

# Preconditioner Updates Applied to CFD Model Problems

Philipp Birken, Jurjen Duintjer Tebbens, Andreas Meister and Miroslav Tůma

October 8, 2007

The work of the first and third author is supported by the German Science Foundation as part of the Sonderforschungsbereich SFB/TR TRR 30 „Prozessintegrierte Herstellung funktional gradierter Strukturen auf der Grundlage thermo-mechanisch gekoppelter Phänomene“, project C2. The work of the second and fourth author is supported by the Program Information Society under project 1ET400300415. The work of the second author is also supported by project number KJB100300703 of the Grant Agency of the Academy of Sciences of the Czech Republic.

## Abstract

This paper deals with solving sequences of nonsymmetric linear systems with a block structure arising from compressible flow problems. The systems are solved by a preconditioned iterative method. We attempt to improve the overall solution process by sharing a part of the computational effort throughout the sequence. Our approach is fully algebraic and it is based on updating preconditioners by a block triangular update. A particular update is computed in a black-box fashion from the known preconditioner of some of the previous matrices, and from the difference of involved matrices. Results of our test compressible flow problems show, that the strategy speeds up the entire computation. The acceleration is particularly important in phases of instationary behavior where we saved about half of the computational time in the supersonic and moderate Mach number cases. In the low Mach number case the updated decompositions were similarly effective as the frozen preconditioners.

*Keywords:* Finite volume methods; Update preconditioning; Krylov subspace methods; Euler equations; Conservation laws

## 1 Introduction

Finite volume methods are standard discretization schemes for both stationary and instationary problems in aerodynamics. As the CFL condition puts a severe restriction on the time step of explicit methods, time integration is often done implicitly. Using Newton's method for the appearing nonlinear equation systems, the problem of solving a partial differential equation numerically is transformed into the problem of solving a sequence of linear equation systems. In general, up to 80% of the CPU time for a flow solver is spent solving the linear systems. Thus, the major bottleneck in numerical simulation is the solution of the sequence of linear systems and there is a continuous demand to improve upon the existing methods.

Popular methods used in solving the large and sparse linear systems involved here include multigrid methods and Krylov subspace methods. Multigrid methods use multiple discretization levels and combine several techniques on the different levels (see, e.g. [26]). For some important classes of problems they are asymptotically optimal, but they can also be sensitive to changes of the problem [11]. Krylov subspace methods are based on projecting the large linear system to subspaces of small dimension (see, e.g. [24]). The subspaces are generated through multiplication of vectors by the system matrix, thus enabling exploitation of sparsity. In favorable cases, dominant properties become apparent at an early stage of computation and a satisfactory approximation to the solution can be obtained in a relatively small number of iterations. In practice, one often combines multigrid methods with Krylov subspace methods by using a method of one class as a preconditioner for a method of the other class (see, e.g. [29]). We will consider here Krylov subspace methods, but the techniques we describe may also be applied to other solvers. For the non-normal linear systems that we have to solve, basically two classes of Krylov subspace methods may be used. In the first class, whose main representative is the GMRES method [25], we find methods that reduce residual norms in every iteration, but that must be restarted for reasons of

storage and computational costs. The second class contains methods like BiCGSTAB [27], working with short recurrences but without guarantee that the process does not start to oscillate or does not break down. Often more important than the choice of the specific Krylov subspace method used is the choice of the preconditioner for the linear systems. For our problems, incomplete factorizations lead to good results that are in many cases hard to improve.

In order to speed up the solution process of the linear systems arising in CFD problems, we will not search for new and even more sophisticated linear solvers or preconditioners in this paper. Instead, we will try to accelerate the existing methods by considering the whole sequence of linear systems and by trying to share some of the computational effort throughout the sequence. In stationary and instationary problems linear systems are often close during many subsequent iterations of the nonlinear process. A well-known way to exploit this is by skipping some evaluations of the Jacobian in Newton's method, changing only the right hand sides. Unfortunately, this leads to weaker convergence of the nonlinear process. Concerning preconditioning, closeness of system matrices has been taken advantage of only in a rather naive way. Very often, a preconditioner is recomputed periodically with some heuristic choice of period, and at a certain point it may be completely frozen [18].

In recent years, a few attempts to *update* preconditioners for large sparse systems have been made in the numerical linear algebra community. The main idea is to derive efficient preconditioners from previous systems of the sequence in a cheap way, thus avoiding the expensive computation of a new preconditioner. For instance, in case of a sequence of linear systems from a quasi-Newton method, straightforward approximate small rank updates can be useful (this is shown in the SPD case in [20], [6]). SPD matrices and updates of incomplete Cholesky preconditioners are considered in [19]. In [3, 7] approximate diagonal and tridiagonal preconditioner updates were introduced for sequences of parametric complex symmetric linear systems. This technique was generalized to approximate (possibly permuted) triangular updates for nonsymmetric sequences in [10]. Finally, recycling of Krylov subspaces by using adaptive information generated during previous runs has been used to update both preconditioners and Krylov subspace iterations (see [22], [15], [21] and [2]). Note that from the mentioned techniques only the last two are designed for sequences of nonsymmetric linear systems.

In this paper we investigate the effect of updating preconditioners on the speed of the solution process for some model problems from CFD. These are chosen from a broad range of Mach numbers to represent different wellknown types of problems. The model problems lead to nonsymmetric linear systems and we will update the corresponding preconditioners based on the technique proposed in [10]. To our knowledge, this kind of strategy is applied to the CFD model problems for the first time. We will describe how we adapted the original technique in order to be used for the model problems. Then we demonstrate that the technique is able to speed up the solution of the involved linear systems, with an acceleration being particularly significant in phases with important changes between subsequent system matrices. In the next section we address the governing equations and the discretization we used for the numerical solution process. In Section 3 we say some words about solving the linear systems in general and then concentrate on the update technique. Among others, we present some new theoretical results and a detailed overview of the modifications for block

systems. In Section 4 we display and discuss the results of numerical experiments with the model problems. Unless otherwise stated,  $\|\cdot\|$  denotes an arbitrary matrix norm.

## 2 Governing Equations and Finite Volume Discretization

### 2.1 The Euler Equations

The equations governing our model problems are the 2D Euler equations. These consist of the conservation laws of mass, momentum and energy, closed by an equation of state. Given an open domain  $D \subset \mathbb{R}^2$ , the equations can be expressed as

$$\partial_t \mathbf{u} + \sum_{j=1}^2 \partial_{x_j} \mathbf{f}_j(\mathbf{u}) = \mathbf{0} \quad \text{in } D \times \mathbb{R}^+,$$

where  $\mathbf{u} = (\rho, m_1, m_2, \rho E)^T$  represents the vector of conserved variables. The flux functions  $\mathbf{f}_j$  are given by

$$\mathbf{f}_j(\mathbf{u}) = \begin{pmatrix} m_j \\ m_j v_1 + \delta_{1j} p \\ m_j v_2 + \delta_{2j} p \\ H m_j \end{pmatrix}, \quad j = 1, 2,$$

with  $\delta_{ij}$  denoting the Kronecker symbol. The quantities  $\rho$ ,  $\mathbf{v} = (v_1, v_2)^T$ ,  $\mathbf{m} = (m_1, m_2)^T$ ,  $E$  and  $H = E + \frac{p}{\rho}$  describe the density, velocity, momentum per unit volume, total energy per unit mass and total enthalpy per unit mass, respectively. The pressure is defined by the equation of state for a perfect gas  $p = (\gamma - 1)\rho(E - \frac{1}{2}|\mathbf{v}|^2)$ , where  $\gamma$  denotes the ratio of specific heats, taken as 1.4 for air.

