Stabilized and Block Approximate Inverse Preconditioners for Problems in Solid and Structural Mechanics

Michele Benzi[*] Reijo Kouhia[†] Miroslav Tůma[‡]

October 30, 2000

Abstract

The solution of linear systems arising in the finite element analysis of shells and solids by the preconditioned conjugate gradient method is considered. Stabilized and block versions of the AINV factorized approximate inverse preconditioner are presented and tested on a variety of difficult problems. Comparisons with other preconditioning methods are also included.

Keywords: preconditioning, conjugate gradient, factorized sparse approximate inverses, block algorithms, finite elements, shells

1 Introduction

In this paper, we address the question of devising robust and efficient preconditioners for the conjugate gradient method applied to problems in solid and structural mechanics. Our main focus is on sparse linear systems of the type

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \tag{1}$$

arising from finite element models of thin shells, which are known to be especially challenging for iterative linear solvers. Problems from the modeling of three-dimensional solids will also be considered. The preconditioners

[*]Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA (benzi@mathcs.emory.edu).

[†]Laboratory of Structural Mechanics, Helsinki University of Technology, PO Box 2100, FIN 02015 HUT, Finland (reijo.kouhia@hut.fi).

[‡]Institute of Computer Science, Academy of Sciences of the Czech Republic, Pod vodárenskou věží 2, 182 07 Prague 8, Czech Republic (tuma@cs.cas.cz).

examined in this paper are mainly sparse approximate inverses in factorized form; however, other techniques, like block Jacobi and point and block versions of incomplete Cholesky preconditioning, will also be considered.

A *sparse approximate inverse preconditioner* is a sparse matrix $\boldsymbol{M}$ that directly approximates the inverse of the coefficient matrix $\boldsymbol{A}$:

$$\boldsymbol{M} \approx \boldsymbol{A}^{-1}.$$

The (left) preconditioned system, therefore, is of the form

$$\boldsymbol{MAx} = \boldsymbol{Mb}.$$

When $\boldsymbol{A}$ is symmetric positive definite and the preconditioned conjugate gradient (PCG) method is being used, the preconditioner must also be symmetric and positive definite. One way to insure this is to take $\boldsymbol{M}$ to be in factorized form, i.e.,

$$\boldsymbol{M} = \boldsymbol{ZZ}^T$$

for a nonsingular sparse matrix $\boldsymbol{Z}$. If $\boldsymbol{Z}$ is triangular, then nonsingularity is trivial to check, and the preconditioner can be regarded as an incomplete factorization of $\boldsymbol{A}^{-1}$. Thus, the preconditioned system takes the form

$$\boldsymbol{Z}^T\boldsymbol{AZu} = \boldsymbol{Z}^T\boldsymbol{b} \tag{2}$$

and the solution to the original system (1) is given by $\boldsymbol{x} = \boldsymbol{Zu}$. Note that the coefficient matrix in (2) is symmetric positive definite. In this paper, we consider exclusively approximate inverses in factorized form.

Several techniques have been proposed to compute factorized sparse approximate inverse preconditioners; see, e.g, [32], [10], [42], [25], [31] and the references therein. There are two major issues in computing a sparse approximate inverse: the first is the choice of an appropriate sparsity pattern for the approximate inverse factor $\boldsymbol{Z}$, and the second is the actual computation of the entries of $\boldsymbol{Z}$. Determining a good sparsity pattern for the approximate inverse can be difficult, especially for unstructured problems. Techniques for guessing a sparsity pattern *a priori* have been investigated, for instance, in [19]. Once a sparsity pattern has been found, least squares techniques can be used to compute an approximate inverse with the given sparsity pattern [32]. Alternatively, the sparsity pattern can be computed dynamically, together with the nonzero entries, by applying a drop tolerance in the course of an inverse factorization algorithm. This is the approach taken in the AINV algorithm [10], and in this paper we will focus on variants of this technique.

Our interest in approximate inverse preconditioners stems in part from the fact that they are comparatively easy to implement on modern high-performance computers, particularly vector and parallel machines, since their application only involves matrix-vector products. This is in contrast with more traditional preconditioners, like incomplete factorization methods, which require highly sequential triangular solves and are difficult to implement efficiently on current parallel architectures, especially for unstructured problems. Another motivation for considering approximate inverse techniques instead of more traditional incomplete factorization preconditioners is related to the issue of numerical stability. In a previous paper [8], we ran into severe instability problems when trying to apply approximate inverse techniques to the stiffness matrices arising in thin shell analysis. Similar problems also plague the incomplete Cholesky methods, on occasions. In this paper we report on recent progress that was made towards ensuring the stable computation of factorized approximate inverse preconditioners for thin shell problems and other highly ill-conditioned matrices. We will argue that the new techniques are inherently more stable than standard incomplete factorization methods. Indeed, factorized approximate inverse preconditioners can be computed without the need for any kind of diagonal correction; this is not true, in general, for incomplete factorizations of $\boldsymbol{A}$.

Another contribution of this paper is to present block algorithms for both the basic and stabilized approximate inverse preconditioners, and to assess their performance on thin shell and solid mechanics problems. These problems have a natural block structure (stemming from the fact that there are several unknowns associated with each node), and exploiting this structure can result in increased robustness and improved convergence rates of the PCG iteration. Furthermore, dense blocks can be exploited to achieve reasonable performance on modern cache-based architectures.

The remainder of the paper is organized as follows. In section 2 we review the basic AINV technique together with its stabilized variants. Block algorithms are described in section 3. Section 4 is devoted to experimental results. We present some conclusions in section 5.

2 A stable approximate inverse algorithm

In this section we recall the basic AINV technique [10] and its stabilized variants [7], [30]. The AINV algorithm builds a factorized sparse approximate inverse of the form

$$\boldsymbol{M} = \boldsymbol{Z}\boldsymbol{D}^{-1}\boldsymbol{Z}^T \approx \boldsymbol{A}^{-1} \tag{3}$$

where $\boldsymbol{Z}$ is a sparse unit upper triangular matrix and $\boldsymbol{D}$ is diagonal.

The algorithm computes $\boldsymbol{Z}$ and $\boldsymbol{D}$ directly from $\boldsymbol{A}$ by means of an incomplete $\boldsymbol{A}$-orthogonalization process (that is, a Gram-Schmidt process with respect to the "energy" inner product $\langle \boldsymbol{x}, \boldsymbol{y} \rangle := \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{y}$) applied to the unit basis vectors $\{\boldsymbol{e}_i\}_{i=1}^n$. In this process, small elements are dropped to preserve sparsity. The resulting vectors $\{\boldsymbol{z}_i\}_{i=1}^n$ are approximately conjugate with respect to $\boldsymbol{A}$, i.e., approximately $\boldsymbol{A}$-orthogonal. Letting $\boldsymbol{Z} = [\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_n]$, then $\boldsymbol{Z}^T \boldsymbol{A} \boldsymbol{Z}$ is approximately diagonal. Letting $\boldsymbol{D} = \mathrm{diag}(p_1, p_2, \ldots, p_n)$ where $p_i = \boldsymbol{z}_i^T \boldsymbol{A} \boldsymbol{z}_i$, it follows that $\boldsymbol{Z}\boldsymbol{D}^{-1}\boldsymbol{Z}^T \approx \boldsymbol{A}^{-1}$. The approximate inverse factor $\boldsymbol{Z}$ is a sparse approximation of the inverse of the $\boldsymbol{L}^T$ factor in the $\boldsymbol{L}\boldsymbol{D}\boldsymbol{L}^T$ decomposition of $\boldsymbol{A}$. When the $\boldsymbol{A}$-orthogonalization process is performed exactly, the diagonal matrix $\boldsymbol{D}$ is the same in the two decompositions, and contains the *pivots* down the main diagonal.

The underlying assumption here is that many entries in $\boldsymbol{L}^{-1}$ are small in magnitude. This is true for many problems of practical interest, particularly for discretizations of boundary value problems for partial differential operators of elliptic type, the Green's function of which is known to exhibit fast decay in many cases. Sparsity can also be achieved by combining dropping of small entries with suitable sparse matrix orderings, such as nested dissection and minimum degree. These orderings are beneficial in that they result in smaller time and space requirements for forming and storing the preconditioner, while at the same time improving the quality of the preconditioner in a significant number of cases; see [11], [14], [25].

Denote by $\boldsymbol{a}_i^T$ the ith row of $\boldsymbol{A}$. The $\boldsymbol{A}$-orthogonalization process for computing $\boldsymbol{Z}$ and $\boldsymbol{D}$ from $\boldsymbol{A}$ can be written in two mathematically equivalent ways. We call these the *left-looking* and the *right-looking* algorithms, respectively.

Algorithm 2.1 Left-looking $\boldsymbol{A}$-orthogonalization algorithm:

Let $\boldsymbol{z}_1^{(0)} = \boldsymbol{e}_1; \quad p_1^{(0)} = a_{11}$
For $i = 2, \ldots, n$
  $\boldsymbol{z}_i^{(0)} = \boldsymbol{e}_i$
  For $j = 1, \ldots, i-1$
    $p_i^{(j-1)} := \boldsymbol{a}_j^T \boldsymbol{z}_i^{(j-1)}$
    $\boldsymbol{z}_i^{(j)} := \boldsymbol{z}_i^{(j-1)} - \left(\dfrac{p_i^{(j-1)}}{p_j^{(j-1)}} \right) \boldsymbol{z}_j^{(j-1)}$
  End
  $p_i^{(i-1)} := \boldsymbol{a}_i^T \boldsymbol{z}_i^{(i-1)}$
End

Algorithm 2.2 Right-looking $\boldsymbol{A}$-orthogonalization algorithm:

$$
\begin{array}{l}
\textit{Let } \boldsymbol{z}_i^{(0)} = \boldsymbol{e}_i \quad (1 \le i \le n) \\
\textit{For } i = 1, 2, \ldots, n \\
\quad \textit{For } j = i, i+1, \ldots, n \\
\quad\quad p_j^{(i-1)} := \boldsymbol{a}_i^T \boldsymbol{z}_j^{(i-1)} \\
\quad \textit{End} \\
\quad \textit{If } i < n \textit{ Then} \\
\quad\quad \textit{For } j = i+1, \ldots, n \\
\quad\quad\quad \boldsymbol{z}_j^{(i)} := \boldsymbol{z}_j^{(i-1)} - \left(\dfrac{p_j^{(i-1)}}{p_i^{(i-1)}} \right) \boldsymbol{z}_i^{(i-1)} \\
\quad\quad \textit{End} \\
\quad \textit{End If} \\
\textit{End}
\end{array}
$$

Both of these algorithms can be used to construct an AINV preconditioner by dropping entries below a prescribed drop tolerance in the $\boldsymbol{z}$-vectors (corresponding to entries above the main diagonal of $\boldsymbol{Z}$) as they are computed. We distinguish between two kinds of dropping criteria, absolute and relative. In absolute dropping, a new entry in the $\boldsymbol{Z}$ matrix is dropped if it is less than a prescribed drop tolerance $\psi > 0$ times the largest entry in $\boldsymbol{A}$. With relative dropping, an entry is dropped if it is less than ψ times the infinity norm of the current row of $\boldsymbol{A}$. In either case, the dropping is applied after each update step for $\boldsymbol{z}_i^{(j)}$ in the left-looking algorithm, and after each update step for $\boldsymbol{z}_j^{(i)}$ in the right-looking algorithm. The two variants (left- and right-looking) of AINV are not computationally equivalent. The numerical values of the entries in the resulting preconditioners are generally not identical, and the different implementations result in different timings for constructing the preconditioner. In our experience, the right-looking algorithm is usually faster on the relatively dense matrices arising from finite element analysis.

