector product $K_i x_i$. After that, it adds the contributions from the sub-domains that share degrees of freedom with its own sub-domain. This requires a data exchange between threads owning adjacent sub-domains. As these data exchanges can be executed in parallel, the communication overhead does not increase with the number of threads, provided that the available communication bandwidth scales with the number of processing cores.

3.3 The two-level preconditioner

The classic additive Schwarz (AS) preconditioner is constructed as follows [17]:

$$P^{-1} = \sum_{i=1}^n R_i^T P_i^{-1} R_i$$

in which P_i^{-1} is a sub-domain preconditioner that is obtained by means of an incomplete Cholesky decomposition in a similar way as the original PLAXIS solver computes the preconditioner [14]. Because the additive Schwarz preconditioner has the same structure as the global stiffness matrix K , a product of the form

$$y = P^{-1}x$$

can be computed in the same way as a matrix-vector product. In fact, the same data exchange procedures can be used.

The AS preconditioner can be formulated in an alternative way by using both the left and right location operators:

$$P^{-1} = \sum_{i=1}^n L_i^T P_i^{-1} R_i$$

This formulation, also known as the restricted additive Schwarz (RAS) preconditioner, is actually (much) more effective than the classic AS preconditioner [2]. It generally results in a higher convergence rate of the solver, and it requires less data to be exchanged between the threads. A drawback, however, is that the RAS preconditioner is not symmetric, even when the sub-domain preconditioners are symmetric. This means that the RAS preconditioner must be combined with a non-symmetric Krylov method like GMRES.

The coarse preconditioner P_0^{-1} is obtained by projecting the stiffness matrix onto a vector space that forms the algebraic coarse grid. If this vector space is spanned by the columns of the matrix V , then

$$P_0^{-1} = V A^{-1} V^T = V \left(V^T K V \right)^{-1} V^T$$

with A the projected stiffness matrix, also called the coarse matrix. The effectiveness of the coarse preconditioner is determined by the construction of the matrix V . A good choice has proven to be:

$$V = [R_1^T N_1 \cdots R_n^T N_n]$$

That is, the columns of V are formed by the columns of the sub-domain matrices N_i , of which the exact definition will be given later. By introducing the coarse restriction operators C_i , the matrix V can also be written as:

$$V = \sum_{i=1}^n R_i^T N_i C_i$$

Each coarse restriction operator has a unique set of non-zero columns so that

$$C_i C_j^T = \begin{cases} 0 & \text{if } i \neq j, \\ I & \text{if } i = j \end{cases}$$

The coarse restriction operators can be used to extract a coarse sub-domain vector from a global coarse vector:

$$x_i = C_i x$$

Note that the length of the coarse sub-domain vector x_i equals the number of columns in the matrix N_i , and that the length of the global coarse vector x equals the number of columns in the matrix V .

The coarse matrix can be constructed the column wise through a series of parallel matrix-vector products and data exchanges. If A^j denotes the j -th column of the coarse matrix, then:

$$A^j = A e^j = V^T K V e^j = V^T K a^j = V^T b^j = \sum_{i=1}^n C_i^T N_i^T R_i b^j \quad (4)$$

with

$$a^j = V e^j = R_i^T N_i C_i e^j, \quad b^j = K a^j = \sum_{i=1}^n L_i^T K_i R_i a^j$$

and with e^j the j -th column of the identity matrix.

The columns of the coarse matrix are not stored as a single entities; each thread i stores only the partial columns A_i^j that are obtained by applying the restriction operator C_i to the columns of the coarse matrix. These partial columns

can be computed in parallel by executing the following algorithm.

1. Form the coarse sub-domain vector $e_i^j = C_i e^j$. Note that this vector is non-zero for only one thread.
2. Compute $\tilde{a}_i^j = N_i e_i^j$.
3. Exchange data to obtain $a_i^j = \sum_{k=1}^n R_i R_k^T \tilde{a}_k^j$.
4. Compute $\tilde{b}_i^j = K_i a_i^j$.
5. Exchange data to obtain $b_i^j = \sum_{k=1}^n R_i L_k^T \tilde{b}_k^j$.
6. Compute $A_i^j = C_i A^j = N_i^T b_i^j$.

Although this algorithm is relatively straightforward to implement, it does not exploit the special structure of the unit vectors e^j . To be precise, it will perform a number of matrix-vector multiplications involving zero vectors. It is possible to implement a more efficient algorithm, but the gain in performance would only be significant for larger numbers of sub-domains.