### 2.2 The Finite Volume Method

We will use here a finite volume discretization. As this approach is covered extensively in the literature [14], [17] we will give only a short summary of the specific concepts used. Our spatial discretization of the time independent physical domain into control volumes or cells  $\sigma_i$  is constructed as a secondary mesh from an underlying Delaunay-triangularization, see figure 1 (left). For a control volume  $\sigma_i$  with volume  $|\sigma_i|$ , let  $N(i)$  denote the set of its neighbors. Then integration of the Euler equations over  $\sigma_i$  and the divergence theorem results in (see figure 1 (right) for the notation)

$$\frac{d}{dt} \mathbf{u}_i(t) = -\frac{1}{|\sigma_i|} \sum_{j \in N(i)} \sum_{k=1}^2 \int_{I_{ij}^k} \sum_{\ell=1}^2 \mathbf{f}_\ell(\mathbf{u}) \mathbf{n}_{ij,\ell}^k ds. \quad (1)$$

We now consider the mean value  $\mathbf{u}_i(t) := \frac{1}{|\sigma_i|} \int_{\sigma_i} \mathbf{u} dx$  in each cell. The line integrals are computed using a second order Gaussian quadrature rule with Gauss point  $x_{ij}^k$  and a

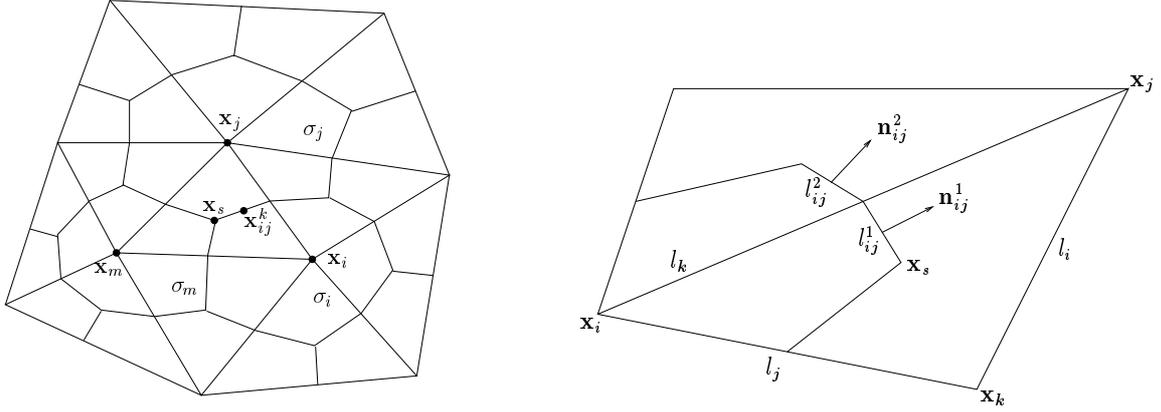


Figure 1: Triangularization and boxes (left). Geometry between boxes (right)

numerical flux function  $\mathbf{H}$ , which we have chosen to be AUSMDV from [28] or for low Mach numbers a Lax-Friedrichs-type flux developed for these cases [16]. Then, we obtain the following evolution equation for the cell averages on  $\sigma_i$ :

$$\frac{d}{dt} \mathbf{u}_i(t) = -\frac{1}{|\sigma_i|} \sum_{j \in N(i)} \sum_{k=1}^2 |l_{ij}^k| \mathbf{H}(\mathbf{u}_i(t), \mathbf{u}_j(t); \mathbf{n}_{ij}^k). \quad (2)$$

To obtain higher order, we use a linear reconstruction technique, combined with the Barth-Jespersen-limiter to reduce the order where necessary.

Implicit time stepping schemes inherently fulfill the CFL stability condition, since the numerical domain of dependence always covers the physical one. In the numerical experiments we will consider the computation of steady states via timestepping with large time steps. Therefore, we employ the implicit Euler scheme and obtain the nonlinear system

$$\mathbf{\Omega} \mathbf{u}^{n+1} = \mathbf{\Omega} \mathbf{u}^n + \Delta t \mathbf{H}(\mathbf{u}^{n+1}),$$

where  $\mathbf{u}$  is the vector of the conservative variables from all cells. Correspondingly,  $\mathbf{H}(\mathbf{u})$  denotes an evaluation of the numerical flux function on the whole grid.  $\mathbf{\Omega}$  is the diagonal matrix of the volumes of the cells, corresponding to the variables in  $\mathbf{u}$ . This equation is solved approximately using one step of Newton's method, which is sufficient for steady state problems. For unsteady problems more steps are often required and the extension of the method is straightforward. The starting value here is  $\mathbf{u}^n$  and the corresponding linear system of equations can be written as (see (2))

$$\mathbf{A} \Delta \mathbf{u} = \mathbf{rhs}(\mathbf{u}^n), \text{ where } \mathbf{A} = \left[ \mathbf{\Omega} + \Delta t \frac{\partial \mathbf{H}(\mathbf{u})}{\partial \mathbf{u}} \right]_{\mathbf{u}^n}, \quad (3)$$

with the update  $\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta \mathbf{u}$ . The matrix  $\mathbf{A} = (\mathbf{A}_{ij})$  has a block structure, where each element  $\mathbf{A}_{ij} \in \mathbb{R}^{4 \times 4}$  vanishes if the corresponding control volumes  $\sigma_i$  and  $\sigma_j$  are not adjacent. Clearly,  $\mathbf{A}$  represents a large and sparse matrix. As the involved grid is in general unstructured, so is the sparsity pattern of  $\mathbf{A}$ . Note that the sparsity pattern of

these matrices remains the same during all time steps. Whereas in some cases at least the pattern is symmetric, usually the matrix itself is nonsymmetric. From (3) we can deduce that the matrix is close to a block diagonal matrix for small time steps and small derivatives of  $\mathbf{H}(\mathbf{u})$ . Diagonal dominance implies some attractive properties of preconditioners and iterative solvers; however, in our problems the dominance is too weak to take advantage of.

### 3 Iterative Solution of the Involved Systems

#### 3.1 Preconditioned Krylov Subspace Methods

As we mentioned in the introduction, we will solve the linear systems from (3) with Krylov subspace methods. For simplicity of notation, we denote linear systems from (3) by  $\mathbf{A}\mathbf{x} = \mathbf{b}$ . For the nonsymmetric matrices we have here, the choices of robust Krylov subspace methods are somewhat limited. A popular and efficient method with low demands on storage costs is the BiCGSTAB method [27]. Whereas the similarly popular GMRES method [25] has some other advantages that we explained in the introduction, we concentrate here on BiCGSTAB because for our finite volume scheme it has turned out to be slightly faster than GMRES. Of major importance for the performance of Krylov subspace methods is the choice of the preconditioner. From experience, right preconditioning seems to be the better choice in the context of compressible flows. Therefore, from now on we assume  $\mathbf{M}$  is a right preconditioner approximating  $\mathbf{A}$  which is applied as

$$\mathbf{A}\mathbf{M}^{-1}\mathbf{x}^P = \mathbf{b}, \quad \mathbf{x} = \mathbf{M}^{-1}\mathbf{x}^P.$$

An overview of preconditioners with special emphasis on application in flow problems can be found in [18] and [8]. In our context, the most appropriate class of preconditioners is that of incomplete LU (ILU) decompositions. Here we focus on ILU(0), which has no additional level of fill beyond the sparsity pattern of the original matrix  $\mathbf{A}$ . This has the obvious advantage that it enables straightforward a priori allocation, and its memory demands are more predictable than for some other incomplete decompositions. Though ILU(0) may not be powerful enough for some difficult problems, for an important number of applications from CFD, including our model problems, it is efficient. In fact, as most problems have a block structure, the used preconditioner is a *block* ILU(0) decomposition (BILU(0)) where pointwise operations are replaced by blockwise operations in the Gaussian elimination process. In our model problems, the blocks correspond to the  $4 \times 4$  units the Jacobian consists of (see (3)). For the involved BILU(0) decompositions we use the following notation. We assume they are computed rowwise, hence the result is a block lower triangular factor denoted by  $\mathbf{L}$  with  $4 \times 4$  identity matrices on the main diagonal and a block upper triangular factor  $\mathbf{U}_{\mathbf{D}}$  with arbitrary nonsingular  $4 \times 4$  matrices on the main diagonal. In addition, we denote by  $\mathbf{D}$  the block diagonal part of  $\mathbf{U}_{\mathbf{D}}$  and let  $\mathbf{U}$  be the matrix  $\mathbf{U}_{\mathbf{D}}$  scaled by  $\mathbf{D}^{-1}$ , i.e.  $\mathbf{U} = \mathbf{D}^{-1}\mathbf{U}_{\mathbf{D}}$ . Then  $\mathbf{U}$  has, like  $\mathbf{L}$ ,  $4 \times 4$  identity matrices on its main diagonal.