It is worth stressing that in order to have an efficient algorithm, the AINV process must be carefully implemented. Not only the sparsity in the rows of $\boldsymbol{A}$ and in the computed columns of $\boldsymbol{Z}$ must be exploited, but it is also important to make sure that of the inner products involving the rows of $\boldsymbol{A}$, only those that are structurally nonzero are actually computed; see, e.g., [15]. Then, for typical values of ψ and assuming an even distribution of nonzero entries in the incomplete inverse factor, all but a few of the innermost update loops in j can be skipped for each i, and the cost of computing the preconditioner scales linearly in the dimension n of the problem.

A crucial issue is that of the stability of the AINV process. For the preconditioner to be positive definite it is necessary that the pivots p_i be positive. In the absence of dropping, one has

$$p_i = \boldsymbol{a}_i^T \boldsymbol{z}_i = \boldsymbol{z}_i^T \boldsymbol{A} \boldsymbol{z}_i > 0,$$

since $\boldsymbol{A}$ is positive definite and $\boldsymbol{z}_i \neq \boldsymbol{0}$. However, when dropping is used the pivots may become non-positive, forcing the algorithm to terminate. This is referred to as a *breakdown*. In [10], it was proved that the AINV process cannot suffer a breakdown if $\boldsymbol{A}$ is an H-matrix (see [5] for a definition of H-matrix), similar to incomplete Cholesky. If $\boldsymbol{A}$ is not an H-matrix, breakdowns are possible; unfortunately, stiffness matrices arising in finite element analysis are not H-matrices in general, and breakdowns are a common occurrence. For instance, in [8] we found that thin shell problems lead to massive breakdowns in the AINV algorithm. (Some incomplete factorization algorithms are also susceptible to break down on these problems.) Several diagonal compensation strategies have been proposed to deal with pivot breakdowns in the context of incomplete Cholesky preconditioning. Of these, the most popular in the structural engineering community is the Ajiz-Jennings variant of incomplete Cholesky [1]. These techniques can be readily extended to the case of factorized approximate inverses like AINV; however, the resulting preconditioner is often of poor quality. Fortunately, it is possible to prevent breakdowns without the need for any diagonal corrections, simply by formulating the $\boldsymbol{A}$-orthogonalization algorithm in a slightly different manner. This reformulation was recently developed, independently, in [7] and [30]; see also [13] for a somewhat different approach. We refer to this algorithm as the SAINV (for stabilized AINV) algorithm. In passing, we also note that there exist other algorithms for constructing factorized approximate inverse preconditioners for positive definite matrices which are guaranteed to be breakdown-free. These include the FSAI preconditioner [32] and the bordering algorithm described in [42].

SAINV is based on a reformulation of the $\boldsymbol{A}$-orthogonalization process, presented below, which in exact arithmetic and in the absence of dropping is mathematically equivalent to the previous ones. For brevity, we describe the right-looking variant only.

Algorithm 2.3 Right-looking $\boldsymbol{A}$-orthogonalization algorithm, version 2:

Let $\boldsymbol{z}_i^{(0)} = \boldsymbol{e}_i \quad (1 \le i \le n)$
For $i = 1, 2, \ldots, n$
  $\boldsymbol{v}_i = \boldsymbol{A}\boldsymbol{z}_i^{(i-1)}$
  For $j = i, i+1, \ldots, n$
    $p_j^{(i-1)} := \boldsymbol{v}_i^T \boldsymbol{z}_j^{(i-1)}$
  End
  If $i < n$ *Then*
    For $j = i+1, \ldots, n$
      $\boldsymbol{z}_j^{(i)} := \boldsymbol{z}_j^{(i-1)} - \left(\frac{p_j^{(i-1)}}{p_i^{(i-1)}}\right) \boldsymbol{z}_i^{(i-1)}$
    End
  End If
End

It is easy to see [7] that when dropping in the $\boldsymbol{z}$-vectors is used in this algorithm, the pivots $p_i = p_i^{(i-1)}$ remain positive, and no breakdown is possible. In principle, breakdowns could still be possible because of round-off errors, but this is not likely to happen and we have never observed such an occurrence in actual computations. A more real danger is that of positive but very small pivots, which could potentially cause overflow and strong fill-in in the $\boldsymbol{z}$-vectors. In this case, some safeguarding might be necessary, such as pivot shifts and imposing a limit on the number of nonzeros allowed in a $\boldsymbol{z}$-vector; however, we have not met any situation in actual computation where these safeguards were needed. We found SAINV to be a robust procedure in practice. With this preconditioner, we were able to solve a number of the problems that we could not solve with the standard AINV preconditioner; see the numerical tests in [7].

The price to pay for this robustness is a somewhat higher cost of computing the preconditioner. This increase, however, is usually modest, and the cost of computing the preconditioner remains reasonable—typically a linear function of the problem dimension n; see [7]. This is especially true if the preconditioner can be reused over different solves, for in this case the set-up cost would be easily amortized.

Although the stabilized approximate inverse algorithm is guaranteed to avoid breakdowns, the quality of the preconditioning as measured by the convergence rates of the conjugate gradient iteration may still be unsatisfactory, particularly for very ill-conditioned problems. The situation can sometimes be improved using block preconditioning, as described in the next section. Diagonal and block diagonal scalings are also recommended. If $\boldsymbol{D}$ denotes the main diagonal of $\boldsymbol{A}$, then the preconditioning is applied to the symmetrically (Jacobi) scaled matrix $\boldsymbol{D}^{-\frac{1}{2}}\boldsymbol{A}\boldsymbol{D}^{-\frac{1}{2}}$. When the matrix has a natural block partitioning, a symmetric block Jacobi scaling is often beneficial. We will say more on scaling in section 4.

3 Block algorithms

A standard technique to improve performance in dense matrix computations is to use blocking; see, e.g., [22]. By partitioning the matrices and vectors into blocks of suitable size (which usually depends on the target architecture) and by making such blocks the elementary entities on which the computations are performed, high-level BLAS can be used for cache efficiency on current architectures with a hierarchical memory structure. As a result of such fine-tuning, computational rates near the theoretical peak are possible for many dense linear algebra calculations.

In contrast, computations involving sparse matrices are much more difficult to optimize, particularly when the matrices are irregularly populated, and computational rates are typically only a small fraction of the theoretical peak. This is largely caused by the presence of indirect addressing in the innermost loops, e.g., in multiplying a sparse matrix by a dense vector. The situation can sometimes be improved by extending the use of blocking to sparse matrices.

Block iterative methods have been popular for many years in the solution of linear systems arising from the discretization of partial differential equations on structured grids; see, e.g., [5] or [38]. Here the blocks arise from some natural partitioning of the problem (grid lines, planes, or subdomains) and they are usually large and sparse. For structured grid problems in 2D, efficient band solvers can be used to invert these blocks, leading to good performance in many cases.

A more interesting case for us is when the blocks are small and dense, as is the case when several variables are associated with a grid point, as with systems of partial differential equations. In this case high-level BLAS can be used as computational kernels and indirect addressing can be removed from the innermost loops, e.g., in the execution of matrix-vector products. For cache-based architectures, this leads to fairly good performance.

This block structure can either be naturally present in the matrix, or it must be imposed. As already mentioned, matrices with a natural block form often arise when the finite element method is used to discretize a partial differential equation or a system of them. In other cases there may be no natural block structure for us to exploit. However, we may still be able to use block algorithms by imposing a suitable blocking on the matrix, for instance by reordering the matrix using row and column permutations. A natural goal of the permutation should be to result in a block partitioned matrix with fairly dense blocks. To this end, any band/profile minimizing heuristic like reverse Cuthill-McKee (RCM) can be used; see [23] or [42]. Another possibility is to use the PABLO algorithm and its variants; see [41], [18]. These schemes result in block partitioned forms in which the dense blocks tend to be clustered around the main diagonal of the matrix. In our case we opted for a block construction which provides a general distribution of the blocks in the matrix. The algorithm to construct the blocks is based on a graph compression procedure due to Ashcraft [4]. We will give now a brief overview of this approach.

The block construction is based on the properties of the underlying graph. The aim of the graph compression is to find *cliques* in the undirected graph of the matrix. Consider the graph $G = (V, E)$ of the symmetric matrix $\boldsymbol{A} = (a_{ij})$. $V = \{1, \ldots, |V|\}$ is the set of vertices that correspond to rows and columns of the matrix. $E \subseteq V \times V$ is the set of its edges, where $(i, j) \in E$ if and only if the entry a_{ij} is nonzero. Then the adjacency set of a vertex v is

$$\mathrm{adj}(v) = \{u | (v, u) \in E\}.$$

Vertices $u \in V$ and $v \in V$ are adjacent if and only if $v \in \mathrm{adj}(u)$ (this is the same condition as $u \in \mathrm{adj}(v)$.) A *clique* in G is a set of vertices which are all mutually adjacent.

The task of finding blocks as large and dense as possible in a matrix $\boldsymbol{A}$ is equivalent to that of finding all cliques which are maximal with respect to inclusion in the graph G of the matrix $\boldsymbol{A}$. Of course, this task could be accomplished by brute force, comparing the adjacency sets of the vertices in G. However, this would result in an unnecessarily time-consuming procedure. In the paper [4], some basic enhancements of this procedure are described. Before actual computation and comparison of the adjacency sets take place, a preprocessing phase is performed. The candidates for inclusion into actual blocks are tested with respect to two easily computed additional quantities: vertex degree (size of an adjacency set) and vertex checksum (a simple function of adjacent vertices). These tests enable one to compute and compare adjacency sets in a few cases only. This contributes to the low cost of the graph compression routine.