The coarse preconditioner is formed by computing the QR-decomposition [6] of the coarse matrix:

$$A = QU$$

in which Q is an orthonormal matrix and U is an upper triangular matrix. They can be obtained by applying a Gram-Schmidt orthonormalization procedure to the columns of A . To execute this procedure in parallel, each thread computes the partial decomposition

$$A_i = C_i A = C_i QU = Q_i U$$

This means that each thread stores the sub-matrix Q_i and a full copy of the matrix U . A drawback of this procedure is that a backward substitution of the form

$$y = U^{-1}x$$

takes a constant time, regardless of the number of threads. A more scalable approach is to explicitly compute the inverse of the upper triangular matrix U so that a backward substitution becomes a matrix-vector product. Such a product can be executed in parallel by dividing the inverse matrix U^{-1} row-wise between the threads in the same way as the matrix Q . That is, each thread stores the sub-matrix

$$\left(U^{-1}\right)_i = C_i U^{-1}$$

The inverse matrix U^{-1} can be computed by executing a parallel algorithm that is similar to the GCR algorithm [4, 9].

The application of the coarse preconditioner involves a series of parallel matrix-vector products, inner products, and data exchanges. This follows from:

$$\begin{aligned} y &= P_0^{-1}x = VA^{-1}V^T x = VU^{-1}Q^T V^T x \\ &= VU^{-1}Q^T a = VU^{-1}b = Vc \\ &= \sum_{i=1}^n R_i^T N_i C_i c \end{aligned}$$

with

$$\begin{aligned} a &= V^T x = \sum_{i=1}^n C_i^T N_i^T R_i x \\ b &= Q^T a = \sum_{i=1}^n Q_i^T C_i a \\ c &= U^{-1}b = \sum_{i=1}^n C_i^T \left(U^{-1}\right)_i b \end{aligned}$$

To evaluate these equations, all threads execute the following procedure in parallel.

1. Compute $a_i = N_i^T x_i$ with $x_i = R_i x$.
2. Compute $b_i = Q_i^T a_i$.
3. Compute $b = \sum_{i=1}^n b_i$. This requires a global sum over all threads. Each thread stores a copy of b .
4. Compute $c_i = \left(U^{-1}\right)_i b$.
5. Compute $\tilde{y}_i = N_i c_i$.
6. Exchange data to obtain $y_i = \sum_{k=1}^n R_i R_k^T \tilde{y}_k$.

The coarse preconditioner is most effective when it removes problematic eigenvalues from the stiffness matrix that cannot be removed by the RAS preconditioner. In geomechanical applications, these eigenvalues are typically related to the rigid body modes of the sub-domains. This implies that each matrix N_i contains six columns representing three rigid translations and three rigid (linearized) rotations of the i -th sub-domain. The six column vectors can be viewed as a series of sextets, one for each node in the sub-domain, representing the three displacements along the global coordinate axes and the rotations about those axes. These sextets are given by:

$$\begin{aligned} \text{translation along the x-axis} &: [1, 0, 0, 0, 0, 0] \\ \text{translation along the y-axis} &: [0, 1, 0, 0, 0, 0] \\ \text{translation along the z-axis} &: [0, 0, 1, 0, 0, 0] \\ \text{rotation about the x-axis} &: [0, -z, y, 1, 0, 0] \\ \text{rotation about the y-axis} &: [z, 0, -x, 0, 1, 0] \\ \text{rotation about the z-axis} &: [-y, x, 0, 0, 0, 1] \end{aligned}$$

where x , y , and z are the coordinates of the node with which the sextet is associated. Note that the rotational degrees of freedom is only present when a shell element is attached to a node. Otherwise, only the translational degrees of freedom will be present.

The rigid body modes do not need to be computed for each sub-domain on an individual basis. An alternative approach is to setup a global matrix N of which the six columns represent the six rigid body modes of the complete finite element model without constraints. The rigid body modes for each sub-domain can then simply be obtained by applying the right restriction operator:

$$N_i = R_i N$$

This approach simplifies the programming interface of the solver as the user only has to provide one matrix containing the global rigid body modes.

In the case that a sub-domain contains degrees of freedom that are associated with different materials, the effectiveness of the coarse preconditioner can be increased by using more than six rigid body modes. To be precise, a set of rigid body modes is used for each group of connected degrees of freedom that is associated with the same material. Although this results in a larger coarse grid, the extra computational effort can be more than offset by the increased effectiveness of the preconditioner. The current implementation of the solver created at most four groups per sub-domain. If the number of groups is larger, than adjacent groups are merged until four groups remain.