The main focus of this paper is efficient preconditioning of the *sequences* of linear systems arising from the scheme described above. Some strategies to share part of the computational effort throughout a sequence were mentioned in the introduction. The two tools we will use

here are periodic *recomputation* of preconditioners combined with *approximate updating*. The idea of periodic recomputation is clear: Computing the preconditioner for every new linear system is time-consuming and unnecessary when the system matrices change slowly. Therefore, we will freeze preconditioners while solving several subsequent systems. Here we will not consider the problem of finding optimal recomputation periods or sophisticated strategies to adapt periods dynamically. This decision is supported by a set of experiments in which we failed to improve a fixed period for recomputation of the frozen preconditioner by simple adaptation guided by a reference number of iterations. The reason was that by simple adaptation to the iteration counts of our preconditioned iterative method we may fail to distinguish what are small/large numbers of iterations with respect to different phases of the problem. Different phases, which may be induced not only by the physics, but also by other adaptive procedures (e.g. for timestepping) may have completely different convergence properties. Therefore, dynamic strategies for preconditioner recomputations should be rather sophisticated. Instead, we will use periodically recomputed frozen preconditioners, which we found to perform rather well.

Our contribution concentrates on a way to update the frozen preconditioners to enhance their power. We believe that our strategy is easy to implement, parameter-free and with a small overhead. The technique we base our updates on is described in [10]. In the next section we have reformulated this strategy for the type of decomposition used here. We present several theoretical statements on the efficiency of the updates for the BILU(0) preconditioning. Furthermore, we give a detailed description of some implementation aspects which are relevant when applying the updates to our applications.

### 3.2 Preconditioner Updates

In addition to a system  $\mathbf{Ax} = \mathbf{b}$  with preconditioner  $\mathbf{M} = \mathbf{LU}_D = \mathbf{LDU}$ , let  $\mathbf{A}^+\mathbf{x}^+ = \mathbf{b}^+$  be a system of the same dimension arising later in the sequence and denote the difference matrix  $\mathbf{A} - \mathbf{A}^+$  by  $\mathbf{B}$ . We search for an updated preconditioner  $\mathbf{M}^+$  for  $\mathbf{A}^+\mathbf{x}^+ = \mathbf{b}^+$ . We have

$$\|\mathbf{A} - \mathbf{M}\| = \|\mathbf{A}^+ - (\mathbf{M} - \mathbf{B})\|,$$

hence the level of accuracy of  $\mathbf{M}^+ \equiv \mathbf{M} - \mathbf{B}$  for  $\mathbf{A}^+$  is the same, in the chosen norm, as that of  $\mathbf{M}$  for  $\mathbf{A}$ . The update techniques from [10] are based on the *ideal* updated preconditioner  $\mathbf{M}^+ = \mathbf{M} - \mathbf{B}$ . If we would use it as a preconditioner, we would need to solve systems with  $\mathbf{M} - \mathbf{B}$  as system matrix in every iteration of the linear solver. Clearly, for general difference matrices  $\mathbf{B}$  the ideal updated preconditioner cannot be used in practice since the systems would be too hard to solve. We will consider cheap approximations of  $\mathbf{M} - \mathbf{B}$  instead.

If  $\mathbf{M} - \mathbf{B}$  is nonsingular, we approximate its inverse by a product of factors which are easier to invert. The approximation consists of two steps. First, we approximate  $\mathbf{M} - \mathbf{B}$  as

$$\mathbf{M} - \mathbf{B} = \mathbf{L}(\mathbf{U}_D - \mathbf{L}^{-1}\mathbf{B}) \approx \mathbf{L}(\mathbf{U}_D - \mathbf{B}), \quad (4)$$

or by

$$\mathbf{M} - \mathbf{B} = (\mathbf{LD} - \mathbf{BU}^{-1})\mathbf{U} \approx (\mathbf{LD} - \mathbf{B})\mathbf{U}. \quad (5)$$

Next we replace  $\mathbf{U}_D - \mathbf{B}$  or  $\mathbf{LD} - \mathbf{B}$  by a nonsingular and easily invertible approximation. In [10] several options are proposed. We have here modified the first option in order to apply it to BILU(0) preconditioners and will approximate as

$$\mathbf{U}_D - \mathbf{B} \approx \text{btriu}(\mathbf{U}_D - \mathbf{B}),$$

or as

$$\mathbf{LD} - \mathbf{B} \approx \text{btril}(\mathbf{LD} - \mathbf{B}),$$

where *btriu* and *btril* denote the block upper and block lower triangular parts (including the main diagonal), respectively. Putting the two approximation steps together, we obtain updated preconditioners in the form

$$\mathbf{M}^+ = \mathbf{L}(\mathbf{U}_D - \text{btriu}(\mathbf{B})) \quad (6)$$

and

$$\mathbf{M}^+ = (\mathbf{LD} - \text{btril}(\mathbf{B}))\mathbf{U}. \quad (7)$$

They can be obtained very cheaply. They ask only for subtracting block triangular parts of  $\mathbf{A}$  and  $\mathbf{A}^+$  (and for saving the corresponding block triangular part of  $\mathbf{A}$ ). In addition, as the sparsity patterns of the factors from the BILU(0) factorization and from the block triangular parts of  $\mathbf{A}$  (and  $\mathbf{A}^+$ ) are identical, both backward and forward substitution with the updated preconditioners are as cheap as with the frozen preconditioner  $\mathbf{LU}_D = \mathbf{LDU}$ .

It is clear from the two approximations we make, that the distance of the proposed updated preconditioners (6) and (7) to the ideal preconditioner is mainly influenced by the following two properties. The first is closeness of  $\mathbf{L}$  or  $\mathbf{U}$  to the identity. If matrices have a strong diagonal, the diagonal dominance is in general inherited by the factors  $\mathbf{L}$  and  $\mathbf{U}$  [5, 3], yielding reasonable approximations of the identity. The second property that helps in approximating the ideal preconditioner is a block triangular part containing significantly more relevant information than the other part. In one of our model problems we emphasize one triangular part by using a numbering of grid cells corresponding to the direction of the flow characteristics. Summarizing, one may expect updates of the form (6) or (7) to be accurate whenever *btril*( $\mathbf{B}$ ) or *btriu*( $\mathbf{B}$ ) is a useful approximation of  $\mathbf{B}$  and when the factor  $\mathbf{L}$  or  $\mathbf{U}$  is close to the identity matrix. The following lemma suggests that under the mentioned circumstances, the updates have the potential to be more accurate than the frozen or any other (possibly recomputed) preconditioner for  $\mathbf{A}^+$ .

**Lemma 1** *Let  $\|\mathbf{A} - \mathbf{LDU}\| = \varepsilon\|\mathbf{A}\| < \|\mathbf{B}\|$  for some  $\varepsilon > 0$ . Then the preconditioner from (7) satisfies*

$$\|\mathbf{A}^+ - \mathbf{M}^+\| \leq \frac{\|\mathbf{U}\| \|\text{bstriu}(\mathbf{B})\| + \|\mathbf{U} - \mathbf{I}\| \|\mathbf{B}\| + \varepsilon\|\mathbf{A}\|}{\|\mathbf{B}\| - \varepsilon\|\mathbf{A}\|} \cdot \|\mathbf{A}^+ - \mathbf{LDU}\|, \quad (8)$$

where *bstriu* denotes the block strict upper triangular part.

