

Updating incomplete factorization preconditioners for model order reduction

Hartwig Anzt¹ · Edmond Chow² · Jens Saak³ ·
Jack Dongarra^{1,4,5}

Received: 18 September 2015 / Accepted: 10 February 2016 / Published online: 23 February 2016
© Springer Science+Business Media New York 2016

Abstract When solving a sequence of related linear systems by iterative methods, it is common to reuse the preconditioner for several systems, and then to recompute the preconditioner when the matrix has changed significantly. Rather than recomputing the preconditioner from scratch, it is potentially more efficient to update the previous preconditioner. Unfortunately, it is not always known how to update a preconditioner, for example, when the preconditioner is an incomplete factorization. A recently proposed iterative algorithm for computing incomplete factorizations, however, is able to exploit an initial guess, unlike existing algorithms for incomplete factorizations. By treating a previous factorization as an initial guess to this algorithm, an incomplete factorization may thus be updated. We use a sequence of problems from model order reduction. Experimental results using an optimized GPU implementation show that updating a previous factorization can be inexpensive and effective, making solving sequences of linear systems a potential niche problem for the iterative incomplete factorization algorithm.

✉ Hartwig Anzt
hantz@icl.utk.edu

¹ Innovative Computing Lab, University of Tennessee, Knoxville, TN, USA

² School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA

³ Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany

⁴ Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

⁵ University of Manchester, Manchester, UK

Keywords Sequence of linear systems · Preconditioner update · Incomplete factorization · Finegrained parallelism · Model order reduction · GPU

1 Introduction

The need to solve sequences of related linear systems arises in several contexts, for example, in the solution of nonlinear systems and time-dependent problems. When the linear systems are large and sparse, preconditioned Krylov subspace solvers are preferred. In this case, it is possible to “recycle” subspace information from solve to solve. Another way to exploit the fact that we are solving a sequence of systems is to reuse the preconditioner for several solves. When the preconditioner becomes stale, it is computed from scratch for the current linear system.

For certain types of preconditioners, it may be possible to update the stale preconditioner, to use for the current linear system, which may be more efficient than computing the preconditioner from scratch. In this paper, we consider updating incomplete factorization preconditioners. These preconditioners are popular because they are effective on a wide range of problems and are problem-independent (they do not need information from the problem besides the matrix). Methods for updating an incomplete factorization have been proposed in the past, mainly for matrices with (possibly complex) diagonal updates, e.g., [11, 18], but also for general updates [19, 25].

In this paper, we consider a different, highly parallel method for updating an incomplete factorization. This is based on a recently proposed iterative algorithm for computing incomplete factorizations that is completely different from the standard algorithm [21]. This algorithm offers fine-grained parallelism suitable for current hardware architectures, including graphics processing units (GPUs). Instead of a Gaussian elimination process, the new incomplete factorization algorithm computes an approximation to the factorization via a fixed-point iteration. A major advantage is that it is possible to compute a factorization starting from an initial guess. By treating a previous factorization as an initial guess to this algorithm, an incomplete factorization may thus be updated. This feature of the algorithm has not been studied before.

To test our ideas, we use a realistic sequence of linear systems arising from the computation of transformation matrices in model order reduction (MOR). The topic of iteratively solving this sequence of problems in MOR has received only limited attention so far. Authors have studied the reuse of subspace information for different shift parameters in the sequence of coefficient matrices [1, 9, 10]. In contrast to this previous work, we focus on the reuse of preconditioner information for these sequences. Regarding parallel computations in MOR, the main development has been on sparse direct solution of the linear systems [6, 7, 29, 30] and matrix sign function approaches using dense matrices. For a selection of papers that either use the same model that we use in the numerical experiments, or also employ GPUs, see, e.g., [8, 12, 13].

This paper is structured as follows. Section 2 introduces model order reduction as one possible origin of a sequence of related linear systems. Section 3

provides background on the iterative algorithm for computing incomplete factorizations. Section 4 reports on experimental tests for solving the linear systems using the iteratively computed incomplete factorizations as preconditioners. We compare the case of updating the preconditioner from a previous factorization, and computing the preconditioner from scratch. We also experiment with reusing such preconditioners for more than one linear system. Section 5 concludes the paper.