The graph compression procedure was proposed as a tool for improving the performance of sparse direct solvers based on triangular factorizations of the system matrix. In this context, the compression contributes in a crucial way to reducing the symbolic overhead associated with the decomposition (reordering, elimination tree manipulation, symbolic factorization, etc.) As for the numeric phase of the factorization, it represents a useful complement to another important blocking strategy used in modern direct solvers, namely, the supernodal technique [35]. Note that the supernodal strategy is actually based on the structure of the triangular factors of $\boldsymbol{A}$, which are usually much less sparse than the system matrix itself. Because we are interested in block iterative algorithms, we focus on block strategies based on the sparsity of the original matrix only, like graph compression.

3.1 Block AINV preconditioning

Although the right-looking AINV algorithm is often faster, in this paper we restrict our attention to the left-looking block AINV scheme, which is easier to implement.

Assume the system matrix has been partitioned in the following block form:

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} & \cdots & \boldsymbol{A}_{1N} \\ \boldsymbol{A}_{21} & \boldsymbol{A}_{22} & \cdots & \boldsymbol{A}_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{A}_{N1} & \boldsymbol{A}_{N2} & \cdots & \boldsymbol{A}_{NN} \end{bmatrix}.$$

Here $\boldsymbol{A}_{ij}$ has order $n_i \times n_j$, where $1 \leq n_i \leq n$. N is the block dimension of the matrix, n is its dimension. Denote $\sum_{j<i} n_j$ by m_i (it is the offset of the ith block). Also, denote the block rows of $\boldsymbol{A}$ by $\boldsymbol{A}_i^T$, $i = 1, \dots, N$. That is,

$$\boldsymbol{A}_i^T = [\boldsymbol{A}_{i1}, \dots, \boldsymbol{A}_{iN}].$$

Note that the diagonal blocks $\boldsymbol{A}_{ii}$ are square symmetric positive definite matrices, and that $\boldsymbol{A}_{ji} = \boldsymbol{A}_{ij}^T$ for $i \neq j$.

The block AINV algorithm computes block partitioned $\boldsymbol{Z}$ and $\boldsymbol{D}$ directly from $\boldsymbol{A}$ based on a block $\boldsymbol{A}$-orthogonalization process applied to blocks of columns of the identity matrix. Note that $\boldsymbol{D}$ is now block diagonal.

Let $\boldsymbol{E}_i$ denote the $n \times n_i$ matrix with all zero rows except for rows $m_i + 1$ through $m_i + n_i$ which correspond to the rows of the $n_i \times n_i$ identity matrix $\boldsymbol{I}_{n_i}$. The block $\boldsymbol{A}$-orthogonalization procedure can be written as follows.

Algorithm 3.1 Block $\boldsymbol{A}$-orthogonalization algorithm:

Let $\boldsymbol{Z}_i^{(0)} = \boldsymbol{E}_i \ \ (1 \leq i \leq N);\ \ \boldsymbol{P}_1^{(0)} = \boldsymbol{A}_{11}$
For $i = 2, \ldots, N$ *do*
  For $j = 1, \ldots, i-1$ *do*
    $\boldsymbol{P}_i^{(j-1)} := \boldsymbol{A}_j^T \boldsymbol{Z}_i^{(j-1)}$
    $\boldsymbol{Z}_i^{(j)} = \boldsymbol{Z}_i^{(j-1)} - \boldsymbol{Z}_j^{(j-1)} [\boldsymbol{P}_j^{(j-1)}]^{-1} \boldsymbol{P}_i^{(j-1)}$
  End do
  $\boldsymbol{P}_i^{(i-1)} = \boldsymbol{A}_i^T \boldsymbol{Z}_i^{(i-1)}$
  End do
End do
Let $\boldsymbol{Z}_i := \boldsymbol{Z}_i^{(i-1)}$ *and* $\boldsymbol{D}_i := \boldsymbol{P}_i^{(i-1)}$, *for* $1 \leq i \leq N$. *Return*
$\boldsymbol{Z} = [\boldsymbol{Z}_1,\ \boldsymbol{Z}_2, \ldots, \boldsymbol{Z}_N]$ *and* $\boldsymbol{D} = \mathrm{diag}(\boldsymbol{D}_1, \boldsymbol{D}_2, \ldots, \boldsymbol{D}_N)$.

The block AINV preconditioner $\boldsymbol{M} = \boldsymbol{Z}\boldsymbol{D}^{-1}\boldsymbol{Z}^T \approx \boldsymbol{A}^{-1}$ is computed by applying this block generalized Gram-Schmidt process incompletely. Blocks with small norms (corresponding to positions above the block diagonal part of the triangular matrix $\boldsymbol{Z}$) are dropped to preserve sparsity after the block updates for $\boldsymbol{Z}_i^{(j)}$. In our implementation, we used the infinity norm to gauge the size of the matrix blocks, but other choices are possible; see [15].

The inverses of the pivot blocks appearing in the update step for $\boldsymbol{Z}_i^{(j)}$ above are computed by means of a full triangular factorization. Notice that because of dropping, the pivot blocks in the incomplete process may not be positive definite, or even symmetric. In our implementation, we adopted the following safeguards. First, an $\boldsymbol{LU}$ factorization of $\boldsymbol{P}_j^{(j-1)}$ is computed (without pivoting). Then, the pivot block is symmetrized by replacing it, if necessary, with $\boldsymbol{L}\boldsymbol{L}^T$, where $\boldsymbol{L}$ is the lower triangular factor of $\boldsymbol{P}_j^{(j-1)}$. Pivot modifications may be used to enforce positive definiteness if needed. In practice, we found that such modifications are rarely needed, and the block AINV procedure is much more robust than the point AINV one.

Nevertheless, stabilization can and should be used to obtain a fully reliable algorithm. The block stabilized AINV algorithm is obtained from Algorithm 3.1 by replacing the expressions for $\boldsymbol{P}_i^{(j-1)}$ and $\boldsymbol{P}_i^{(i-1)}$ with $[\boldsymbol{Z}_j^{(j-1)}]^T \boldsymbol{A} \boldsymbol{Z}_i^{(j-1)}$ and $[\boldsymbol{Z}_i^{(i-1)}]^T \boldsymbol{A} \boldsymbol{Z}_i^{(i-1)}$, respectively. The pivot blocks are now guaranteed to be symmetric and positive definite, at the price of a somewhat higher construction cost.

Notice that the block AINV algorithms are rich in dense matrix-matrix operations, hence BLAS-3 kernels can be used to attain high performance.

In practice, we found that the performance of the block AINV schemes can be significantly enhanced by explicitly applying a symmetric block Jacobi scaling to the coefficient matrix $\boldsymbol{A}$ prior to performing the block AINV process. This preprocessing is based on the computation of the Cholesky factorizations $\boldsymbol{A}_{ii} = \boldsymbol{L}_i \boldsymbol{L}_i^T$ for $1 \leq i \leq N$, followed by the explicit formation of the block scaled matrix

$$\hat{\boldsymbol{A}} = \boldsymbol{G}^{-1} \boldsymbol{A} \boldsymbol{G}^{-T}$$

where $\boldsymbol{G} = \mathrm{diag}(\boldsymbol{L}_1, \ldots, \boldsymbol{L}_N)$. The additional costs incurred by this scaling are usually small compared to the total solution costs.

The use of blocking in the context of the AINV algorithm is not entirely new. Challacombe [16], [17] has used a natural blocking of the problem to achieve high performance when applying the AINV transformation in quantum chemistry applications; however, the approximate inverse itself was computed using the point version of AINV (Algorithm 2.1). Bridson and Tang [15] developed a block variant of AINV which is also guaranteed not to break down; our approach differs from theirs in the blocking strategy, in the $\boldsymbol{A}$-orthogonalization scheme used and in the implementation, and appears to result in more robust preconditioning. Indeed, we are able to solve some problems for which they report failure in [15]. A different kind of block approximate inverse preconditioner, based on Frobenius norm minimization, has been recently described in [6]. Other recent references include [21] and [40] for the use of blocking within incomplete factorization preconditioning. It should be noticed that the effect of blocking is quite different for different preconditioners. For instance, for the case of the SPAI preconditioner, blocking has the effect of reducing the time to compute the preconditioner compared to the unblocked version, but at the cost of slower convergence rates [6]. As we will see in the next section, this is not usually the case for AINV. In all cases, however, blocking results in improved cache-efficiency on hierarchical memory computers.

4 Numerical experiments

In this section several numerical examples are presented. Special emphasis is given to problems arising in the modeling of thin shells, which are known to be hard for preconditioned iterative methods. Different parametric studies have been carried out. The effect of the regularizing parameter in the Hughes-Brezzi formulation on the quality of the preconditioner is investigated. The effects of point and block scaling and equation ordering are discussed. In addition, results for some 3-D solid mechanics problems are given.

The computations shown have been performed with a single processor on an SGI Origin 2000 (R12k processor) and on a Compaq AlphaServer GS140 (EV6 processor) at the Center for Scientific Computing, Espoo, Finland. In all the test runs, we used the zero vector as the initial guess and we stopped the iteration when the Euclidean norm of the initial residual had been reduced by five or ten orders of magnitude depending on whether realistic or artificial right-hand-side vectors are used, respectively. Reported CPU timings are in seconds and correspond to the SGI Origin 2000 computations if not otherwise stated.

We also present results obtained with some incomplete Cholesky-type preconditioners, e.g., pointwise and block versions of symmetric ILUT [42] (denoted here by ICT) and of the robust Ajiz-Jennings incomplete factorization [1].

4.1 Shell examples

In this subsection, we report on results of computations for some commonly used shell test problems. In these models, the drilling rotation is accomplished by using the Hughes-Brezzi formulation [26] including an Allman-type displacement field [2]. The plate bending part of the element is based on the stabilized MITC theory [36] or the discrete-Kirchhoff concept. In the MITC formulation the stabilization parameter has been 0.4. The value of the regularizing penalty parameter γ used in the formulation of Hughes and Brezzi, unless otherwise stated, has been $\gamma = 10^{-3}G$, where G is the shear modulus. This value affects the condition number of the stiffness matrix and thus the convergence of the PCG iteration. The effect of the parameter γ on the spectral condition number as well as the convergence of the PCG iteration is demonstrated in [33] and will be further studied in the following examples.

4.1.1 On symmetric diagonal scaling and the effect of ordering

Diagonal scalings have a strong influence on the behavior of the approximate inverse preconditioning with (S)AINV. The reason is that the decay in the inverse factors can be significantly affected by symmetric Jacobi scaling. To illustrate this effect for the pointwise SAINV approach, we consider the analysis of a simple flat plate. For flat plates there is no coupling between the membrane and bending deformation and they depend on different nodal unknowns, thus by using a specific numbering of the unknowns, the global stiffness matrix can be assembled in a block form

$$\begin{bmatrix} \boldsymbol{A}_m & \boldsymbol{O} \\ \boldsymbol{O} & \boldsymbol{A}_b \end{bmatrix},$$

where $\boldsymbol{A}_m$ and $\boldsymbol{A}_b$ are the global stiffness matrices associated with the membrane and the bending action, respectively.