This result is a straightforward modification of Lemma 2.1 in [10]; a similar statement can be obtained for updates of the form (6). Having a reference preconditioner  $\mathbf{LDU}$  which is not too weak we may assume that  $\varepsilon\|\mathbf{A}\|$  is small. Then the multiplication factor before

$\|\mathbf{A}^+ - \mathbf{LDU}\|$  in (8) is dominated by the expression  $\|\mathbf{U}\| \frac{\|bstriu(\mathbf{B})\|}{\|\mathbf{B}\|} + \|\mathbf{U} - \mathbf{I}\|$ , which may become smaller than one when  $btril(\mathbf{B})$  contains most of  $\mathbf{B}$  and when  $\mathbf{U}$  is close to the identity matrix. It is possible to show that also the stability of the updates benefits from situations where  $btril(\mathbf{B})$  contains most of  $\mathbf{B}$  and where  $\mathbf{U}$  is close to identity. In our context, the stability is measured by the distance of the preconditioned matrix to identity. This conforms to the treatment of the stability in [9]. Note that the problem of stability in ILU-type of preconditioners was introduced in the classical paper [12]. It was shown in [4] how this problem can be alleviated by some matrix reorderings. Theorem 2.2 in [10], which addresses this stability, can easily be adopted for our case with preconditioning from the right instead of from the left and with block-wise factorization.

The next result is more specific to the situation we are interested in here. It presents a simple sufficient condition for superiority of the update in the case where the frozen preconditioner is a BILU(0) factorization. The result exploits the fact that the BILU(0) preconditioner is an exact decomposition with the sparsity pattern of the matrix it preconditions. It is formulated here for the update (6), but has, of course, an analogue for (7). The matrix  $\mathbf{E}$  denotes the error  $\mathbf{E} \equiv \mathbf{A} - \mathbf{LDU}$  of the BILU(0) preconditioner and  $\|\cdot\|_F$  stays for the Frobenius norm.

**Lemma 2** *If*

$$\sqrt{\|\mathbf{E}\|_F^2 + \|bstril(\mathbf{B})\|_F^2} < \frac{1 - \|\mathbf{I} - \mathbf{L}\|_F^2}{2\|\mathbf{I} - \mathbf{L}\|_F} \|btriu(\mathbf{B})\|_F, \quad (9)$$

where  $bstril$  denotes the block strict lower triangular part of a matrix, then the accuracy of the updated preconditioner  $\|\mathbf{A}^+ - \mathbf{L}(\mathbf{DU} - btriu(\mathbf{B}))\|_F$  is higher than the accuracy  $\|\mathbf{A}^+ - \mathbf{LDU}\|_F$  of the frozen preconditioner.

**P r o o f :** We have

$$\begin{aligned} \|\mathbf{A}^+ - \mathbf{L}(\mathbf{DU} - btriu(\mathbf{B}))\|_F^2 &= \|\mathbf{A} - \mathbf{LDU} - \mathbf{B} + \mathbf{L} \cdot btriu(\mathbf{B})\|_F^2 = \\ &= \|\mathbf{E} - bstril(\mathbf{B}) - (\mathbf{I} - \mathbf{L})btriu(\mathbf{B})\|_F^2 \leq \\ &= (\|\mathbf{E} - bstril(\mathbf{B})\|_F + \|(\mathbf{I} - \mathbf{L})btriu(\mathbf{B})\|_F)^2 = \\ &= \|\mathbf{E} - bstril(\mathbf{B})\|_F^2 + 2\|\mathbf{E} - bstril(\mathbf{B})\|_F \|(\mathbf{I} - \mathbf{L})btriu(\mathbf{B})\|_F + \|(\mathbf{I} - \mathbf{L})btriu(\mathbf{B})\|_F^2. \end{aligned}$$

Note that the sparsity patterns of  $\mathbf{A}$  and  $\mathbf{E}$  are disjoint. Hence, with the assumption (9),

$$\begin{aligned} \|\mathbf{E} - bstril(\mathbf{B})\|_F^2 + 2\|\mathbf{E} - bstril(\mathbf{B})\|_F \|(\mathbf{I} - \mathbf{L})btriu(\mathbf{B})\|_F + \|(\mathbf{I} - \mathbf{L})btriu(\mathbf{B})\|_F^2 &\leq \\ \|\mathbf{E} - bstril(\mathbf{B})\|_F^2 + 2\|\mathbf{E} - bstril(\mathbf{B})\|_F \|(\mathbf{I} - \mathbf{L})\|_F \|btriu(\mathbf{B})\|_F + \|(\mathbf{I} - \mathbf{L})\|_F^2 \|btriu(\mathbf{B})\|_F^2 &\leq \\ \|\mathbf{E} - bstril(\mathbf{B})\|_F^2 + (1 - \|\mathbf{I} - \mathbf{L}\|_F^2) \|btriu(\mathbf{B})\|_F^2 + \|(\mathbf{I} - \mathbf{L})\|_F^2 \|btriu(\mathbf{B})\|_F^2 &\leq \\ \|\mathbf{E} - bstril(\mathbf{B})\|_F^2 + \|btriu(\mathbf{B})\|_F^2 &= \\ \|\mathbf{A}^+ - \mathbf{LDU}\|_F^2 - \|btriu(\mathbf{B})\|_F^2 + \|btriu(\mathbf{B})\|_F^2 &= \|\mathbf{A}^+ - \mathbf{LDU}\|_F^2. \end{aligned}$$

□

Lemmas 1 and 2 may be used in practice to predict what type of update, (6) or (7), will perform better. For example, one may compare the multiplication factor before  $\|\mathbf{A}^+ -$

$\mathbf{LDU}$  in (8) when using (6) or (7) or compare the differences between the left and right hand side in (9) for the choice (6) and the choice (7). However, inequality (9) cannot be satisfied when the numerator is negative, which is very probable in large dimensions. Also, our experience is that the factor before  $\|\mathbf{A}^+ - \mathbf{LDU}\|$  in (8) is larger than one in many cases.

Because of this we present a result which is based on the same idea as (9) but it is stronger. The price for getting a significantly tighter bound is using a less transparent assumption. The result also reveals that the quality of the updates is influenced by further, and more subtle properties than only by closeness of triangular factors to the identity matrix and by the dominance of one triangular part of  $\mathbf{B}$ .

**Lemma 3** *Let*

$$\rho = \frac{\|btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F (2 \cdot \|\mathbf{E} - bstriu(\mathbf{B})\|_F + \|btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F)}{\|btril(\mathbf{B})\|_F^2} < 1.$$

*Then the accuracy  $\|\mathbf{A}^+ - (\mathbf{LD} - btril(\mathbf{B}))\mathbf{U}\|_F$  of the updated preconditioner (7) is higher than the accuracy of the frozen preconditioner  $\|\mathbf{A}^+ - \mathbf{LDU}\|_F^2$  with*

$$\|\mathbf{A}^+ - (\mathbf{LD} - btril(\mathbf{B}))\mathbf{U}\|_F \leq \sqrt{\|\mathbf{A}^+ - \mathbf{LDU}\|_F^2 - (1 - \rho)\|btril(\mathbf{B})\|_F^2}. \quad (10)$$

*P r o o f :* We have, by assumption,

$$\begin{aligned} \|\mathbf{A}^+ - (\mathbf{LD} - btril(\mathbf{B}))\mathbf{U}\|_F^2 &= \|\mathbf{A} - \mathbf{LDU} - \mathbf{B} + btril(\mathbf{B})\mathbf{U}\|_F^2 \\ &= \|\mathbf{E} - bstriu(\mathbf{B}) + btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F^2 \\ &\leq (\|\mathbf{E} - bstriu(\mathbf{B})\|_F + \|btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F)^2 \\ &= \|\mathbf{E} - bstriu(\mathbf{B})\|_F^2 + \rho\|btril(\mathbf{B})\|_F^2. \end{aligned}$$