The coarse preconditioner can be combined with the AS or RAS preconditioner in an additive or a multiplicative way. In the former approach, the coarse preconditioner is simply added to the AS or RAS preconditioner. In the latter approach, the two preconditioners are applied as follows:

$$\begin{aligned}\tilde{y} &= P^{-1}x \\ \tilde{r} &= x - K\tilde{y} \\ y &= P_0^{-1}\tilde{r}\end{aligned}$$

This scheme has proven to be much more efficient, even though it requires one extra matrix-vector product per solver iteration. Note that the combined preconditioner is not symmetric, even when using a symmetric AS preconditioner. A symmetric preconditioner can be obtained by applying the AS preconditioner a second time in a multiplicative way, but this raises the computational cost of the total preconditioner considerably.

4 Implementation of the solver

The solver has been implemented in C as a stand-alone, portable library with programming interfaces for Fortran 77 and Fortran 90. The library is more or less a black box that can be linked with a serial program. It provides a small set of functions for defining a sparse matrix; for creating a preconditioner associated with that matrix; for solving a system of equations; for setting runtime parameters and options; and

for retrieving information about the state and progress of the solver.

4.1 The message passing programming model

The solver library is based on a multi-threaded message passing programming model to execute the solution algorithms in parallel. When a program defines a sparse matrix, the solver library spawns a number of threads, creates the sub-domains, and assigns one sub-domain to each thread. A custom-made and light-weight message passing module is used to exchange data between adjacent sub-domains and to perform global reduction operations such as a global sum. Because the threads share a common memory space, messages sent by one thread to another can be copied directly from the source memory area to the destination memory area without intervention of the operating system, and without using intermediate memory buffers. To synchronize the execution of the threads, the message passing module uses POSIX mutexes and condition variables, or Windows serial sections and event objects. It tries to reduce system-call overhead by spinning for a short while when one thread needs to wait for another. Note that spinning is turned off when the number of threads exceeds the number of processor cores.

A significant part of the implementation is dedicated to error handling. If a system error or a numerical error occurs, the library will try very hard to return to a well-defined state and report the error to the calling program. A complicated situation occurs when a thread needs to abort an operation because of an error that is local to that thread. If this happens, the library will signal all other threads so that they can abort their operation too. In this way, a thread will not be waiting forever on an event that never occurs because another thread has aborted an operation.

By using threads and an integrated message passing module, the solver library can be incorporated easily into an existing serial program; there is no need to install external libraries and set-up a special runtime environment. A drawback of a thread-based implementation is that the solver cannot make use of distributed memory systems. However, it would not be much work to replace the integrated message passing module by an external message passing library that supports distributed memory systems (one based on MPI, for instance).

4.2 Some implementation details

Both the GMRES algorithm and the coarse grid preconditioner make use of the Gram-Schmidt orthogonalization procedure. The solver implements a combination of the classic and the modified Gram-Schmidt procedure to find

a balance between numerical accuracy and low communication overhead. The hybrid algorithm applies the modified Gram-Schmidt procedure to bundles of four (or six in case of the coarse preconditioner) vectors at a time, as shown in Algorithm 1. Because the four scalar vector products in the first inner for-loop are independent, only one global reduction operation is required to sum the partial products computed by the threads.

Algorithm 1 Create a vector u that is orthonormal to a set of n vectors v_i . This algorithm combines the classic and modified Gram-Schmidt procedures.

```

1: for  $i = 1$  to  $n$  step 4 do
2:   for  $j = i$  to  $i + 3$  do
3:      $a_j = u^T v_j$ 
4:   end for
5:   for  $j = i$  to  $i + 3$  do
6:      $u = u - a_j v_j$ 
7:   end for
8: end for

```

To improve the performance of the GMRES algorithm, the solver bundles each consecutive set of four Krylov vectors into one dense matrix of which the elements are stored row wise in memory. This makes it possible to replace the four scalar vector products in the first inner loop of hybrid Gram-Schmidt procedure by a dense matrix-vector product. This, in turn, increases the utilization of the memory cache and, therefore, reduces the execution time of the solver. Note that if the number of Krylov vectors is not a multiple of four, then up to three zero vectors are added to fill a bundle of four Krylov vectors.

The solver speeds up the construction of the coarse matrix by computing bundles of four columns at the same time. To see how this helps, consider the main operation in the computation of one coarse matrix column: a matrix-vector product involving the global stiffness matrix. This




Fig. 1 The geomechanical model that has been used to analyze the performance of the solver

Table 1 The stiffness moduli associated with the different parts of the model

Part	Stiffness modulus [MPa]
Yellow-colored soil	1.0
Green-colored soil	$1.0 \cdot 10^3$
Concrete block	$1.0 \cdot 10^5$

operation requires that all matrix elements are transported from the (slow) main memory to the processor. The cache memory is typically not effective because the amount of data stored in the stiffness matrix tends to be much larger than the size of the cache. The execution time of one matrix-vector product is therefore dictated primarily by the memory bandwidth and much less by the processor speed. The only way to overcome this bottleneck is to reuse each matrix element once it has been fetched from the main memory. This is exactly what happens when multiple, independent, matrix-vector products are executed in one pass.