The plate is discretized by an uniform quadrilateral 10×10 element mesh with 4-node shell elements with drilling degrees of freedom, resulting in 610 unknowns.

In Figs. 1 and 2 the sparsity patterns of the stabilized AINV preconditioner (incomplete inverse factor $\boldsymbol{Z}^T$) with and without Jacobi scaling are shown. The drop tolerance is chosen in such a way that the density of the preconditioners in Figs. 1 and 2 will be approximately the same. It is clearly seen that without Jacobi scaling the dropping strategy will result in a preconditioner which almost totally ignores the bending deformational part of the system. In constructing these figures the absolute dropping criterion was used; however, similar patterns were also obtained with the relative criterion.

In curved shells the coupling between membrane and bending deformations complicates the construction of optimal dropping strategies. However, the numerical experiments strongly suggest to apply Jacobi scaling when pointwise dropping is used in the preconditioner construction.[1]

To illustrate the effect of point and block scalings, we show some results from an analysis of a pinched cylidrical shell. In [8] the basic AINV strategy was tested for such matrices with modest results. As already mentioned, the main reason for the observed failures of the AINV-preconditioned conjugate gradient iteration in [8] was massive pivot breakdowns. If a static AINV preconditioner with the same sparsity pattern as that of $\boldsymbol{A}$, denoted by AINV(0), is applied to an *a priori* shifted matrix $\boldsymbol{A} + \alpha \text{diag}(\boldsymbol{A})$, breakdowns can be avoided and the PCG iteration converges. However, the number of iterations can be rather high; see Table 1, where the matrix `S3RMQ4M1` is used as a test case. This matrix can be obtained from the Matrix Market [39]. Furthermore, the use of a drop tolerance (instead of a static sparsity pattern) often resulted in instabilities and no PCG convergence [8].

Considerable improvements can be achieved with the stabilized pointwise algorithm SAINV with Jacobi scaling, as was shown in [7]. However, even better results are obtained with the block AINV method, particularly with

[1] The dropping strategy defined by Ajiz and Jennings [1] will implicitly do the same.

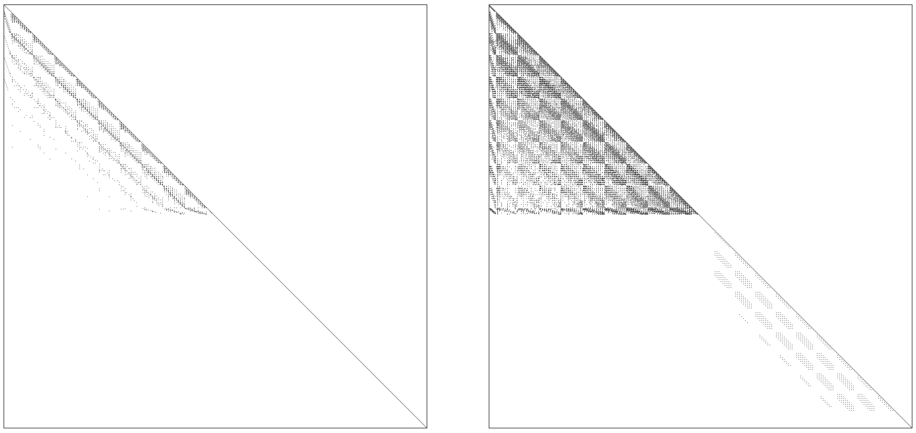

(a) (b)

Figure 1: Sparsity patterns of the SAINV preconditioner with two drop tolerances: (a) $\psi = 0.2 \max |a_{ij}|, \quad nz(\boldsymbol{Z}^T) = 3513$; (b) $\psi = 0.02 \max |a_{ij}|, \quad nz(\boldsymbol{Z}^T) = 20355$.

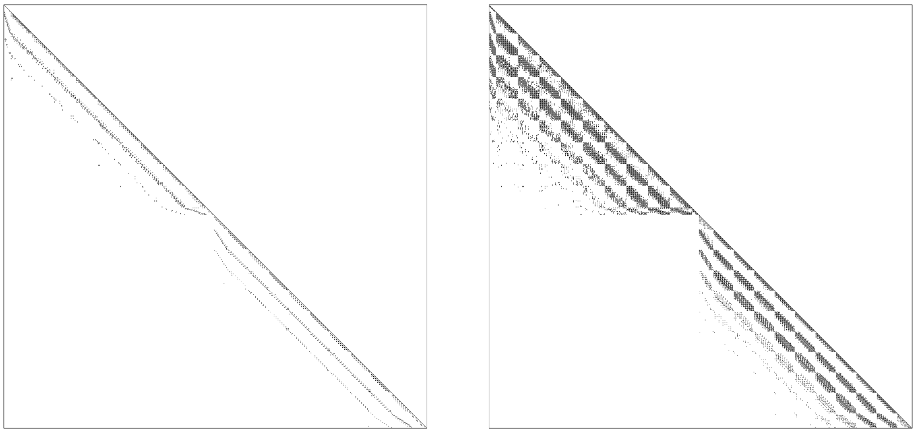

(a) (b)

Figure 2: Sparsity patterns of the SAINV preconditioner with two drop tolerances applied to the Jacobi scaled matrix: (a) $\psi = 0.2, \quad nz(\boldsymbol{Z}^T) = 3521$; (b) $\psi = 0.02, \quad nz(\boldsymbol{Z}^T) = 21623$.

Table 1: Matrix S3RMQ4M1, standard AINV algorithm, MMD nodal ordering.

preconditioner	scaling	ψ	ρ	iterations	notes
static AINV(0)	-	-	1.00	1834	$\alpha = 0.01$
block AINV	-	0.01	0.90	956	
	-	0.005	2.89	440	
	-	0.0025	4.82	294	
block AINV	Jacobi	0.5	1.18	873	
	Jacobi	0.25	2.07	568	
	Jacobi	0.1	4.67	298	
block AINV	block-Jacobi	0.1	1.21	424	
	block-Jacobi	0.05	1.73	331	
	block-Jacobi	0.02	3.66	186	

block Jacobi scaling. This is shown in Table 1, where ρ denotes the relative density of the preconditioner (i.e., the ratio of the number of nonzeros in the preconditioner and the number of nonzeros in the coefficient matrix). The point Jacobi scaling has little effect on the performance of the block AINV preconditioner, but the block scaling is very beneficial.

In Fig. 3 we show the effect of block Jacobi scaling on the performance of the block SAINV algorithm for another thin shell problem, S3RMT3M3, also available from the Matrix Market. This is an unstructured matrix corresponding to a graded mesh with 1666 triangles, 5357 dof's, and a spectral condition number $\text{cond}_2(\boldsymbol{A}) \approx 2.4 \cdot 10^{10}$. Again, it is clear that block Jacobi scaling significantly improves the quality of the preconditioning.

The results above were obtained with a minimum degree ordering of the nodes. It is well known [24] that the numbering of unknowns may have significant influence on the performance of preconditioned iterations. For point versions of the AINV preconditioner, this has been investigated in [14] and [11]. Here we consider nodal orderings, corresponding to orderings of the block partitioned matrix or, equivalently, of the compressed graph. This is standard practice in structural mechanics. All the unknowns in a node are numbered consecutively before entering the next node. How to order the nodes depends on which kind of preconditioning will be used. Implicit incomplete factorization preconditioners tend to benefit from reorderings which try to minimize the RMS bandwidth, like RCM. On the other hand,

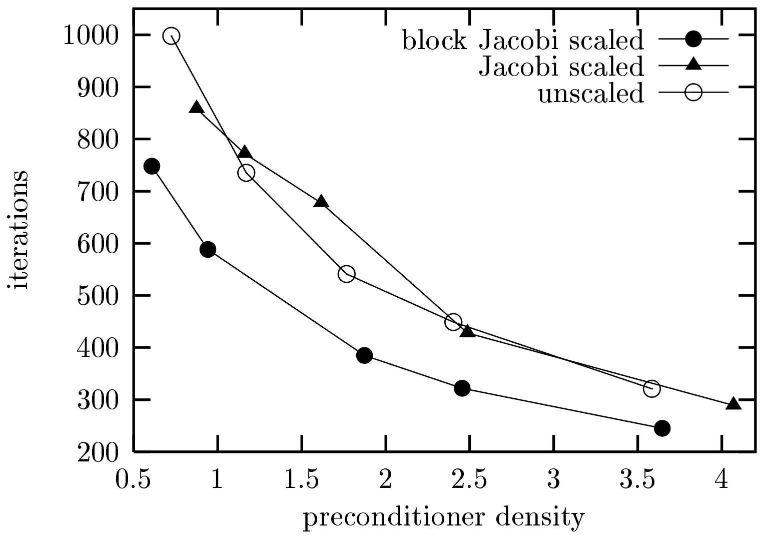

Figure 3: Test case S3RMT3M3, effect of scaling for the block stabilized AINV.

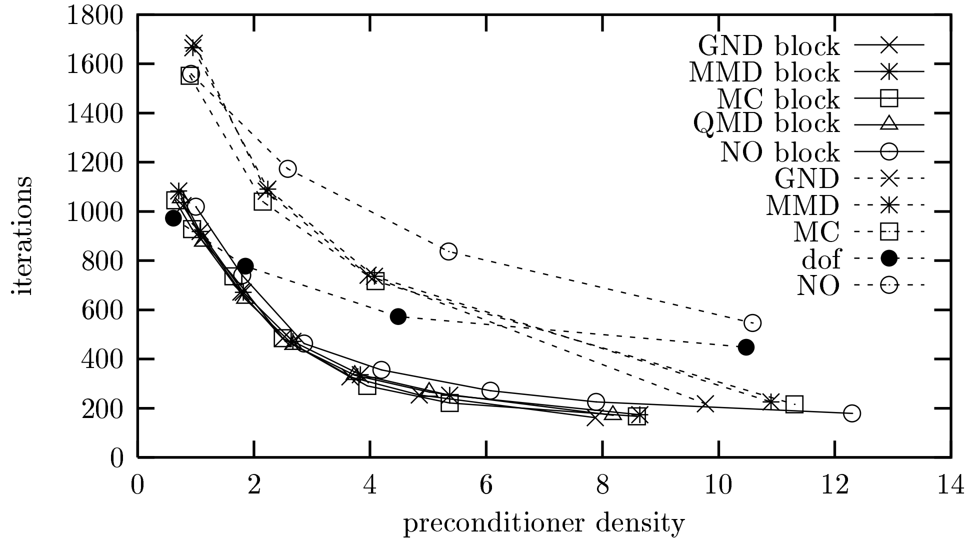

Figure 4: Test case S3RMT3M3, comparison with pointwise stabilized AINV and the block stabilized AINV, different orderings.

factorized approximate inverse preconditioners do poorly with such orderings, and in [14] and [11] it was found that orderings like minimum degree or nested dissection, which result in short and wide elimination trees and sparse inverse factors, should be preferred.