Because the sparsity patterns of  $\mathbf{A}$  and  $\mathbf{E}$  are disjoint,

$$\|\mathbf{E} - bstriu(\mathbf{B})\|_F^2 + \|btril(\mathbf{B})\|_F^2 = \|\mathbf{E}\|_F^2 + \|\mathbf{B}\|_F^2 = \|\mathbf{E} - \mathbf{B}\|_F^2 = \|\mathbf{A}^+ - \mathbf{LDU}\|_F^2.$$

Hence

$$\|\mathbf{E} - bstriu(\mathbf{B})\|_F^2 + \rho\|btril(\mathbf{B})\|_F^2 = \|\mathbf{A}^+ - \mathbf{LDU}\|_F^2 - (1 - \rho)\|btril(\mathbf{B})\|_F^2.$$

□

With (10), the value of  $\rho$  may be considered a measure for the superiority of the updated preconditioner over the frozen preconditioner. However, interpretation of the value of  $\rho$  is not straightforward. We may write  $\rho$  as

$$\rho = \left( \frac{\|btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F}{\|btril(\mathbf{B})\|_F} \right)^2 + 2 \frac{\|\mathbf{E} - bstriu(\mathbf{B})\|_F}{\|btril(\mathbf{B})\|_F^2}, \quad (11)$$

where the ratio

$$\frac{\|btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F}{\|btril(\mathbf{B})\|_F} \quad (12)$$

shows an interesting dependence of  $\rho$  on the extent to which  $btril(\mathbf{B})$  is reduced after its postmultiplication by  $(\mathbf{I} - \mathbf{U})$ . This is something slightly different from the dependence of the quality of the update on the closeness of  $\mathbf{U}$  to identity. In general, also the second term in (11) should be taken into account; only when the lower triangular part clearly dominates and when  $\mathbf{LDU}$  is a powerful factorization, one may concentrate on (12). Computation of  $\rho$  is not feasible in practice because of the expensive product in  $\|btril(\mathbf{B})(\mathbf{I} - \mathbf{U})\|_F$  but it offers some insight in what really influences the quality of the update. As the proof of the lemma uses only one inequality, one may expect (10) to be a tight bound. We confirm this in the section with numerical experiments.

We will now describe how we exploit updated preconditioners of the form (6) and (7) in the solution process of the problems introduced in the previous section. A first issue is the choice between (6) and (7). We can use some of the previous lemmas to make this choice but we prefer simpler strategies. Just as the ideal preconditioner is approximated in two steps, there are basically two types of simple criteria that can be used. The first criterion compares the closeness of the factors to identity, namely the norms  $\|\mathbf{L} - \mathbf{I}\|$  and  $\|\mathbf{U} - \mathbf{I}\|$ . If the former norm is smaller, then we may expect the approximation made in (4) is better than the one in (5) and we prefer to update the upper triangular part of the decomposition as given in (6); if, on the contrary,  $\mathbf{U}$  is closer to identity in some norm, we update the lower triangular part according to (7). Note that a factor close to identity also leads to stable back or forward substitution with the factor. Therefore, an important consequence of choosing the factor which is closest to identity is that we keep, in the update, the more stable part of the initial decomposition. Due to the lack of diagonal dominance in our applications, stability of the factors is a relevant issue. We call this criterion the *stable update criterion*. On the other hand, it is clear that the quality of the approximation  $\mathbf{U}_{\mathbf{D}} - btriu(\mathbf{B})$  of  $\mathbf{U}_{\mathbf{D}} - \mathbf{B}$  (or  $\mathbf{LD} - btril(\mathbf{B})$  of  $\mathbf{LD} - \mathbf{B}$ ) may have a decisive influence on the power of the preconditioner. The second criterion consists of comparing of  $\|btril(\mathbf{B})\|$  and  $\|btriu(\mathbf{B})\|$ . We assume the most important information is contained in the dominating block triangular part and therefore we update with (6) if  $btriu(\mathbf{B})$  dominates  $btril(\mathbf{B})$  in an appropriate norm. Otherwise, (7) is used. This rule is denoted by *information flow criterion*. Note that in our implementation we always used the Frobenius norm to evaluate the criteria.

Our model problems lead to systems with a block structure and for efficiency reasons, this block structure should be exploited whenever possible. In order to solve linear systems blockwise and, in particular, work with BILU(0) decompositions, we have adapted the original updating technique to updates of the form (6) and (7). Blockwise decompositions, however, make the switch between (6) and (7) a slightly more complicated than in the case of classical pointwise decompositions. Using the update (6) is straightforward but note that in order to obtain  $\mathbf{U}$  and to apply (7) we need to scale  $\mathbf{U}_{\mathbf{D}}$  by  $\mathbf{D}^{-1}$ , as we explained in Section 3.1. Scaling with inverse block diagonal matrices does have, in contrast with inverse diagonal matrices, some influence on overall performance and should be avoided as much as possible. Note that our stable update criterion compares  $\|\mathbf{L} - \mathbf{I}\|$  with  $\|\mathbf{U} - \mathbf{I}\|$  where both factors  $\mathbf{L}$  and  $\mathbf{U}$  have a block diagonal consisting of identity blocks. This means that in order to use the criterion we need to scale  $\mathbf{U}_{\mathbf{D}}$ , even if the criterion decides for (6) and scaling would not have been necessary. We may circumvent this possible inefficiency by considering  $\mathbf{U}_{\mathbf{D}}$  and  $\mathbf{LD}$  instead of  $\mathbf{U}$  and  $\mathbf{L}$ . More precisely, we would compare  $\|\mathbf{D} - \mathbf{U}_{\mathbf{D}}\|$  with  $\|\mathbf{LD} - \mathbf{D}\|$ . We call this third criterion the *unscaled stable update criterion*.

A related issue is the frequency of deciding about the update type based on the chosen criterion. On one hand, there may be important differences in the performance of (6) and (7); on the other hand, switching between the two types implies some additional costs like, for instance, storage of both triangular parts of  $\mathbf{B}$ . Consequently, we believe that the criterion query should not be repeated too often. We adopted the following strategy. After every recomputation of the BILU(0) decomposition, which takes place periodically, we perform one query and then use the chosen type of update throughout the whole period. With the information flow criterion we compare  $\|btril(\mathbf{B})\|$  with  $\|btriu(\mathbf{B})\|$  for the first difference matrix  $\mathbf{B}$  generated after recomputation, i.e. just before solving the system following the system for which we used a new BILU(0) decomposition. For the two stable update criteria we may decide immediately which update type should be used for the next couple of iterations as soon as the new BILU(0) decomposition was computed. Note that as soon as the update type is chosen, we need to store only one triangular part of the old reference matrix  $\mathbf{A}$  (and two triangular factors of the reference decomposition).