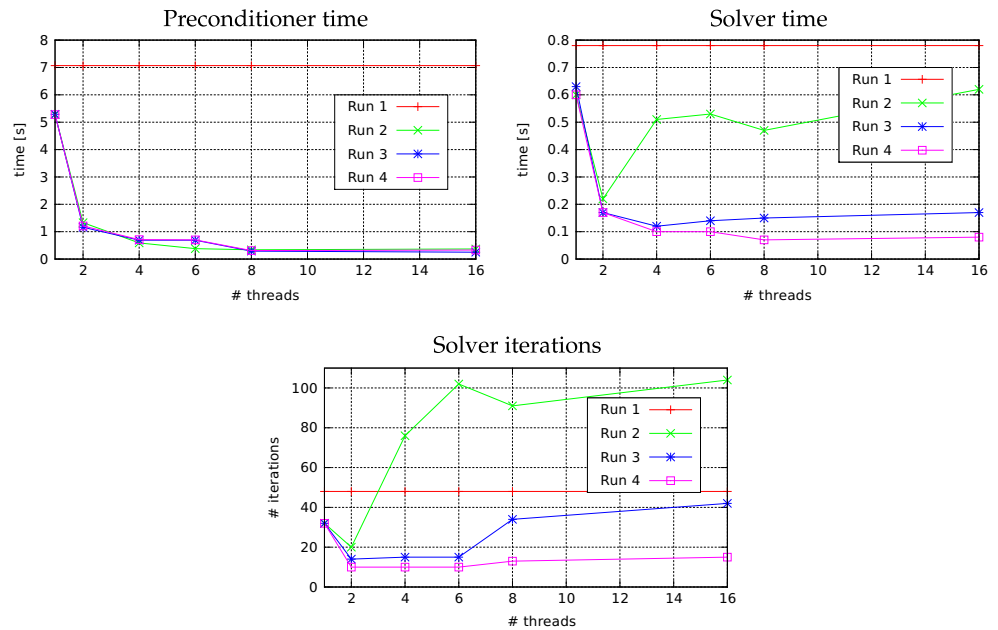
The maximum number of Krylov vectors is an important parameter in the GMRES algorithm. If this parameter is too large, then the solver may run out of memory or the time per GMRES iteration may become excessively large. If this parameter is too small, on the other hand, then GMRES may be forced to make many restarts with a low convergence rate as a result. As the average user can probably not make an informed choice, the solver sets the maximum number of Krylov vectors so that the space needed for storing the Krylov vectors is at most equal to the space required for storing the stiffness matrix and the preconditioner. This automatic approach seems to strike a good balance between performance and memory usage.

Although the original solver in PLAXIS and the new solver use essentially the same Cholesky preconditioner, their implementations differ in three ways. These differences affect the way in which the preconditioners perform, as will be shown in Section 5.

The first difference concerns the algorithm for ordering the degrees of freedom before the incomplete factors are computed. The original solver uses the reverse Cuthill–McKee algorithm [5] in combination with an algorithm that tries to reduce the bandwidth of the stiffness matrix by sorting the nodes according to their X, Y, or Z coordinates. The new solver, on the other hand, uses a slightly modified reverse Cuthill–McKee algorithm in which the vertex degrees are updated during the execution of the algorithm instead of keeping those degrees constant.

The second difference is that the drop tolerance is increased in a different way when the number of fill-ins exceeds a given limit. While the new solver uses the method described by Saukh [14], the original solver maintains a

Fig. 2 The performance results that have been obtained for the four solver runs listed in Table 2. The upper left graph shows the wall-clock time required to create the preconditioner as function of the number of threads. The upper right graph shows the time required to solve one system of equations as function of the number of threads. The lower graph shows the number of solver iterations as function of the number of threads. Note that the number of threads equals the number of sub-domains



cumulative histogram that keeps track of the number of non-zero entries in the Cholesky factors as function of their magnitude. This has the advantage that the solver can quickly determine a new drop tolerance when the number of non-zero entries exceeds a predefined limit.

The third difference concerns the method for computing the maximum number of fill-ins allowed per row/column. The original solver imposes no maximum; it only drops entries that are smaller than the drop tolerance. The new solver does the same thing, but it also limits the total number of fill-ins per row/column to a user-defined maximum. If the number of non-zero entries is larger than that maximum, say m , the solver only keeps the m largest entries.