2 Model order reduction and shifted linear systems

Model order reduction is a technique for replacing a large and usually sparse dynamical system

$$E\dot{x}(t) = Hx(t) + Bu(t), \quad y(t) = Cx(t), \tag{1}$$

where $E, H \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times p}$, and $C \in \mathbb{R}^{m \times n}$ ($m, p \ll n$), by a substantially smaller one of the form

$$\hat{E}\hat{x}(t) = \hat{H}\hat{x}(t) + \hat{B}u(t), \quad \hat{y}(t) = \hat{C}\hat{x}(t), \tag{2}$$

in order to perform fast simulations of the input-output behavior of the system, e.g., in control and optimization tasks. This means one searches for matrices $\hat{E}, \hat{H} \in \mathbb{R}^{r \times r}$, $\hat{B} \in \mathbb{R}^{r \times p}$, and $\hat{C} \in \mathbb{R}^{m \times r}$ with $r \ll n$, with the property that when using the same input u in both systems, the deviation of the outputs $\|y - \hat{y}\|$ can be bounded in some suitable norm. In projection based MOR, the reduced order matrices are formed by a Petrov-Galerkin-projection approach, i.e.,

$$\hat{E} = W^T E V, \quad \hat{H} = W^T H V, \quad \hat{B} = W^T B, \quad \hat{C} = C V. \tag{3}$$

Several reduction techniques exist that only differ in the way how the tall and thin matrices $V, W \in \mathbb{R}^{n \times r}$ are computed. For a textbook discussion of the system theoretically motivated methods see, e.g., the book by Antoulas [3]. Many of the system theoretic methods are in the one way or the other related to rational Krylov subspaces. This is very obvious for moment matching and rational interpolation methods like the prominent iterative rational Krylov algorithm (IRKA) [27]. In each step of IRKA, the matrices V and W are chosen as orthogonal bases of $\text{span}\{(\sigma_1 E - H)^{-1} B, \dots, (\sigma_r E - H)^{-1} B\}$ and $\text{span}\{(\sigma_1 E - H)^{-1} C^T, \dots, (\sigma_r E - H)^{-1} C^T\}$, respectively, and the set $\Sigma = \{\sigma_1, \dots, \sigma_r\}$ is adapted for the next step until the method converges.

Other methods like Balanced Truncation [31, 42] use the solutions of two Lyapunov equations, the so called Gramians, as the basis for computing V and W . For very large n and small m and p , these Gramians typically have low numerical rank. This property can be used by modern solvers like the low rank alternating directions implicit (LR-ADI) iteration [14], or the rational (block-)Krylov subspace method (RKSM) [23] by representing the low rank factors of the Gramians in a rational Krylov basis.

Whenever rational Krylov subspaces are involved, the main effort in computing the basis is the solution of linear systems of equations of the form

$$(\sigma E - H) X = F, \quad (4)$$

or, with the notation $A_\sigma = (\sigma E - H)$,

$$A_\sigma X = F, \quad (5)$$

for varying shifts $\sigma \in \mathbb{C}$. Here, the matrices H and E are the large and sparse coefficients in (1), and X , F may consist of multiple columns. In the LR-ADI and IRKA cases, for example, the dimensions coincide with either those of B or C^T .

We focus on the linear systems of equations arising in the LR-ADI, when E is symmetric and positive definite (SPD), and H is symmetric and negative definite. This implies that the system (1) is asymptotically stable. Also, all appearing σ are real and positive, such that the matrix in (5) is SPD, and a preconditioned conjugate gradient (PCG) method can be used for fast iterative solution. The linear systems arise consecutively such that the solution of the previous system contributes to the right-hand side of the current system. Therefore, the systems can not be solved simultaneously, and an efficient preconditioner allows for reusing information across the sequence of problems. The iterative algorithm for computing incomplete factorizations that we describe next provides this feature by updating a previously computed factorization for varying shifts σ .