Another property of the applications we are interested in here, is that the solution process typically contains heavily instationary phases followed by long nearly stationary phases. This is reflected by parts of the sequence of linear systems with large entries in the difference matrices and other parts where system matrices are very close. Obviously, in the latter parts we may expect a frozen preconditioner to be powerful for many subsequent systems. Our experiments confirm this: In stationary phases we typically observe a deterioration of only 2 to 5 iterations with respect to the iterations needed to solve the system for which the frozen preconditioner was used. Updating the frozen preconditioner in these cases would be counterproductive; it would add some overhead which cannot be compensated by the few savings of iterations. In fact, in these cases there is even a risk that updates produce more iterations, especially when the frozen preconditioner is particularly stable. We therefore apply a very simple technique to avoid unnecessary updating. We start every period by freezing the preconditioner. Denote the number of iterations of the linear solver needed to solve the first system of the period by  $iter_0$ . If for the  $(j + 1)$ st system the corresponding number of iterations  $iter_j$  satisfies

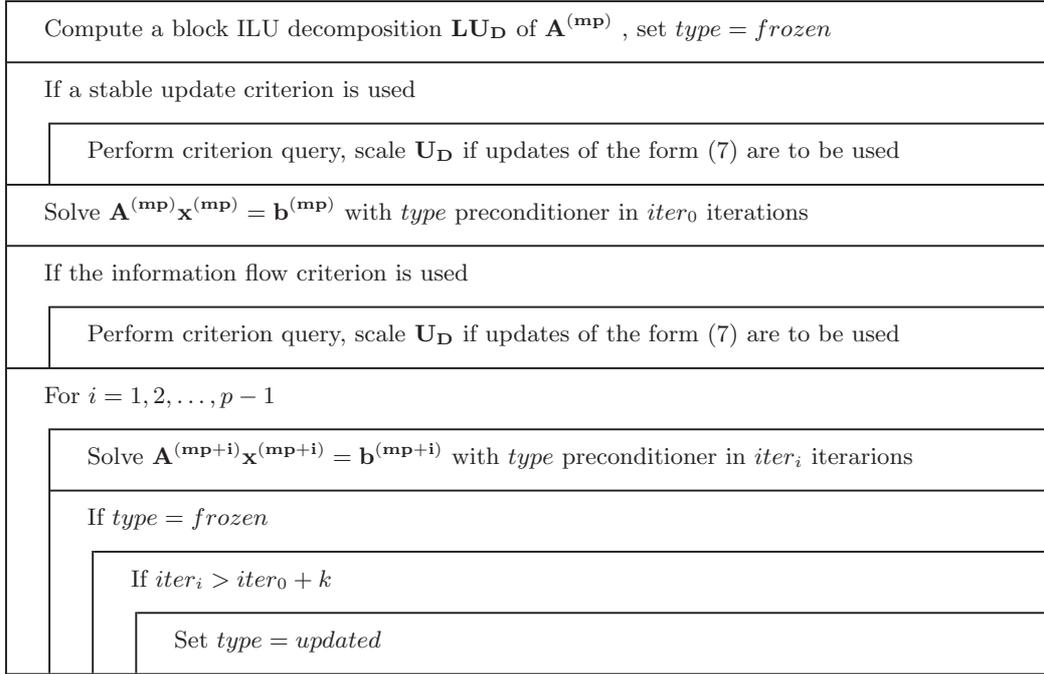
$$iter_j > iter_0 + k, \tag{13}$$

with some threshold  $k \in \mathbb{N}$ , then we use updates for all remaining systems of the period. In accordance with our observations, we used  $k = 3$ . To get a clearer impression of the code decisions to be made we have added a flow diagram. Here,  $p$  denotes the recomputation period and  $m = 0, 1, 2, \dots$

## 4 Numerical Experiments

In this section we demonstrate the behavior of the update technique on some well known steady state test cases. The corresponding linear equation systems are solved until the initial residual has dropped by a factor of  $10^7$ . We always compare periodic refactorization without updating to periodic refactorization with updating, where also the three criteria for deciding whether to use upper or lower updating are compared. The total number of BiCGSTAB iterations as well as the total CPU time for the whole run are recorded. Our primary indi-

**Flow diagram** — preconditioner update decisions after every recomputation



ator to evaluate performance is the CPU time, as a small number of BiCGSTAB iterations may be due to the block preconditioner that takes tremendous amount of computational time. All computations were performed on a Pentium IV with 2.4 GHz.

#### 4.1 Supersonic flow past a cylinder

The first model problem is frontal flow at Mach 10 around a cylinder, which leads to a steady state. 3000 steps of the implicit Euler method are performed. The grid consists of 20994 points, whereby only a quarter of the domain is discretized, and system matrices are of dimension 83976. The number of nonzeros is about  $1.33 \cdot 10^6$  for all matrices of the sequence. For the initial data, freestream conditions are used. Thus, in the beginning, a strong shock detaches from the cylinder, which then slowly moves backward through the domain until reaching the steady state position. Therefore, the linear systems are changing only very slowly during the last 2500 time steps and all important changes take place in the initial phase of 500 time steps. The initial CFL number is 5, which is increased up to 7 during the iteration. The solution is shown in figure 2.

As the flow is supersonic, the characteristics point mostly in one direction. The performance of the linear equation solver can be improved by choosing a numbering of the grid cells that respects the direction of the flow, thereby making the matrix more triangular in nature. This is achieved by numbering first the cells from the inflow boundary, then the cells in direction of the characteristics and by continuing in this manner repeatedly, see [18]. Renumbering reduces the total number of BiCGSTAB iterations by about thirty percent. Furthermore,

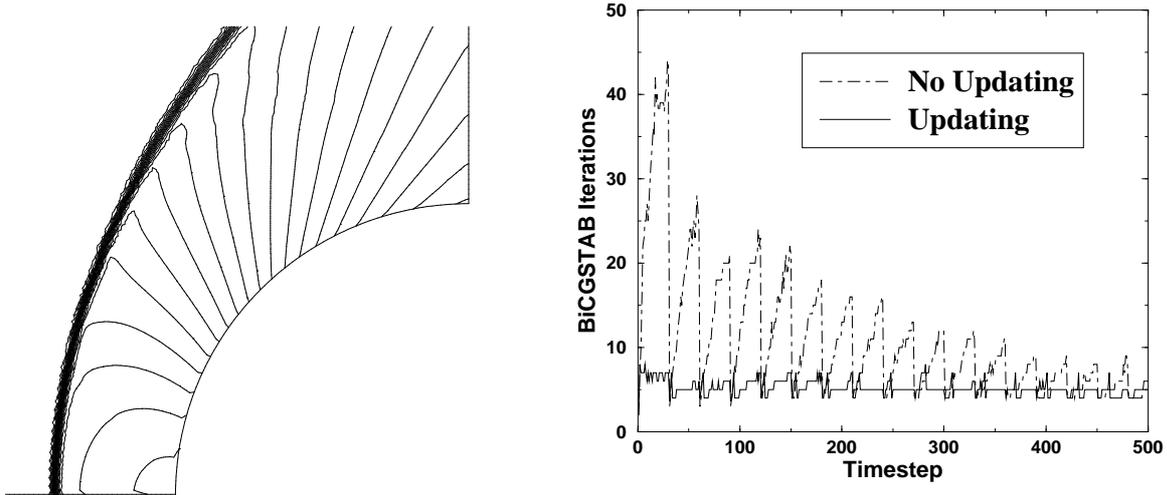


Figure 2: Pressure Isolines (left) and BiCGSTAB iterations per time step (right) for the Cylinder problem

$i$	$\ A^{(i)} - LDU\ _F$	$\ A^{(i)} - M^{(i)}\ _F$	Bound from (10)	$\rho$ from (10)
2	37.454	34.277	36.172	0.571
3	37.815	34.475	36.411	0.551
4	42.096	34.959	36.938	0.245
5	50.965	35.517	37.557	0.104
6	55.902	36.118	38.308	0.083

Table 1: Accuracy of the preconditioners and theoretical bounds

dominance of one of the two triangular parts is exactly the situation in which we expect the update technique to work well. Recall that Lemmas 1, 2 and 3 all suggest that the updated preconditioner is favorably influenced by matrices with a dominating triangular part. In figure 2 excellent performance of the updates is shown for the initial unsteady phase of the first 500 time steps. As subsequent linear systems change heavily, frozen preconditioners produce rapidly deteriorating numbers of BiCGSTAB iterations (with decreasing peaks demonstrating the convergence to steady state). Updating, on the other hand, yields a nearly constant number of iterations per time step. The recomputing period here is thirty and the criterion used is the stable update criterion but other periods and criteria give a similar result. With freezing, 5380 BiCGSTAB iterations are performed in this part of the solution process, while the same computation with updating needs only 2611 iterations.

In Table 1 we explain the superior performance of the updates with the quantities from Lemma 3 for the very first time steps; they demonstrate the general trend for the whole instationary phase. Here,  $M^{(i)}$  denotes the update (7) for the  $i$ th linear system. As the upper bound (10) on the accuracy of the updates is very tight, we conclude that in this problem the power of the updates is essentially due to the small values of  $\rho$ .