4.3 Hardware considerations

To achieve good parallel performance, the solver requires sufficient memory bandwidth for all processor cores combined. In each solver iteration, all threads need to fetch their sub-domain matrix and preconditioner from the main memory as the cache memory is typically much too small to store all the data. Unfortunately it is not possible to reduce the memory traffic by using (parts of) a sub-domain matrix or preconditioner multiple times within one iteration.

The available memory bandwidth must be sufficient to provide each thread full-speed access to the main memory. Otherwise, one thread will have to wait for another while it is fetching data from memory. This used to be a real problem in shared memory machines as the processing cores had to share one memory bus. Fortunately, both AMD and Intel have adopted a memory architecture that can scale with the number of processor cores. That is, each CPU (containing multiple processing cores) has its own memory connection; adding more CPUs will increase the number of memory connections. Note that some computer vendors will try to reduce costs by routing multiple memory connections through one CPU. This will reduce the available memory bandwidth and may reduce the parallel performance of the solver. One simple way to maximize the memory bandwidth is to fill all available memory slots so that all memory connections are used.

5 Performance of the solver

The performance of the solver has been measured by solving various systems of equations on a workstation with two quad-core Intel Xeon E5620 processors, 24 GB memory

Table 2 Description of the four solver runs

Run	Description
1	The original PLAXIS solver.
2	The new solver, without the physics-based domain decomposition method and without the coarse grid preconditioner.
3	The new solver, with the physics-based domain decomposition method but without the coarse grid preconditioner.
4	The new solver, with the physics-based domain decomposition method and with the coarse grid preconditioner.

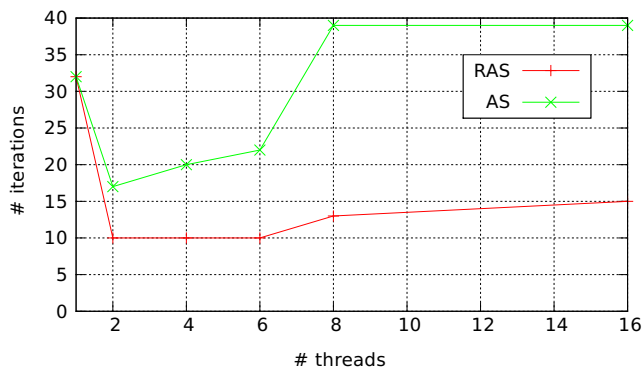


Fig. 3 The number of solver iterations with the AS and RAS preconditioners versus the number of threads

and running Windows 7. Both processors can access the main memory through independent interconnects so that the maximum possible memory bandwidth can be obtained. The systems of equations have been obtained from various geomechanical problems using the commercial finite element program PLAXIS. Note that all reported execution times are wall-clock times.

The next sub-section uses an academic problem to analyze the various components making up the solver. The goal here is to gain insight in the way that those components effect the overall performance of the solver. The second sub-section shows what kind of performance can be obtained for two representative geomechanical problems. The third and last sub-section shows the performance of the solver for a difficult problem. The aim of this sub-section is to highlight the current limitations of the solver.

5.1 Analysis of the performance

Figure 1 shows the geomechanical model that has been used to analyze the performance of the various solver components. The model comprises about 1,400 15-node wedge elements and 10,000 degrees of freedom, and consists of a rectangular block of soil that is composed of five layers with large variations in stiffness. A uniform load (indicated by the blue arrows) is applied to an embedded block of concrete. A simple linear elastic material model (based on Hooke’s elasticity law) is used to describe the deformations

Fig. 4 The two representative models that have been used to test the real-world performance of the solver. The colors indicate different material properties; see Table 3




Table 3 The stiffness moduli associated with the different parts of the two models shown in Fig. 4

Part	Color	Stiffness modulus [MPa]
Concrete shell	Blue	$3.0 \cdot 10^4$
Soil layer 1	Yellow	$1.0 \cdot 10^1$
Soil layer 2	Green	1.5
Soil layer 3	Orange	$4.0 \cdot 10^1$
Soil layer 4	Light Green	9.0
Soil layer 5	Purple	$3.0 \cdot 10^2$

The soil layers are listed from top to bottom

of the soil and the concrete. The relevant material parameter in this model is the stiffness modulus that differs five orders of magnitude between the different parts of the model; see Table 1.

Figure 2 shows the performance results that have been obtained for the four solver runs listed in Table 2. The first run makes use of the original solver in PLAXIS that is based on the CG algorithm in combination with an incomplete Cholesky decomposition as preconditioner. The second run uses the new solver but without the physics-based domain decomposition method; it simply uses the METIS [8] partitioning library to create the sub-domains without taking advantage of the underlying model. The third run uses the new solver with the physics-based domain decomposition method, but without the algebraic coarse grid preconditioner. This preconditioner is added in the fourth run.