Another possibility is to adopt an ordering by degrees of freedom (dof-wise), where unknowns of similar type are ordered sequentially sweeping all the nodes in the mesh.

In Fig. 4 we show the iteration count as a function of the preconditioner density for two types of preconditioners (block SAINV and point SAINV) and different orderings for the shell problem `S3RMT3M3`. The ordering strategies are multiple minimum degree (MMD), generalized nested dissection (GND), multicoloring (MC) and quotient minimum degree (QMD), besides the original ordering (NO). The block SAINV algorithm is scarcely sensitive to the ordering of the nodes. This is unlike point SAINV, for which we can see that the ordering makes a difference. Here, for reasonable preconditioner densities, the dof-wise ordering clearly outperforms the other orderings; however, the performance is not as good as with block SAINV. Experiments with other thin shell problems indicate that on average, MMD nodal ordering is slightly superior to the other orderings when block (S)AINV is being used.

4.1.2 On the value of the regularizing parameter in the Hughes-Brezzi formulation

In-plane rotational degrees of freedom, "drilling degrees of freedom", are particularly convenient in the analysis of shells. Typical shell elements have three translational and two rotational degrees of freedom at a node. This results in many difficulties of model construction and of programming and numerical ill-conditioning for certain types of element assemblages. Thus the presence of all three rotations at a node is advantageous from a practical point of view.

Early attempts to construct membrane elements with drill rotations were unsuccessful. Hughes and Brezzi [26] presented a simple variational formulation, which employs an independent rotation field and is also stable in the discrete case. They also proved that elements based on their formulation are convergent for all standard interpolations including equal order interpolation for displacements and rotation. Numerical experiments reported by Hughes *et al.* [27] confirm the *a priori* theoretical convergence estimates. They compared linear and quadratic triangular and quadrilateral elements and also bilinear elements with incompatible modes. Incompatible modes significantly improve the coarse mesh accuracy of the bilinear element. On the other hand, the static condensation needed to eliminate the nodeless generalized displacements is awkward, especially in nonlinear problems.

Ibrahimbegović *et al.* [29] amended the displacement interpolation of a four node quadrilateral element by the Allman-type [2] quadratic modes in order to improve the coarse mesh accuracy of the element. This considerably improves the bending behavior of the element. They also added a hierarchical bubble interpolation mode to the displacement field.

The bilinear form of the variational equation of the Hughes-Brezzi formulation is

$$\int_\Omega \delta\boldsymbol{\varepsilon} : \boldsymbol{C} : \boldsymbol{\varepsilon} d\Omega + \gamma \int_\Omega (\mathrm{skew}\nabla\delta\boldsymbol{u} - \delta\theta)(\mathrm{skew}\nabla\boldsymbol{u} - \theta) d\Omega \tag{4}$$

where $\boldsymbol{\varepsilon}$ is the strain tensor, i.e., the symmetric part of the displacement gradient, $\boldsymbol{\varepsilon} = \mathrm{symm}\,\nabla\boldsymbol{u}$, θ is a skew-symmetric tensor representing the in-plane rotation and δ denotes the variation. The fourth order tensor $\boldsymbol{C}$ contains the material parameters.

An appropriate value of the regularizing penalty parameter γ is chosen in accordance with the ellipticity condition. For the isotropic case the value $\gamma = G$ (shear modulus) seems to balance the terms in the estimate and thus seems reasonable [28]. The method is insensitive to the choice of the penalty parameter in the region $0 < \gamma \leq G$ if discretization accuracy is considered. Since the above formulation is extensively studied in the papers [26], [27] and [28], additional details are not repeated here.

However, little is known about the influence of the regularizing parameter on the conditioning of the stiffness matrix. In [33] the effect of the γ parameter on the condition number of a cylindrical shell modelled by a regular 30×30 quadrilateral element mesh (four noded stabilized MITC type elements) was determined as well as the convergence of the IC(0)-PCG iteration. In the following, similar results are reported for a cylindrical shell with an irregular mesh modelled by three node triangular elements. The effect of symmetric Jacobi scaling is also studied.

In Fig. 5 the effect of the γ parameter, symmetric Jacobi scaling and different drill-rotation formulations on the spectral condition number of the stiffness matrix are shown. The matrix corresponds to a pinched cylindrical shell discretized by means of an irregular mesh with 1666 triangular three-node stabilized MITC-type shell elements. Note that the unscaled case with $\gamma = 10^{-3}G$ and with Allman displacement field corresponds to the CYLSHELL matrix `S3RMT3M3` in the Matrix Market.

When iterative methods are used, the recommended value for the γ parameter lies in the interval $10^{-6} < \gamma/G < 10^{-2}$, see Fig 5 and Table 2, where

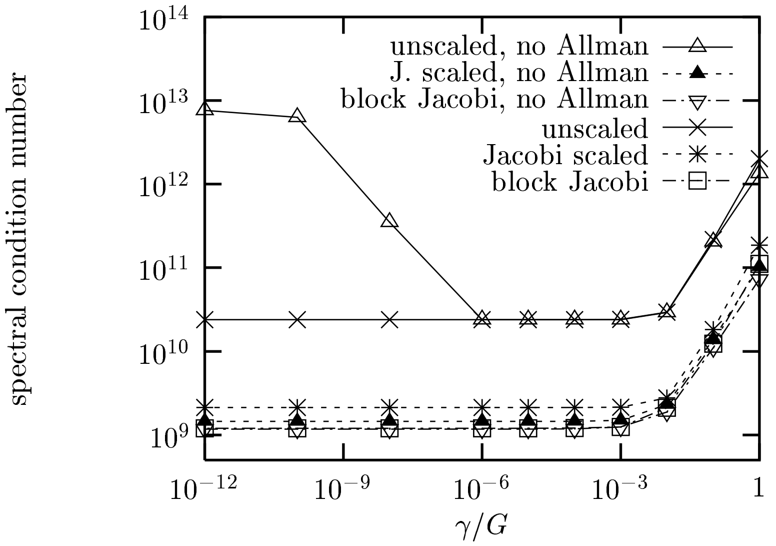

Figure 5: Effect of the γ parameter, symmetric Jacobi scaling and different drill-rotation formulations on the spectral condition number. Pinched cylindrical shell, irregular mesh with 1666 triangular three node stabilized MITC type shell elements, $n = 5357,\ N = 938$.

the number of SAINV-PCG iterations is shown. Jacobi scaling gives an advantageous effect in lowering the spectral condition number by two orders of magnitude. Block Jacobi scaling gives consistently the best performance of the PCG iteration, although the difference in the condition number in comparison to the pointwise Jacobi scaling is not significant.

If an Allman-type displacement field is not used, it is usually recommended that the γ parameter should not be too small, owing to the fact that the condition number grows again if $\gamma < 10^{-6}G$. However, Jacobi or block Jacobi scaling seems to inhibit the growth of the condition number for small γ values.

4.1.3 Performance in shell problems

The performance of several preconditioning techniques in the case of the CYLSHELL matrix `S3DKT3M2` [39] is shown in Table 3. It corresponds to a pinched cylindrical shell model ($t/R = 10^{-3}$) discretized by a uniform 150×100 mesh of triangular three node discrete Kirchhoff-type elements

Table 2: Effect of the regularizing parameter on the convergence of block SAINV preconditioned iterations (MMD nodal ordering). Irregular mesh with 1666 triangular three node stabilized MITC type elements. The drill-formulation includes Allman-type displacement field.

	block SAINV, unsc.			bl. SAINV, J. scaled			bl. SAINV, bl. J. scaled		
γ/G	ψ	ρ	iter	ψ	ρ	iter	ψ	ρ	iter
1	0.5	1.98	>10000	0.5	1.95	5905	0.25	1.52	3969
10^{-1}	0.5	1.75	1892	0.5	1.47	1499	0.25	1.02	1535
10^{-3}	0.5	3.12	379	0.5	1.34	672	0.25	0.76	689
10^{-5}	0.5	2.76	500	0.5	1.33	688	0.25	0.76	673

(90499 unknonws). It is also the most ill-conditioned of the CYLSHELL matrices, having a spectral condition number $\approx 3.63 \cdot 10^{11}$.

In this particular problem, the incomplete Cholesky-type preconditioners perform rather well, partly due to the narrow band of the system matrix. RCM nodal ordering is used for IC-type methods (RMS bandwidth 497, maximum bandwidth 609) and MMD ordering for AINV. Although the block ICT preconditioner gives the fastest execution in this single-processor experiment, it is not easy to find a parameter pair $(\psi, lfil)$ for which the decomposition exists without diagonal modifications, which usually destroy the convergence properties of the preconditioned iteration. (Here ψ denotes, as usual, the drop tolerance and $lfil$ the maximum number of additional nonzero blocks allowed in each block column of the incomplete factor.) This is true for almost all shell problems we have tried to solve with the ICT preconditioner. Therefore this method cannot be considered robust for thin shell problems.

For this particular problem, block SAINV preconditioning is not competitive with the block version of the reliable Ajiz-Jennings IC preconditioner. For $\psi = 0.1$, however, it is faster than block Jacobi, so it may be useful in a parallel implementation.

The efficiency of the block methods is also studied by recording the Mflop rates for a sequence of cylindrical shell problems starting from a 30×30 element mesh (matrix `S2RMT3M1`) having 5489 unknowns to a fine discretization by 180×180 element mesh with 194939 unknowns. Results are shown in Table 4. Here, B-A-J denotes the block Ajiz-Jennings incomplete Cholesky preconditioner. In all cases the number of nonzeros in the preconditioner is approximately the same as the number of nonzeros in the lower triangular

Table 3: Cylindrical shell S3DKT3M2, $n = 90499$, $N = 15251$.

	ψ	ρ	iter	P-time	I-time	notes
Jacobi	-		>20000	-	-	
block Jacobi	-		8149	-	959	
IC(2)	-	1.83	1428	4	910	$\alpha = 2 \cdot 10^{-3}$
block Ajiz-Jennings	0.05	1.27	987	2	277	
	0.01	1.82	672	4	236	
	0.002	2.90	386	7	196	
block ICT	0.1	0.75	3369	1	718	$lfil = 10$
	0.001	2.31	248	4	107	$lfil = 10$
block SAINV	0.2	0.77	4338	5	911	
	0.15	1.17	3783	7	960	
	0.1	2.44	1903	15	782	
	0.05	4.01	1472	28	867	
	0.02	8.44	932	83	1050	

part of the stiffness matrix, denoted by $nz(\boldsymbol{A})$.