3 Iterative incomplete factorization algorithm

An incomplete factorization is the approximate factorization of a nonsingular sparse matrix A into the product of a sparse lower triangular matrix L and a sparse upper triangular matrix U , i.e.,

$$A \approx LU,$$

where nonzeros or fill-in computed by a Gaussian elimination process is somehow truncated to a sparsity pattern, S . For an overview, see [41]. We use incomplete factorizations to precondition the matrices A_σ described in the previous section. A limitation to the use of these preconditioners is the difficult-to-parallelize factorization in the preconditioner setup phase [16, 22, 33, 39]. In particular, strategies providing some parallelism fail to leverage the fine-grained parallelism of current HPC architectures [17].

The recently proposed iterative algorithm for computing incomplete factorizations is different from previous approaches, as it does not try to parallelize the Gaussian elimination process [21]. Instead, it iteratively approximates the incomplete factors using the property

$$(LU)_{ij} = a_{ij}, \quad (i, j) \in S, \quad (6)$$

where $(LU)_{ij}$ denotes the (i, j) entry of the product of the computed factors L and U , and a_{ij} is the corresponding entry in the matrix A . In particular, the iterative

incomplete factorization algorithm computes the unknowns

$$\begin{aligned}
 l_{ij}, & \quad i > j, \quad (i, j) \in S, \\
 u_{ij}, & \quad i \leq j, \quad (i, j) \in S
 \end{aligned}$$

using the constraints

$$\sum_{k=1}^{\min(i,j)} l_{ik}u_{kj} = a_{ij}, \quad (i, j) \in S \tag{7}$$

which corresponds to enforcing the property (6). The formulation

$$l_{ij} = \frac{1}{u_{jj}} \left(a_{ij} - \sum_{k=1}^{j-1} l_{ik}u_{kj} \right), \quad i > j, \tag{8}$$

$$u_{ij} = a_{ij} - \sum_{k=1}^{i-1} l_{ik}u_{kj}, \quad i \leq j \tag{9}$$

suggests solving for the unknowns using a fixed-point iteration $x = G(x)$ where x is the vector containing the unknowns l_{ij} and u_{ij} for $(i, j) \in S$. It can be proven that the iteration is locally convergent for standard (synchronous) and asynchronous iterations [26]. Global convergence (from an arbitrary initial guess) can be proven for standard iterations in certain conditions, but only with a sufficiently large number of steps. However, convergence for an arbitrary initial $x^{(0)}$ in the asynchronous case cannot be guaranteed [21].

Practically, the rate of convergence depends on the norm of the Jacobian of G . For standard Laplacian discretizations, and using standard synchronous iterations, Fig. 1(left) plots the 1-norm of the Jacobian, showing that it remains less than 1, and thus the fixed-point mapping is a contraction. The middle and right figures show the convergence of factorization in the nonlinear residual norm and in the Frobenius norm of $A - LL^T$, where L is the approximation to the incomplete Cholesky factor. The latter figure shows that convergence to a steady value is rapid.

Previously it was shown that scaling A symmetrically (as a preprocessing step) to have a unit diagonal is an aid to converging the fixed-point iterations [21]. Also, the fixed-point iterations may be started from an initial guess for x formed from the lower




Fig. 1 Results for matrices from (2D) 5-point and (3D) 7-point centered-difference approximations of the Laplacian; *left*: Jacobian 1-norm; *middle*: Relative nonlinear residual norm; *right*: Frobenius norm of $A - LL^T$. For the Jacobian 1-norm and the relative nonlinear residual norm, the results are independent of the problem size. The Laplacian was symmetrically scaled to have a unit diagonal, and the initial L was the lower triangular part of this matrix

and upper triangular parts of A . This is called the “standard initial guess”. We will later define an alternative initial guess when computing a sequence of factorizations of related matrices.