Note that the original solver uses only one thread; the results obtained with the original solver have been plotted as horizontal lines in the three graphs. Also note that the runs with the new solver are based on the GMRES algorithm in combination with the RAS preconditioner. Finally, note that the maximum number of threads (16) exceeds the number of processor cores (eight). The reason for this is that a larger number of threads—and therefore sub-domains—provides some insight into the numerical scalability of the solver.

One can draw the following four conclusions from Fig. 2. First, the original solver is less efficient than the new solver when using a single thread or sub-domain. This is because the original solver implements the incomplete Cholesky

Table 4 Performance results obtained for the two models

Model 1				
Solver	# Threads	Precon [s]	Solve [s]	# Iter
PARDISO	8	200	25	1
Original	1	320	130	140
	1	140	75	134
New	2	84	29	80
	4	43	29	111
	8	23	21	113
Model 2				
Solver	# Threads	Precon [s]	Solve [s]	# Iter
PARDISO	8	71	12	1
Original	1	170	19	32
	1	58	19	62
New	2	22	15	74
	4	19	4.8	39
	8	14	3.9	45

The first column specifies the solver for which the results have been obtained. The second column lists the number of threads or sub-domains. The remaining three columns list the wall-clock time required to set-up the preconditioner and to solve the system of equations, and the number of solver iterations. Note that in the case of PARDISO, the third column specifies the time required to factor the matrix

preconditioner in a slightly different—and apparently less efficient—way than the new solver. Note, however, that this is not always the case; for some other problems the situation is reversed and the original solver is more efficient.

The second conclusion is that the physics-based domain decomposition method is absolutely essential to get decent performance with the new solver; without it, the number of solver iterations varies wildly with the number of sub-domains. Indeed, the maximum number of iterations is about a factor five larger than the minimum number of iterations.

The third conclusion is that the coarse grid preconditioner raises the efficiency of the solver by a significant measure. It significantly lowers both the number of solver iterations and the solution time, without adding much to the time that is required to set-up the preconditioner.

The final conclusion is that the new solver scales very well with the number of threads. Indeed, both the time required to set-up the preconditioner and the time required to solve the system of equations are more than eight times smaller when using eight threads instead of one. This is partly caused by a numerical speedup: the sub-domain wise incomplete Cholesky preconditioner requires less retries to find the appropriate drop tolerance and the solver requires less iterations to satisfy the convergence criterion. The remainder of the speedup is the result of using more than one processor core.

Figure 3 shows the number of solver iterations when the RAS preconditioner is replaced by the AS preconditioner. It is clear that the RAS preconditioner is much more efficient in this case than the AS preconditioner. Indeed, the effectiveness of the RAS preconditioner justifies the use of GMRES instead of CG.

5.2 Results for two representative models

The new solver has been the standard iterative solver in PLAXIS since the end of 2010 and it has been used to solve a wide range of geomechanical problems. Two representative models (see Fig. 4) have been selected to show what kind of performance can be obtained in practice. Model 1 involves 680,000 degrees of freedom and is composed of a concrete foundation on top of several soil layers with a varying stiffness. Model 2 involves 414,000 degrees of freedom and is similar to Model 1, except that the foundation has been replaced by a tunnel. Table 3 lists the stiffness moduli associated with the different parts of the two models.

The soil is modeled with 10-node tetrahedral elements with quadratic shape functions. The concrete foundation and




Fig. 5 The parallel speedup obtained for the two models. The curve labeled “Precon” depicts the parallel speedup of the preconditioner. The curve labeled “Solve” depicts the parallel speedup of the solver,

excluding the time required to set up the preconditioner. The curve labeled “Total” depicts the total speedup that is based on the time for setting up the preconditioner and for solving the system of equations

the tunnel walls are modeled with 6-node triangular Mindlin elements that can handle bending moments, normal forces, and shear deformations. Both types of elements incorporate a linear elastic material model based on Hooke's elasticity law.

Table 4 shows the performance results that have been obtained for the two models. It also shows the results obtained with the original solver and with PARDISO [15], a fast parallel direct solver that is available on a commercial basis. Note that the original solver uses a single processing core and PARDISO uses all eight processing cores. Figure 5 shows the parallel speedup that has been obtained with the new solver. The speedup is defined as the ratio between the execution time with n threads over the execution time with one thread.

Four observations can be made when looking at these results. First, the performance of the new solver varies considerably between the two models, especially when looking at the number of solver iterations versus the number of threads. One reason is that the solver is not able to separate the different materials well enough when the number of threads (and sub-domains) is relatively small. Another reason is that the solver needs to make a trade-off between a well-balanced work load and sub-domains that do not cut through material boundaries, resulting in a sub-optimal configuration of sub-domains. Because of this, the solver is allowed to use twice as many threads as processor cores, unless the number of threads has been explicitly set by the user.