The maximum performance of the SGI Origin 2000 R12k processor is 600 Mflops, thus for problems small enough to fit in the processor's external cache (8 Mb = one million double precision real numbers) we can achieve over 25% of the processor's maximum performance.[2] For bigger problems we can see a rather rapid decay of performance with respect to flop rates. Interestingly, the flop rate for the preconditioner construction phase is scarcely affected by the dimension of the problem. For block SAINV on the Origin, this was typically between 200 and 212 Mflops (between 260 and 310 Mflops on the AlphaServer), regardless of problem size.

A test case where block SAINV tends to outperform the incomplete Cholesky-type preconditioners, although it has a small bandwidth, is a simply supported T-beam problem from [34], where it has been used as a test example of interactive buckling phenomena. The structured mesh consist of 880 quadrilateral discrete-Kirchhoff type shell elements, resulting in 5563 unknowns. A strongly orthotropic layer is used to model the overlap between flange and web, having high bending rigidity in the transverse direction; see [34] for details. This leads to a highly ill-conditioned equation system, the spectral condition number of the stiffness matrix being $\approx 1.87 \cdot 10^{14}$.

Results are shown in Table 5. In this particular case the ICT precondi-

[2]For the Compaq AlphaServer GS140 EV6 processor, the external cache is 4 MB.

Table 4: Cylindrical shell, Mflop rates.

			Origin 2000†		AlphaServer GS140‡	
mesh	n	$nz(\boldsymbol{A})$	B-SAINV	B-A-J	B-SAINV	B-A-J
30×30*	5489	112505	173	128	217	180
40×40	9719	201605	160	128	134	110
50×50	15149	316505	139	109	101	90
60×60	21779	457205	101	92	81	78
70×70	29609	623705	83	80	76	75
80×80	38639	816005	77	73	75	70
90×90	49136	1034105	73	67	75	65
100×100	60299	1278005	71	67	74	62
140×140	118019	2511605	65	57	68	53
180×180	194939	4158005	65	62	67	53

† compiled with -mips4 -64 -r12k -OPT:Olimit=8000 -03 command
‡ compiled with -fkapargs='-o=5 -so=3' command
* case `S2RMT3M1`

tioner is completely useless. A stable preconditioner is obtained only with a very small drop tolerance and large $lfil$, resulting in nearly complete factorizations. The block Ajiz-Jennings preconditioner is stable, but it is generally inferior to block SAINV for reasonable amounts of fill.

4.2 Solid models

In this section some typical results with 3-D solid models are given. A standard displacement-based element formulation is used.

4.2.1 Surface mounted transistor

The first 3-D solid finite element model is from a thermal stress analysis of a surface mounted transistor, see Fig. 6. Only half of the transistor is modelled with a piece of a printing wiring board (PWB). The model consists of 1704 reduced triquadratic 20 node brick elements, resulting in 25710 unknowns. The number of nonzeros in the lower triangular part of the matrix is $nz(\boldsymbol{A}) = 1889447$.

Table 5: T-beam, $n = 5563$, $N = 945$.

block SAINV			
ψ	ρ	iter	
0.8	0.15	>10000	= block Jacobi
0.7	0.29	3669	
0.6	0.36	3689	
0.5	0.52	2851	
0.4	0.76	>10000	
0.3	1.11	4943	
0.25	1.27	4298	
0.24	1.39	4426	
0.23	1.43	4029	
0.225	1.47	731	
0.22	1.49	800	
0.2	1.66	699	
0.1	3.11	948	
0.05	4.57	550	
0.02	6.89	191	
0.01	9.13	73	
block Ajiz-Jennings IC			
0.1	1.00	>10000	
0.01	2.10	5403	
0.005	2.64	4405	
0.001	4.99	1221	
0.0005	5.46	840	
0.0002	5.61	308	

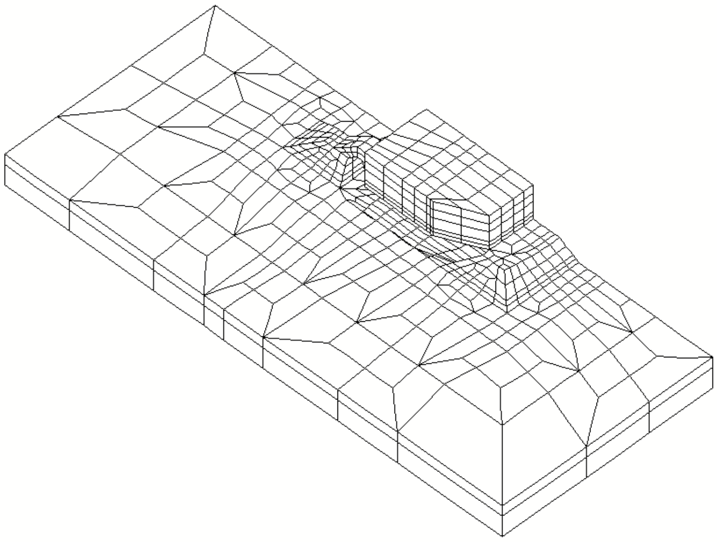

Figure 6: Surface mounted transistor, 1704 triquadratic elements, 8744 nodes.

Table 6: Surface mounted transistor, $n = 25710$, $nz = 1889447$, $N = 8744$.

preconditioner	ψ	ρ	iterations	P-time	I-time	P+I
block IC(0)	-	1.00	308	14	80	94
block Ajiz-Jennings	0.02	0.99	422	9	103	112
block Ajiz-Jennings	0.01	1.19	286	13	92	105
block SAINV	0.2	0.20	917	34	111	145
block SAINV	0.1	0.44	633	112	99	211
SAINV	0.03	0.43	379	26	99	125
SAINV	0.02	0.63	307	43	89	132
SAINV	0.01	1.17	218	96	85	181

In this problem the standard IC and AINV approaches need shifting to behave well, and only block IC(0) succeeds without shifts. The ICT preconditioner was unsuccessful; without shifting the matrix, it is impossible to find a working pair of parameters ψ and $lfil$. Therefore only robust preconditioners, i.e. the Ajiz-Jennings IC and SAINV approaches are used in the comparison.

There is no big difference between nodal reorderings like MMD, QMD and GND, in terms of preconditioner density and number of iterations. The computing times are also comparable, thus only results with MMD reordering are shown in Table 6.

The Gibbs-King reordering gave the smallest RMS bandwidth (1231) and it was used with incomplete factorization-based preconditioners.

It should be mentioned that the iteration CPU-times for the pointwise SAINV can be reduced if the matrix is stored in the block format although the preconditioner is formed in a pointwise fashion. When using the block format one matrix-vector multiply takes 0.095s while in the pointwise Compressed Storage by Columns (CSC) format the time is 0.18s, thus 25-36 seconds can be substracted from the iteration times of the pointwise SAINV results in Table 6. In this example the SAINV preconditioner results in similar convergence rates to the IC-based preconditioners although the preconditioner is only half as dense as the IC-type ones. Hence, SAINV gives the fastest iteration time of the compared methods even in this single-processor experiment. This makes SAINV quite attractive, especially in a parallel environment.

4.2.2 Engine head

In Fig. 7 a finite element model of an engine head is shown. Unstructured discretization in 151483 linear 4-node tetrahedral elements results in 143571 unknowns (number of nonzero elements in the lower triangular part of the stiffness matrix $nz(\boldsymbol{A}) = 2424822$). The loading used in the computations corresponds to a temperature increase by one degree centigrade.

This problem is an easy one to solve with preconditioned iterative methods. The standard AINV approach as well as IC preconditioners can be constructed without shifting. Nevertheless, the standard AINV performs worse than its stabilized variant, therefore such results are omitted. MMD reordering of nodes has been used in computations with the AINV algorithms, and the Gibbs-King profile reduction resulted in the smallest RMS bandwidth (2566), and it was used for the IC-type preconditioners.

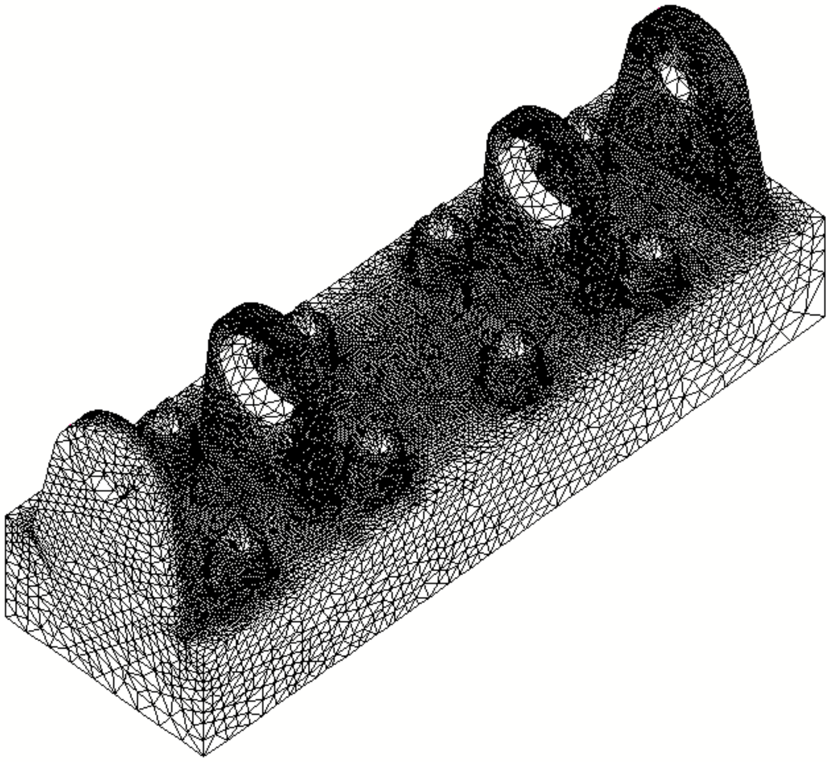

Figure 7: Engine head, 151483 linear tetrahedral elements, 48989 nodes.

Table 7: Engine head, $n = 143571, nz = 2424822, N = 47857$.

preconditioner	ψ	ρ	iterations	P-time	I-time	P+I	$lfil$
Right-looking SAINV preconditioners							
SAINV	0.1	0.29	568	15	226	241	
SAINV	0.05	0.63	427	26	209	235	
SAINV	0.025	1.24	375	51	236	287	
Left-looking SAINV preconditioners							
block SAINV	0.2	0.48	695	14	218	232	
block SAINV	0.1	1.02	485	33	202	235	
Incomplete factorization-based preconditioners							
IC(0)	-	1.00	312	1	162	163	
block IC(0)	-	1.00	320	3	148	151	
block A-J	0.02	1.38	275	7	129	136	
block A-J	0.01	1.69	209	10	109	119	
block ICT	0.05	0.81	560	3	206	209	10
block ICT	0.01	1.34	211	6	90	96	10

It can be seen from the results in Table 7 that block SAINV gives slightly slower convergence rates in comparison to the pointwise version.