Algorithm 1 gives the pseudocode for the fixed-point iterations for solving the equations (7) in the case of an SPD factorization, $A \approx U^T U$, which is called the “iterative incomplete Cholesky” or “iterative IC” algorithm. In the remainder of the paper we use this algorithm as we address sequences of SPD systems. Each “sweep” (a term we will use repeatedly in this paper) corresponds to one fixed-point iteration updating all the unknowns in the factor U . In contrast to the iterative IC algorithm, we call the standard Gaussian elimination-based incomplete Cholesky algorithm the “exact IC” algorithm, as it computes the exact IC factors.

Algorithm 1 Iterative IC algorithm

```

1  Set unknowns  $u_{ij}$  to initial values
2  for sweep = 1, 2, ... until convergence do
3    parallel for  $(i, j) \in S_U$  do
4       $s = a_{ij} - \sum_{k=1}^{i-1} u_{ki}u_{kj}$ 
5      if  $i \neq j$  then
6         $u_{ij} = s/u_{ii}$ 
7      else
8         $u_{ii} = \sqrt{s}$ 
9      end
10   end
11 end

```

We note that Algorithm 1 assumes a fixed sparsity pattern for the incomplete factorization. It is possible to design a procedure, similar to threshold-based incomplete factorizations, that produces a factorization with a dynamically generated sparsity pattern, i.e., one that is generated during the factorization itself. For example, starting with an initial sparsity pattern, nonzero entries may be added or deleted after each fixed-point sweep, depending on different criteria. This work will be reported in a future publication.

We use a GPU implementation of Algorithm 1 described in [20]. In this implementation, subsets of the components of x are assigned to GPU thread blocks that are scheduled onto the GPU multiprocessors. Each thread block updates the components of x assigned to it. Usually, there are more thread blocks than multiprocessors, which implies that some thread blocks are processed before others, and thus update their components of x before others. Every thread block uses the latest available values for components. Hence, some updates within one fixed-point sweep may use newer data than others. As there is in general no pre-defined order in which thread blocks are scheduled onto multiprocessors, the iteration scheme may be considered as “block-asynchronous” [4]. In [20], several unconventional optimization strategies are proposed that allow for better reuse of GPU cache and control of update order, important for the overall efficiency of the fixed point iteration.

Based on this previous study, we use the following strategies as the default setup for the iterative IC algorithm applied in this paper:

- The CUDA kernel performs one sweep of the iterative IC algorithm. Additional sweeps are performed by launching this kernel multiple times.
- The inputs to the kernel are arrays corresponding to the matrices A and U . The input value of U is the initial guess for the U factor. The output value of U is the computed incomplete factor.
- We use the reverse Cuthill-McKee (RCM) ordering for the system matrix, which is a preferred ordering for incomplete factorizations [24]. Further, RCM ordering reduces the bandwidth of the systems, which is beneficial for cache reuse in the iterative IC algorithm [20].
- The linear systems are scaled such that A has a unit diagonal, which improves the convergence of the iterative IC algorithm.

4 Experimental results

In Section 4.1, we describe a sequence of linear systems arising from a model order reduction problem that serves as our benchmark application. In Section 4.2, we show the performance of the iterative IC preconditioner applied to solving these linear systems. Sections 4.3 and 4.4 are the main contribution of this paper, showing results for the case of computing a factorization by updating a previous factorization. Section 4.5 demonstrates how sparse triangular solves can also be performed efficiently in parallel for these problems.

The experimental platform is a two socket Intel Xeon E5-2670 (Sandy Bridge) system accelerated by an NVIDIA Tesla K40c GPU. The host system has a theoretical peak of 333 GFLOP/s and the GPU has a theoretical peak of 1,682 GFLOP/s (double precision). The main memory size is 64 GB with a theoretical peak memory bandwidth of 51 GB/s. On the K40 GPU, 12 GB of main memory is available at a theoretical peak bandwidth of 288 GB/s. The implementation of the iterative incomplete factorization GPU kernels are implemented in CUDA [35] version 6.5 [36] using a default thread block size of 128. The PCG linear solver is from the MAGMA open-source software library, version 1.6.0 [28]. It is able to solve with multiple right-hand sides simultaneously. The sparse triangular solve routines are taken from the NVIDIA cuSPARSE library [38]. All computations use double precision arithmetic.