The second observation is that the new solver is faster than the original one when using a single thread. The reason is that the new solver implements a somewhat different and more efficient incomplete Cholesky decomposition algorithm. This algorithm does not only lower the time to set-up the preconditioner, but also to solve the system of equations, even when the new solver requires more iterations.

The third observation is that the construction of the preconditioner scales well with the number of threads. This is not surprising as the construction of the preconditioner requires relatively a little communication between the threads; a large fraction of the computations is local to the threads. When solving the system of equations, on the other




Fig. 6 A model of the Toulon tunnel in France. Part of the soil is not shown so that the tunnel is visible

Table 5 Performance results obtained with PARDISO, the original solver, and the new solver for the model of the Toulon tunnel

Solver	# Threads	Precon [s]	Solve [s]	# Iter
PARDISO	8	16.5	6.8	1
Original	1	290	17	33
	1	110	19	59
New	2	31	50	214
	4	54	38	150
	8	17	27	139

hand, the communication between the threads makes up a larger fraction of the execution time. In addition, there is less scope for reusing data that resides in the local caches of the processor cores so that the memory bandwidth becomes a more limiting factor.

The final observation is that the new solver is faster than PARDISO when eight threads are used, both in preconditioning time and solve time. The total performance gain in the given examples ranges from 4.6 to 5.

5.3 Seeking the limits of the solver

The new solver still has some limitations that can become apparent when solving difficult problems. One of those problems involves a nonlinear model for simulating the construction of a tunnel at Toulon in France; see Fig. 6. The model has about 750,000 degrees of freedom and consists of solid elements, shell elements, and embedded pile elements with a stiffness modulus that ranges from 1.0 to $2.1 \cdot 10^5$ MPa. A complete simulation of the construction process requires the solution of multiple linear systems of equations. Only the first system is considered here because the subsequent systems of equations are slightly different when using different solvers or different numbers of threads.

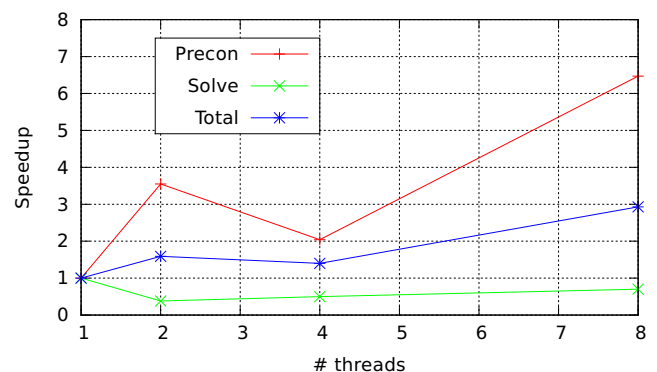


Fig. 7 The parallel speedup obtained for the model of the Toulon tunnel

The soil and the tunnel walls are modeled with the same type of elements that have been used in the previous two models. The embedded piles are modeled with 3-node Mindlin line elements that can handle normal, shear, and moment deformations. Because the piles can pass through solid elements at arbitrary locations, interface elements are used to model the interaction between the piles and the solid elements.

The problem is difficult to solve for at least three reasons. First, there is the large variation in the stiffness modulus, but this is also the case for the two models in the previous sub-section. A second reason is that the simulation of the construction process involves cutting away part of the soil, resulting in a very small elements. A third reason is that the beam elements only have a very small rotational stiffness to prevent the system of equations from becoming singular. All these three factors lead to a stiffness matrix of which the diagonal elements differ 14 orders of magnitude. This does not need to be a problem per se, but it does indicate that this is a non-standard problem.

Table 5 and Fig. 7 show the performance results that have been obtained for the model of the Toulon tunnel. It is clear that the new solver does not perform very well because the number of solver iterations increases substantially when using more than one thread. The exact cause of this is not known at this moment. One possibility is that a node-based domain decomposition algorithm is not able to separate the different materials and element types into different sub-domains. In that case, switching to an element-based decomposition algorithm would help, but this would make it more difficult to use the solver as a black box.

Note that in this case PARDISO is much faster than both the new and the old solver.

6 Conclusions

The solver proposed in this paper can solve geomechanical problems efficiently on multi-core machines, even when they involve materials with very different properties, for instance a stiffness modulus that varies five orders of magnitude. As the solver is not tightly coupled to the geomechanical nature of the underlying problem, it is reasonable to expect that the solver can also be used efficiently for solving other types of solid mechanics problems involving large variations in material properties.