In Table 2 we display the performance of the updates for the whole sequence. To evaluate

Period	No updating		Stable update		Unscaled stable update		Information flow	
	Iter.	CPU in s	Iter.	CPU in s	Iter.	CPU in s	Iter.	CPU in s
10	10683	7020	11782	7284	11782	7443	11782	7309
20	12294	6340	12147	6163	12147	6300	12147	6276
30	13787	7119	12503	5886	12503	5894	12503	5991
40	15165	6356	12916	5866	12916	5835	12916	5856
50	16569	6709	13139	5786	13139	5715	13139	5740

Table 2: Total iterations and CPU times for supersonic flow example

the results, first note that the reduction of the BiCGSTAB iterations happens primarily in the first 500 time steps. After 500 time steps, freezing is a very efficient strategy and actually gains again on updating. Thus the visual success of updating is somewhat damped by the long stationary tail of this model problem. The different updating strategies lead to nearly identical results, whereby the stable update criterion is the best, except for the last two periods. As expected, the update criterions all choose to update the lower triangular part according to (7), as the upper triangular part is close to identity due to the numbering of the unknowns and the high Mach number. Therefore, they all obtain the same iteration numbers. Updating is clearly better than freezing if the recomputing period is at least 20. For recomputing periods of 30 or greater, the performance of the updating strategy does not much depend on the period. The CPU time is decreased by about 10 % in general; with the recomputing period 50 it reaches up to 20 %. For longer recomputing periods, the number of iterations is reduced by even more than 20 %. For the period 10 the frozen preconditioner does not deteriorate very much during the periods and achieves lower overall numbers of iterations (and timings) than any updates. This must be caused by the fact that the frozen preconditioner is more stable than the updates. However, the recomputing period 10 is easily beaten by longer periods. If the BILU(0) decomposition would have been recomputed in *every* step, only 11099 BiCGSTAB iterations would be needed, but 28583 seconds of CPU time.

## 4.2 Flow past a NACA0012 airfoil

The second model problem corresponds to the NACA0012 profile at an angle of attack of two degrees on a grid with 4605 cells at different Mach numbers. System matrices are of dimension 18420 and the number of non-zeroes is about  $5 \cdot 10^5$  for all matrices of the sequence. For the initial data, freestream conditions are used.

At first we consider a reference Mach number of  $M = 0.8$ . 1000 steps of the implicit Euler method are performed. The initial CFL number is 5, which is increased up to 30 during the process. For the solution, see figure 3 (left). Transition to steady state is such that after the shock on the airfoil has formed, the rate of convergence slows down, even though the CFL number is increased. Similarly as in the supersonic model, the equation systems differ much from step to step at first, but are very close towards the end. In fact, this behavior is here even more extreme: With decisions based on (13), updating is applied during the

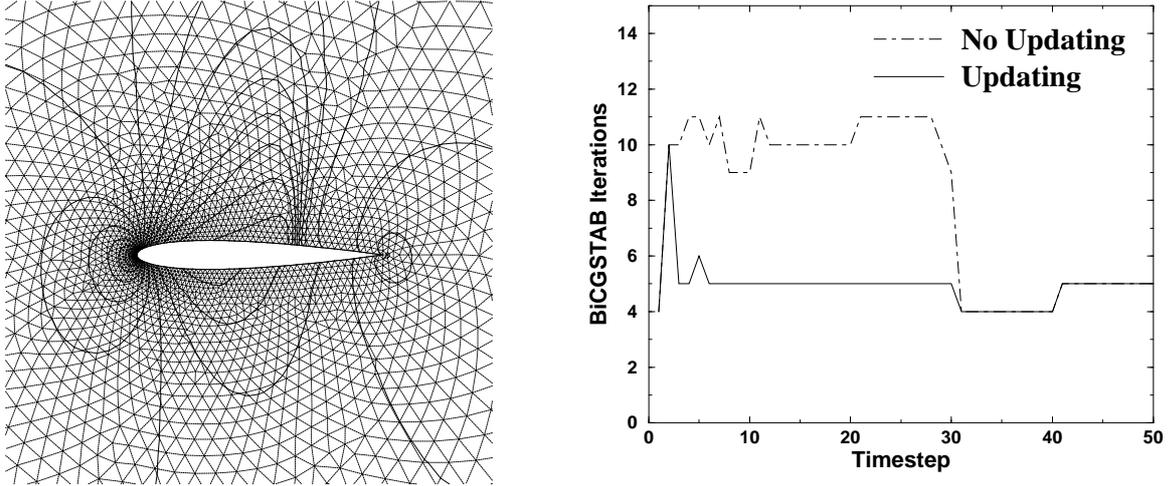


Figure 3: Pressure Isolines and grid (left) and BiCGSTAB iterations per time step (right) for NACA profile with Mach 0.8

Period	No updating		Stable update		Unscaled stable update		Information flow	
	Iter.	CPU in s	Iter.	CPU in s	Iter.	CPU in s	Iter.	CPU in s
10	5375	543	5336	498	5336	494	5336	483
20	5454	497	5364	469	5364	468	5364	459
30	5526	491	5379	464	5379	467	5379	453
40	5558	491	5411	456	5411	462	5411	452
50	5643	525	5413	466	5413	470	5413	448

Table 3: Total iterations and CPU times for transonic flow example

very first period only. To illustrate this, figure 3 (right) compares, for recomputation with a period of 30 time steps, classical freezing with our strategy. Clearly, increasing BiCGSTAB iteration numbers of the frozen preconditioner can be corrected with the updates. But after the first period, there is no need to correct anymore.

The entire process is shown in Table 3. As we can see, the number of iterations decreases if the recomputation period is shortened. This is not true for the CPU time, as recomputations are costly. For the strategy without updating, the CPU time decreases at first, but increases again, as the benefit of fewer recomputations is balanced by the increase in BiCGSTAB iterations. As for the different updating strategies, all lead to both fewer iterations and shorter computing times. As we explained before, the reduction of iterations must be solely due to the very first time steps where updates are applied. The information flow criterion provides the fastest results, whereas the stable update criterion and the unscaled stable update criterion lead to somewhat higher total timings, but still faster than without any updates. All three criteria lead to an identical number of BiCGSTAB iterations, because they always choose the same triangular part to update. If the BILU(0) decomposition would have been recomputed in every step, only 5333 BiCGSTAB-iterations would be needed, but 964 seconds of CPU time. Thus the number of iterations with updating often comes close to

the number with refactorization in every single step. The differences in CPU time come from the cost of selecting the appropriate triangular part and all in all, the computation of the steady state is improved up by about 7 to 15%. Note that again, the CPU time depends less on the choice of the recomputation period with updates than is the case without updating.

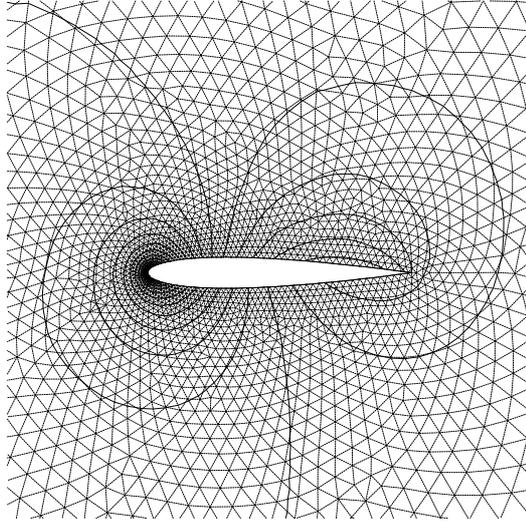


Figure 4: Pressure Isolines and Grid for NACA profile at Mach 0.001.