One matrix-vector multiply takes 0.25s or 0.14s depending on whether the pointwise or block storage format is used, respectively. As explained in the previous example, 41 to 62 seconds could be saved from the pointwise SAINV iteration times if the matrix were stored in the block format.

4.2.3 Sensitivity to element aspect ratio

Current approaches for analyzing complex laminated structures with solid elements will usually result in element gometries where one dimension is small compared with the others. The convergence rate of the PCG iteration deteriorates with increasing aspect ratio.

An elastic block composed by three material layers and occupying the region (in cartesian coordinates) $0 < x, y < L$ and $0 < z < H$ is considered. The material interfaces are horizontal layers parallel to the xy-plane, and having positions $z = \frac{2}{5}H$ and $z = \frac{3}{5}H$. The stack models a ceramic (AlN) to metal (Ti) joint brazed together with Ag-Cu filler alloy. The constitutive parameters used for these materials are given in Fig. 8. A uniform $20\times20\times20$

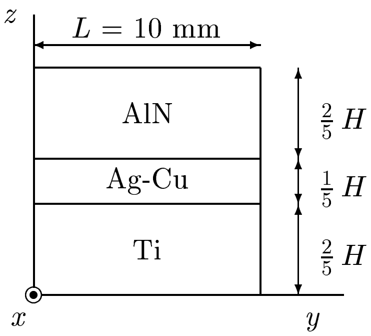

	Ti	Ag-Cu	AlN
E/GPa	118.6	94.4	315
ν	0.36	0.36	0.25
$10^6\alpha$	6.441	17.995	2.85

Figure 8: Three material solid block: geometry and material data.

element mesh with eight node trilinear brick elements is used. Loading is temperature change defined by $\Delta T = \Delta T_0 xyz/L^3$. Minimal constraints which prevent the rigid body motion are imposed.

Table 8: Trimetallic block

	SAINV			BSAINV		
L/H	ψ	ρ	iter	ψ	ρ	iter
1	0.025	0.48	185	0.1	0.46	244
10	0.025	0.75	493	0.1	0.82	745
100	0.025	0.71	3036	0.1	1.03	9987

Computations are carried out with three different aspect ratios L/H. As usual, the convergence of the preconditioned iterations deteriorates with increasing aspect ratio distortions. The block preconditioner suffers more severely from high aspect ratios than the pointwise one. Only if the element aspect ratios are in a feasible range, say $L/H < 10$, the performance of the block preconditioner is reasonable. This is exactly the opposite of what is observed in thin shell applications, where the block preconditioners are more robust also with respect to mesh distorsions. Hence, the relative behavior of block and pointwise versions of SAINV can be quite different depending on whether they are applied to shells or to solids.

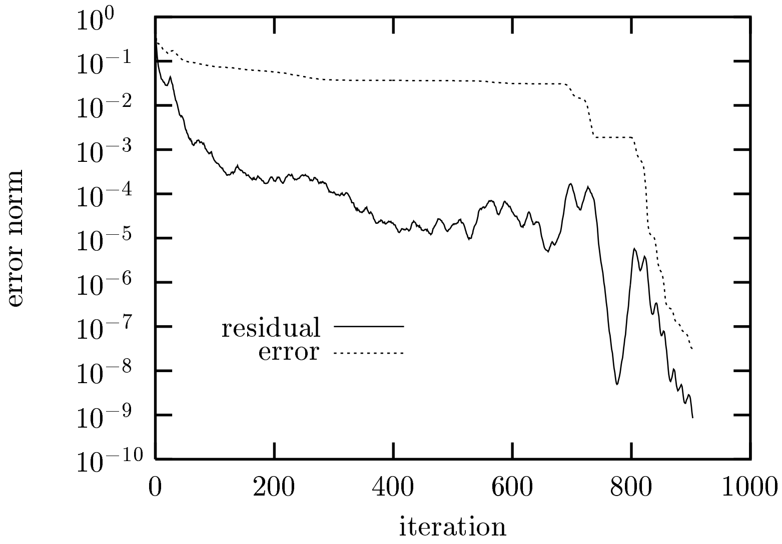

Figure 9: Relative Euclidean norms for CG residual and the true solution error when artificial rhs-vector is used. Pinched cylindrical test case `S3RMT3M3`.

4.3 On the use of artificial right-hand sides

A frequent problem when testing iterative solvers is the difficulty in procuring physically realistic right-hand side vectors $\boldsymbol{b}$. Especially when using matrices from sparse matrix collections, such as the Matrix Market [39], corresponding right-hand sides may not be provided nor there are any hints for the construction of such vectors. As a consequence, developers of linear solvers are often forced to use artificially constructed right-hand sides, having no relation to the problem from which the coefficient matrix arises, and devoided of physical meaning. A typical choice is to compute $\boldsymbol{b}$ as the product of $\boldsymbol{A}$ times a given vector $\boldsymbol{x}$. In this way, the exact solution is known and the accuracy of the computed solution can be checked. Here we would like to emphasize the pitfalls involved in the use of artificial right-hand sides when testing iterative methods.

The convergence histories in Fig. 9 illustrate the fact that when using an artificial right-hand side, declaration of convergence based on the relative residual norm can be completely misleading. As it can be seen from the figure, the relative residual drops rapidly to the level of $10^{-4} - 10^{-5}$, even though there is practically no progress in the accuracy of the solution. However, if a physically realistic load vector is used, the relative residual behaves like the true solution error in Fig. 9. Nevertheless, residual tolerances can still lead to inaccurate results in some cases; more reliable stopping criteria have been proposed, e.g., in [3] and [37].

4.4 On the use of post-filtration

Dropping selected entries in the computed approximate inverse factor $\boldsymbol{Z}$, usually referred to as *a posteriori* filtration, may result in considerable savings and a more efficient solution strategy. The idea is to reject entries that contribute little to the quality of the preconditioner, thereby reducing the cost of applying the preconditioner without too much affecting the rate of convergence of the PCG iteration. A simple criterion is to drop entries below a prescribed drop tolerance. *A posteriori* filtration has been shown to improve the efficiency of the FSAI preconditioned iteration in [31]. Here we consider post-filtration for the SAINV preconditioner. Because SAINV, unlike FSAI, is already computed using a drop tolerance, the usefulness of applying post-filtration is not immediately obvious.

Effects of post-filtration on point and block SAINV preconditioning are shown in Fig. 10 for the thin shell test example `S3RMT3M3`. The solid markers represent computations with unfiltered preconditioners while the open circles correspond to filtered cases obtained from the most dense preconditioners of the unfiltered case. Drop tolerances for the unfiltered case (solid markers) are 0.1, 0.05, 0.04, 0.03, 0.025, 0.02 for block SAINV and 0.06, 0.035, 0.025, 0.02, 0.015, 0.01 for pointwise SAINV. The filtration tolerances are 0.05, 0.1, 0.15, 0.2, 0.3 for block SAINV and 0.06, 0.1, 0.15, 0.2 and 0.4 for pointwise SAINV. It can be seen that pointwise SAINV retains its convergence properties even though 50 % of its elements are dropped. However, block SAINV is scarcely affected by filtering.

Table 9: Surface mounted transistor, effect of post-filtration.

preconditioner	ψ	$\psi_{\rm pf}$	ρ	iterations	P-time	I-time	P+I
SAINV	0.01	0.05	0.53	244	96	67	163
SAINV	0.02	0.07	0.33	341	43	82	125
SAINV	0.02	0.1	0.23	375	43	88	131

$\psi_{\rm pf}$: post filtration tolerance
ρ: density of the preconditioner after post-filtration

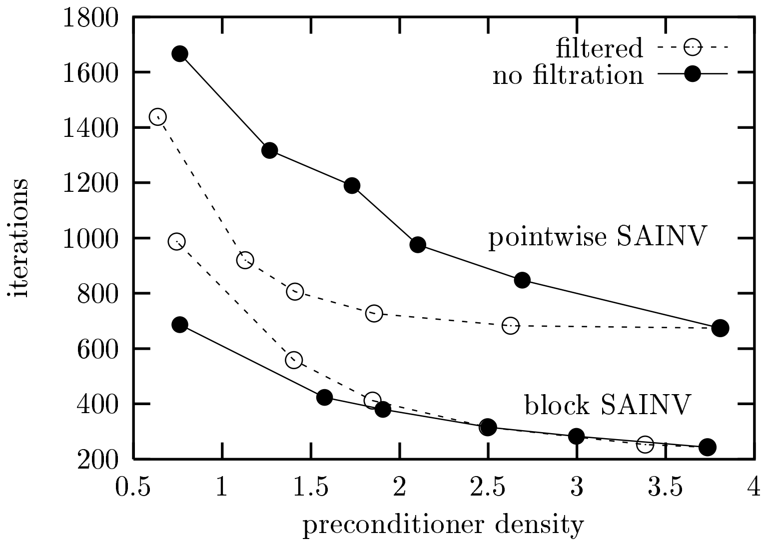

Figure 10: Effect of *a posteriori* filtration on behavior of SAINV preconditioned CG. Test case `S3RMT3M3`.

In 3-D solid examples the behavior appears to be somewhat different. For the engine head problem there is hardly no improvement in iteration times when post-filtration is used.

Slight improvements are obtained for the surface mounted transistor case. However, if the total time for a single solve is considered, the higher preconditioner set up times appears to off-set the gain in iteration time; compare Tables 6 and 9.

5 Conclusions

While direct solution techniques, like skyline or multifrontal solvers, are still widely preferred for solid and structural mechanics problems, preconditioned iterative methods are rapidly gaining in popularity. As computer technology advances, the use of increasingly detailed and sophisticated models results in problems of ever growing size. While direct solvers are very effective for problems of moderate size, the use of iterative methods becomes unavoidable for very large systems. Indeed, the storage requirements for large problems make direct methods unfeasible.

Iterative methods require less storage than direct ones in most cases and can be more easily implemented on parallel architectures. Unfortunately, iterative solvers still lack the reliability of direct methods, and they can fail on challenging problems. Models of thin shells typically result in very ill-conditioned linear systems which are difficult to solve with iterative methods.

In order to increase the reliability of iterative solvers, more robust preconditioning techniques are needed. In this paper we have introduced some new preconditioning techniques for the conjugate gradient method and examined their performance on difficult shell and solid mechanics problems. Besides the SAINV preconditioner, which was first described in [7] and [30], we have developed new block variants of AINV and SAINV, as well as block versions of various incomplete Cholesky preconditioners, including ones based on drop tolerances and the stabilized variant of Ajiz and Jennings. Of these methods, the pointwise and block versions of SAINV are the only ones that do not require any diagonal modifications, and are always well-defined.