The physics-based domain decomposition approach plays an important role in the efficiency of the solver. Examples with straight-forward domain decomposition of “layered” finite element models showed a low convergence rate, whereas physics-based domain decomposition

clearly improves the efficiency and robustness of the solver. This means that the solver can lower the runtime significantly even when using only one processor core.

The new solver has been the standard iterative solver in PLAXIS since 2010 and it has shown good performance for a wide range of problems in geomechanics. Indeed, the new solver is between seven and eight times faster than the original solver on an eight-core machine for two representative problems discussed in this paper.

The new solver has its limitations and in some cases it performs worse than the original solver. This has been illustrated with a difficult problem involving the simulation of the construction of a tunnel at Toulon in France. The reason for the sub-optimal performance is not yet completely clear; this is the subject of future research.

The solver has been implemented as a kind of black box that can easily be linked with a serial program without needing additional libraries or a special runtime environment. To achieve this, the solver is currently limited to shared memory machines. However, the use of a message passing programming model makes it possible to port the solver to distributed memory systems without much effort.

References

1. Brinkgreve, R.B.J., Swolfs, W.M., Engin, E.: PLAXIS 3D 2011 User Manual. Plaxis BV, Delft, The Netherlands (2011)
2. Cai, X.-C., Dryja, M., Sarkis, M.: Restricted additive schwarz preconditioners with harmonic overlap for symmetric positive definite linear systems. *SIAM J. Numer. Anal.* **41**, 1209–1231 (2003)
3. Dongarra, J.J., Duff, L.S., Sorensen, D.C., van der Vorst, H.A.: Numerical linear algebra for high-performance computers. SIAM, Philadelphia (1998)
4. Eisenstat, S.C., Elman, H.C., Schultz, M.H.: Variational iterative methods for non-symmetric systems of linear equations. *SIAM J. Num. Anal.* **20**, 345–357 (1983)
5. George, J.A., Liu, J.W.-H.: Computer solution of large sparse positive definite systems Prentice-Hall (1981)
6. Golub, G.H., van Loan, C.F.: Matrix computations. The John Hopkins University Press, Baltimore (1996)
7. Jönsthövel, T.B., van Gijzen, M.B., Vuik, C., Kasbergen, C., Scarpas, A.: Preconditioned Conjugate Gradient method enhanced by deflation of rigid body modes applied to composite materials. *Comput. Model. Eng. Sci.* **47**, 97–118 (2009)
8. Karypis, G., Kumar, V.: Metis. A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices Version 4, Vol. 0. Technical report, University of Minnesota, Department of Computer Science / Army HPC Research Center, Minneapolis, USA (Sep 1998)
9. Lingem, F.J.: Design of an object oriented finite element package for parallel computers. PhD thesis, Delft University of Technology. Delft, The Netherlands (2000)

10. Lingen, F.J., Bonnier, P.G., Brinkgreve, R.B.J., van Gijzen, M.B., Vuik, C.: A parallel linear solver exploiting the physical properties of the underlying mechanical problem Technical Report, Vol. 12-12. Delft University of Technology, Delft, The Netherlands (Sep 2012)
11. Meijerink, J.A., Van der Vorst, H.A.: An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-matrix. *Math. Comp.* **31**, 148–162 (1977)
12. Saad, Y.: Iterative methods for sparse linear systems, Second Edition. SIAM, Philadelphia (2003)
13. Saad, Y., Schultz, M.H.: GMRES A Generalized Minimal Residual algorithm for solving non-symmetric linear systems. *SIAM J. Sci. Stat. Comput.* **7**, 856–869 (1986)
14. Saikh, S.: Incomplete Cholesky factorization in fixed memory with flexible drop-tolerance strategy (Sep 2003)
15. Schenk, O., Gärtner, K.: Solving unsymmetric sparse systems of linear equations with PARDISO. *Journal of Future Generation Computer Systems* **20**(3), 475–487 (2004)
16. Tang, J.M., MacLachlan, S.P., Nabben, R., Vuik, C.: A comparison of two-level preconditioners based on multigrid and deflation. *SIAM. J. Matrix Anal. and Appl.* **31**, 1715–1739 (2010)
17. Toselli, A., Widlund, W.: Domain decomposition methods. *Computational Mathematics*, Vol. 34. Springer, Berlin (2005)
18. van der Vorst, H.A.: High performance preconditioning. *SIAM J. Sci. Stat. Comput.* **10**, 1174–1185 (1989)
19. Vuik, C., Segal, A., Meijerink, J.A.: An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients. *J. Comp. Phys.* **152**, 385–403 (1999)