4.1 Benchmark application

The benchmark application is the cooling process in rolling mills for steel rail production. A uniform temperature distribution across cross-sections is needed before applying the final rolling step, to avoid degrading material properties. A uniform temperature distribution can be generated via applying active cooling to the surface of the rail. To this end, the boundary of the cross-section is split into seven disjoint parts, where different controls can be applied. These controls are realized as coolant sprays with different temperatures or intensities. Key for a dynamic cooling process

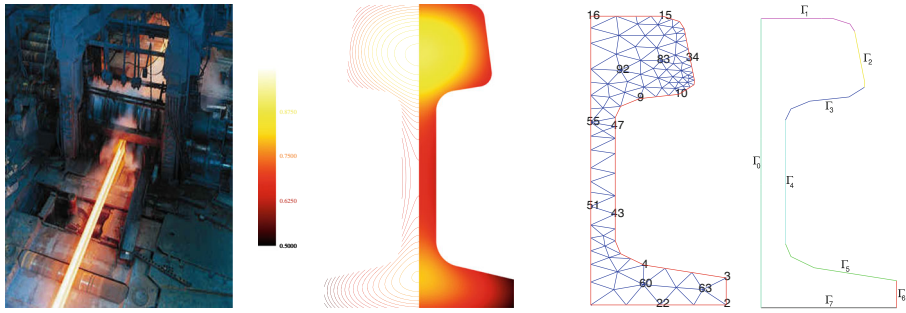


Fig. 2 The steel cooling example: process picture (*left*), sample heat distribution (*center*), initial grid and boundaries (*right*)

is the use of real-time temperature estimators. For this purpose, model order reduction is applied to a spatial discretization using finite elements (see Fig. 2) of the partial differential equation model. In this model, both the stiffness and mass matrices are symmetric, so the PCG solver can be employed to solve equation (5). The shift parameters are chosen according to the optimal procedure suggested for the ADI method by Wachspress (see, e.g., [43]).

The discretization we use is a re-implementation of the steel profile example [15] from the Oberwolfach benchmark collection for model order reduction using the FEniCS [32] finite element software package in version 1.2.0. Like in the original benchmark, we discretize using increasing mesh granularities via global bisection refinement. For every granularity, we obtain a set of matrices H , E and F for equation (4). In Table 1 we list some key properties for these matrices.

The main effort (usually more than 90 % of the execution time in the case of the LR-ADI iteration mentioned in Section 2) towards the computation of the projection matrices V and W (see equation (3)) is to solve equation (5) for different shifts σ and block right-hand sides, each containing 7 vectors. We attempt to accelerate this stage of the MOR process by providing a fast preconditioning strategy for the sequence of shifted linear systems. The length of the sequence of shifts, as well as the shift

Table 1 Properties of the discretizations of the rail problem, showing matrix dimension n , number of nonzeros nnz , number of shifts or matrices in the sequence, the maximum and minimum shift values, and the minimum and maximum condition numbers of matrices in the sequence

problem	size n	$nnz(A_i)$	$\#\sigma_i$	$\max_i(\sigma_i)$	$\min_i(\sigma_i)$	$\min_i(\kappa(A_i))$	$\max_i(\kappa(A_i))$
DISC_1	371	2,343	22	1.02e+00	7.90e-05	3.57e+01	5.58e+02
DISC_2	1,357	8,997	25	4.59e+00	7.91e-05	4.67e+01	2.46e+03
DISC_3	5,177	35,241	28	2.00e+01	7.91e-05	5.26e+01	1.05e+04
DISC_4	20,209	139,473	30	8.35e+01	7.90e-05	5.47e+01	4.01e+04
DISC_5	79,841	554,913	33	3.42e+02	7.90e-05	5.53e+01	1.77e+05
DISC_6	317,377	2,213,697	35	1.39e+03	7.98e-05	5.56e+01	7.10e+05
DISC_7	1,265,537	8,842,881	39	5.64e+03	1.23e-05	5.58e+01	2.44e+06