For the block methods, the block structure naturally comes from the finite element formulation. Exploitation of this block structure improves robustness and performance on current cache-based architectures.

The importance of point and block diagonal scalings and the effect of different (block) orderings were investigated. In general, scalings and reorderings are found to significantly increase the robustness and performance of the preconditioning. For shell problems, we found that scalings also reduce the influence on the condition number of parameters like the regularizing parameter γ in the Hughes-Brezzi formulation.

Concerning the performance of approximate inverse techniques, we found that the block methods performed better overall than the pointwise ones on thin shell models, where the block size is typically six. On the other hand, for problems in solid mechanics, where the typical block size is only three, we found that the block methods, while reliable, were no better than the pointwise SAINV algorithm. We also found the block methods to be more sensitive to bad aspect ratios than the point ones, exactly the opposite than for the shell models.

Our experiments indicate that while incomplete Cholesky-type methods are often faster than approximate inverse techniques, it may be difficult to find parameter values that will result in a stable incomplete factorization. As for the Ajiz-Jennings robust IC factorization, the diagonal modifications needed to stabilize the decomposition can lead to slow convergence in some cases (as in the T-beam example). In contrast, robust approximate inverse preconditioners like SAINV and block SAINV cannot break down, do not need any diagonal modifications, and often result in convergence rates that are comparable with those for incomplete Cholesky-type preconditioners or even better. Thus, approximate inverse techniques can be a useful alternative to the popular incomplete Cholesky-type preconditioners, even on serial computers. However, they are expected to be especially useful on parallel computers, since their application only requires matrix-vector products.

Parallel implementations of the basic AINV algorithm have been described in [9] and [12], with good results on a range of scalar PDE problems from the modeling of diffusion and transport phenomena on both structured and unstructured grids. Future work will focus on developing parallel implementations of the SAINV and block SAINV preconditioners for large-scale problems in solid and structural mechanics.

Acknowledgements The second author gratefully acknowledges the hospitality of the Institute of Computer Science of the Academy of Sciences of the Czech Republic where part of this work was performed. The work of the third author was supported by Grant Agency of the Academy of Sciences of the Czech Republic grant No. 2030801 and Grant Agency of the Czech Republic grant No. 201/00/0080. Part of this work was done during a visit to the Laboratory of Structural Mechanics at Helsinki University of Technology: The hospitality of the Laboratory is greatly appreciated.

References

[1] M. A. Ajiz and A. Jennings, *A robust incomplete Cholesky conjugate gradient algorithm*, Int. J. Numer. Methods Engrg., 20 (1984), 949–966.

[2] D. J. Allman, *A compatible triangular element including vertex rotations for plane elasticity analysis*, Computers and Structures, 19 (1984), pp. 1–8.

[3] M. Arioli, I. S. Duff and D. Ruiz, *Stopping criteria for iterative solvers*, SIAM J. Matrix Anal. Applic., 13 (1992), pp. 138–144.

[4] C. Ashcraft, *Compressed graphs and the minimum degree algorithm*, SIAM J. Sci. Comput., 16 (1995), pp. 1404–1411.

[5] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, Cambridge, 1994.

[6] S. T. Barnard and M. J. Grote, *A block version of the SPAI preconditioner*, in B. A. Hendrickson et al., editors, Proceedings of the Ninth SIAM Conference on Parallel Processing for Scientific Computing, SIAM, Philadelphia, 1999 [CD-ROM].

[7] M. Benzi, J. K. Cullum and M. Tůma, *Robust approximate inverse preconditioning for the conjugate gradient method*, , SIAM J. Sci. Comput., 22 (2000), pp. 1318–1332.

[8] M. Benzi, R. Kouhia and M. Tůma, *An assessment of some preconditioning techniques in shell problems*, Comm. Numer. Methods Engrg., 14 (1998), pp. 897–906.

[9] M. Benzi, J. Marín and M. Tůma, *A two-level parallel preconditioner based on sparse approximate inverses*, in Iterative Methods in Scientific Computation IV, D. R. Kincaid and A. C. Elster, eds., IMACS Series in Computational and Applied Mathematics, Vol. 5, IMACS, New Brunswick, NJ, 1999, pp. 167–178.

[10] M. Benzi, C. D. Meyer and M. Tůma, *A sparse approximate inverse preconditioner for the conjugate gradient method*, SIAM J. Sci. Comput., 17 (1996), pp. 1135–1149.

[11] M. Benzi and M. Tůma, *Orderings for factorized approximate inverse preconditioners*, SIAM J. Sci. Comput., 21 (2000), pp. 1851–1868.

[12] L. Bergamaschi, G. Pini and F. Sartoretto, *Approximate inverse preconditioning in the parallel solution of sparse eigenproblems*, Numer. Linear Algebra Appl., 7 (2000), pp. 99–116.

[13] R. Bridson, *Multi-Resolution Approximate Inverses*, M.Sc. thesis, University of Waterloo, Ontario, Canada, 1999.

[14] R. Bridson and W.-P. Tang, *Ordering, anistropy, and factored sparse approximate inverses*, SIAM J. Sci. Comput., 21 (1999), pp. 867–882.

[15] R. Bridson and W.-P. Tang, *Refining an approximate inverse*, J. Comput. Appl. Math, to appear.

[16] M. Challacombe, *A simplified density matrix minimization for linear scaling self-consistent field theory*, J. Chem. Phys., 110 (1999), pp. 2332–2342.

[17] M. Challacombe, *A general parallel sparse-blocked matrix multiply for linear scaling SCF theory*, Comp. Phys. Comm., 128 (2000), pp. 93–107.

[18] H. Choi and D. B. Szyld, *Application of threshold partitioning of sparse matrices to Markov chains*, in Proc. IEEE International Computer Performance and Dependability Symposium, IPDS'96, Urbana-Champaign, Illinois, September 4–6, 1996, pp. 158–165, IEEE Computer Society Press, Los Alamitos, CA, 1996.

[19] E. Chow, *A priori sparsity patterns for parallel approximate inverse preconditioners*, SIAM J. Sci. Comput., 21 (2000), pp. 1804–1822.

[20] E. Chow and M. A. Heroux, *An object-oriented framework for block preconditioning*, ACM Trans. Math. Software, 24 (1998), pp. 159–183.

[21] E. Chow and Y. Saad, *Approximate inverse techniques for block-partitioned matrices*, SIAM J. Sci. Comput., 18 (1997), pp. 1657–1675.

[22] J. J. Dongarra, I. S. Duff, D. C. Sorensen and H. A. van der Vorst, *Numerical Linear Algebra for High-Performance Computers*, SIAM, Philadelphia, 1998.

[23] I. S. Duff, A. M. Erisman and J. K. Reid, *Direct Methods for Sparse Matrices*, Oxford University Press, Oxford, 1986.

[24] I. S. Duff and G. Meurant, *The effect of ordering on preconditioned conjugate gradients*, BIT, 29 (1989), pp. 635–657.

[25] M. R. Field, *Improving the Performance of Factorised Sparse Approximate Inverse Preconditioner*, Technical Report HDL-TR-98-199, Hitachi Dublin Laboratory, Dublin, Ireland, 1998.

[26] T. J. R. Hughes and F. Brezzi, *On drilling degrees of freedom*, Comput. Methods Appl. Mech. Engrg., 72 (1989), pp. 105–121.

[27] T. J. R. Hughes, F. Brezzi, A. Masud and I. Harari, *Finite elements with drilling degrees of freedom: Theory and numerical evaluations*, in R. Gruber, J. Periaux and R. P. Shaw, editors, Proceedings of the fifth International Symposium on Numerical Methods in Engineering, pp. 3–17, Springer-Verlag, Berlin and New York, 1989.

[28] T. J. R. Hughes, A. Masud and I. Harari, *Numerical assessment of some membrane elements with drilling degrees of freedom*, Computers and Structures, 55 (1995), pp. 297–314.

[29] A. Ibrahimbegović, R. L. Taylor and E. L. Wilson, *A robust quadrilateral membrane finite element with drilling degrees of freedom*, Int. J. Numer. Methods Engrg., 30 (1990), pp. 445–457.

[30] S. A. Kharchenko, L. Yu. Kolotilina, A. A. Nikishin and A. Yu. Yeremin, *A Reliable AINV-Type Preconditioning Method for Constructing Sparse Approximate Inverse Preconditioners in Factored Form*, Preprint, Russian Academy of Sciences, Moscow, 1999.

[31] L. Yu. Kolotilina, A. A. Nikishin and A. Yu. Yeremin, *Factorized sparse approximate inverse preconditioning IV: Simple approaches to rising efficiency*, Numer. Linear Algebra Applic., 6 (1999), pp. 515–531.

[32] L. Yu. Kolotilina and A. Yu. Yeremin, *Factorized sparse approximate inverse preconditioning I: Theory*, SIAM J. Matrix Anal. Applic., 14 (1993), pp. 45–58.

[33] R. Kouhia, *Using iterative solvers in non-linear continuation algorithm*, in J. Aalto and T. Salmi, editors, Proceedings of the 6th Finnish Mechanics Days, pp. 253–267, 1997.

[34] R. Kouhia and M. Mikkola, *Tracing the equilibrium path beyond compound critical points*, Int. J. Numer. Methods Engrg., 46 (1999), pp. 1049–1074.

[35] J. W. H. Liu, E. G. Ng and B. W. Peyton, *On finding supernodes for sparse matrix computations*, SIAM J. Matrix Anal. Appl., 14 (1993), 242–252.

[36] M. Lyly, R. Stenberg and T. Vihinen, *A stable bilinear element for the Reissner-Mindlin plate model*, Comput. Methods Appl. Mech. Engrg., 110 (1993), pp. 343–357.

[37] G. Meurant, *Numerical experiments in computing bounds for the norm of the error in the preconditioned conjugate gradient method*, Numer. Algorithms, 22 (1999), pp. 353–365.

[38] G. Meurant, *Computer Solution of Large Linear Systems*, North-Holland-Elsevier, Amsterdam etc., 1999.

[39] National Institute of Standards, *Matrix Market*, available online at `http://math.nist.gov/MatrixMarket`.

[40] E. Ng, B. W. Peyton and P. Raghavan, *A blocked incomplete Cholesky preconditioner for hierarchical-memory computers*, in D. R. Kincaid and A. C. Elster, editors, Iterative Methods in Scientific Computation IV, IMACS Series in Computational and Applied Mathematics, vol. 5, pp. 211–221, 1999.

[41] J. O'Neil and D. B. Szyld, *A block ordering method for sparse matrices*, SIAM J. Sci. Stat. Comput., 11 (1990), pp. 811–823.

[42] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Co., Boston, 1996.