In the last test case we use a Mach number of  $M = 0.001$ . This problem is much more stiff than the transonic problem. Consequently, the linear systems are harder to solve. Furthermore, for the same CFL number, the time steps should be much smaller due to the larger maximum eigenvalues of the involved matrices. We computed 750 time steps, starting with the CFL number of 0.5, which was increased to its final value equal to two. For the solution, see figure 4, for the comparison of updating techniques see Table 4. In this case, the linear systems do not differ very much among the time steps, not even in the beginning. Thus, the freezing strategy works well and the number of iterations needed increases very slowly in one recomputation cycle. Therefore, even if updating is used, the criterion (13) is seldom fulfilled and the updating strategy has only a small effect in decreasing the iteration numbers, but essentially none on the CPU time. Nevertheless, it is not worse than the classic strategy, which is mainly due to the inclusion of criterion (13): Otherwise, the method would compute an update in every step to no effect. Note that if the BILU(0) decomposition would have been recomputed in every step, 19609 BiCGSTAB-iterations would be needed, but 1437 seconds of CPU time.

## 5 Conclusions

We employed an updating method for block ILU preconditioners for sequences of non-symmetric linear systems in the context of compressible flow. The updating method was motivated by the need to improve frozen preconditioners in order to obtain preconditioners similarly powerful as if they would have been recomputed. For the model problems considered here we showed that as soon as the frozen preconditioners yield high numbers of

Period	No updating		Stable update		Unscaled stable update		Information flow	
	Iter.	CPU in s	Iter.	CPU in s	Iter.	CPU in s	Iter.	CPU in s
10	19444	1189	19288	1158	19398	1121	19289	1129
20	19584	1105	19492	1135	19451	1117	19375	1094
30	19641	1144	19412	1122	19531	1158	19544	1112
40	19622	1104	19521	1112	19594	1114	19523	1107
50	19622	1127	19265	1129	19339	1086	19396	1139

Table 4: Total iterations and CPU times for low Mach flow example

iterations of the linear solver, the updates indeed succeed in reducing the number to the normal level. Whereas the derivation of the updates assumes diagonal dominance of system matrices, the present experiments imply the technique is efficient with rather poor diagonal dominance as well. Note that the success of the new strategy may be significantly enhanced if the time for recomputations becomes prohibitive, which was not our case.

Based on the number of Krylov subspace method iterations, our implementation decides whether updating is necessary. In this way we obtained a preconditioning strategy that is faster than the standard strategy of periodic recomputing for well-known test cases and it is even close to recomputing in every step with respect to iteration numbers. In contrast to periodic recomputations without updates, our method is rather insensitive to the chosen recomputation period.

The method is particularly successful in the phases where the solution process exhibits some kind of instationary behavior and thus it is promising for the computation of instationary flows. In our tables we willingly chose to display results for the whole solution process including long stationary phases of the problems. If we would restrict ourselves to the phases where the updates were actually applied, the results would be even more convincing.

## References

- [1] O. AXELSSON, *Iterative Solution Methods*, Cambridge University Press, Cambridge, 1994.
- [2] J. BAGLAMA, D. CALVETTI, G. H. GOLUB, AND L. REICHEL, *Adaptively preconditioned GMRES algorithms*, SIAM J. Sci. Comput., 20 (1998), pp. 243–269.
- [3] M. BENZI AND D. BERTACCINI, *Approximate inverse preconditioning for shifted linear systems*, BIT, 43 (2003), pp. 231–244.
- [4] M. BENZI, D. B. SZYLD AND A. VAN DUIN, *Orderings for incomplete factorization preconditioners of nonsymmetric problems*, SIAM J. Sci. Comput., 20 (1999), pp. 1652–1670.
- [5] M. BENZI AND M. TUMA, *Orderings for factorized sparse approximate inverse preconditioners*, SIAM J. Sci. Comput., 21 (2000), pp. 1851–1868.

- [6] L. BERGAMASCHI, R. BRU, A. MARTÍNEZ, AND M. PUTTI, *Quasi-Newton preconditioners for the inexact Newton method*, ETNA, 23 (2006), pp. 76–87.
- [7] D. BERTACCINI, *Efficient preconditioning for sequences of parametric complex symmetric linear systems*, Electronic Transactions on Numerical Mathematics, 18 (2004), pp. 49–64.
- [8] A. CHAPMAN, Y. SAAD, AND W. L., *High-order ILU preconditioners for CFD problems*, Int. J. Numer. Methods Fluids, 33 (2000), pp. 767–788.
- [9] E. CHOW AND Y. SAAD, *Experimental study of ILU preconditioners for indefinite matrices*, J. Comp. Appl. Math., 86 (1997), pp. 387–414.
- [10] J. DUINTJER TEBBENS AND M. TUMA, *Efficient preconditioning of sequences of non-symmetric linear systems*, SIAM J. Sci. Comput., to appear (2007).
- [11] H. C. ELMAN AND A. RAMAGE, *Fourier analysis of multigrid for the two-dimensional convection-diffusion equation*, BIT Numer. Math., online version (May, 2006).
- [12] H. C. ELMAN, *A Stability Analysis of Incomplete LU Factorization*, Mathematics of Computation, 47 (1986), pp. 191–218.
- [13] J. FRANK AND C. VUIK, *On the construction of deflation-based preconditioners*, SIAM J. Sci. Comput., 23 (2001), pp. 442–462.
- [14] R. J. LEVEQUE, *Finite volume methods for hyperbolic problems*, Cambridge Texts in Applied Mathematics, Cambridge University Press, Cambridge, 2002.
- [15] D. LOGHIN, D. RUIZ, AND A. TOUHAMI, *Adaptive preconditioners for nonlinear systems of equations*, J. Comput. Appl. Math., 189 (2006), pp. 326–374.
- [16] ———, *Asymptotic based preconditioning technique for low mach number flows*, Z. Angew. Math. Mech., 83 (2003), pp. 3–25.
- [17] A. MEISTER AND T. SONAR, *Finite-volume schemes for compressible fluid flow*, Surveys Math. Indust., 8 (1998), pp. 1–36.
- [18] A. MEISTER AND C. VÖMEL, *Efficient preconditioning of linear systems arising from the discretization of hyperbolic conservation laws*, Adv. Comput. Math., 14 (2001), pp. 49–73.
- [19] G. MEURANT, *On the incomplete Cholesky decomposition of a class of perturbed matrices*, SIAM J. Sci. Comput., 23 (2001), pp. 419–429.
- [20] J. MORALES AND J. NOCEDAL, *Automatic preconditioning by limited-memory quasi-Newton updates*, SIAM J. Opt., 10 (2000), pp. 1079–1096.
- [21] M. L. PARKS, E. DE STURLER, G. MACKEY, D. D. JOHNSON, AND S. MAITI, *Recycling Krylov subspaces for sequences of linear systems*, Technical Report UIUCDCS-R-2004-2421, University of Illinois, 2004.
- [22] ———, *Recycling Krylov subspaces for sequences of linear systems*, SIAM J. Sci. Comput., 28 (2006), pp. 1651–1674.

- [23] Y. SAAD, *ILUT: a dual threshold incomplete LU factorization*, Numer. Linear Algebra Appl., 1 (1994), pp. 387–402.
- [24] ———, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, second ed., 2003.
- [25] Y. SAAD AND M. H. SCHULZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
- [26] V. V. SHAIĐUROV, *Multigrid methods for finite elements*, vol. 318 of Mathematics and its Applications, Kluwer Academic Publishers Group, Dordrecht, 1995. Translated from the 1989 Russian original by N. B. Urusova and revised by the author.
- [27] H. A. VAN DER VORST, *Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of non-symmetric linear systems*, SIAM J. Sci. Stat. Comput., 12 (1992), pp. 631–644.
- [28] Y. WADA AND M.-S. LIOU, *A flux splitting scheme with high-resolution and robustness for discontinuities*, AIAA Paper, 94-0083 (1994).
- [29] C.-T. WU AND H. C. ELMAN, *Analysis and comparison of geometric and algebraic multigrid for convection-diffusion equations*, SIAM J. Sci. Comput., 28 (2006), pp. 2208–2228.