Fig. 3 Plot of the sequence of shifts σ_i for the distinct discretizations (left) and the condition numbers for the matrices $(\sigma_i E - A)$ for the DISC_5 problem (right)

values themselves, depend on the discretization, the accuracy bound we demand for the model order reduction, and the method we employ for computing V and W . In this paper, we use shifts according to a bound of 10^{-5} for the error $\|y - \hat{y}\|$ when using the same input u in the original and the reduced system and applying LR-ADI based Balanced Truncation. Figure 3 (left) shows the shift values for the different discretizations. The variation of the matrix condition numbers along the sequence of shifts shows a characteristic pattern which is similar for the different discretizations. The condition numbers are plotted for the example problem DISC_5 in Fig. 3 (right).

4.2 Results on iterative incomplete Cholesky factorization

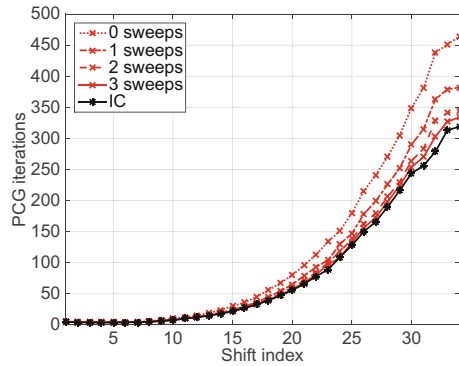
In this section, we first consider the entire set of shifted matrices for DISC_5 and show the PCG iteration counts when using IC preconditioners computed approximately using the iterative IC algorithm. The right-hand side of each system is the block of 7 vectors denoted as matrix F in Eq. (5). The systems were diagonally scaled on the left and right such that the matrices remain symmetric but have unit diagonal. The iterations start with a zero initial guess, and they are stopped when the residual norm relative to the initial residual norm has been reduced below 10^{-6} for the first vector. We always use the level 0 incomplete Cholesky factorization in this paper, but this is not a limitation of the iterative IC algorithm. We recall that the standard initial guess is used for computing the iterative IC factorization, which is the upper triangular part of the scaled coefficient matrix.

Table 2 shows the PCG iteration counts when the iteratively computed IC factorization, using 0 to 5 sweeps, is used as a preconditioner. (Zero sweeps corresponds to only using the initial guess, and resembles using symmetric Gauss–Seidel preconditioning.) Solver iteration counts are also shown for the incomplete Cholesky preconditioner computed exactly using a Gaussian elimination process (labelled “Exact IC”), and for no preconditioning (labelled “No Prec”). The results for the iterative IC algorithm are averaged over 10 runs. Note that the increasing iteration counts for decreasing shift magnitudes are consistent with the increasing condition numbers depicted in Fig. 3. Figure 4 plots the PCG iteration count data for up to 3 sweeps of

Table 2 Solver iteration counts using the preconditioner generated via an exact IC factorization or with up to 5 sweeps of the iterative counterpart. The linear systems are the set of shifted matrices for the DISC_5discretization

Shift σ_i	Exact		Number of sweeps					
	No Prec	IC	0	1	2	3	4	5
-3.42e+02	13	5	5	5	5	5	5	5
-2.74e+02	13	4	5	5	4	4	4	4
-1.90e+02	13	4	5	4	4	4	4	4
-1.22e+02	12	4	5	4	4	4	4	4
-7.60e+01	12	4	5	4	4	4	4	4
-4.68e+01	11	4	4	4	4	4	4	4
-2.87e+01	13	4	5	4	4	4	4	4
-1.76e+01	17	5	6	5	5	5	5	5
-1.08e+01	22	6	8	7	6	6	6	6
-6.58e+00	27	8	10	8	8	8	8	8
-4.02e+00	35	10	12	10	10	10	10	10
-2.46e+00	44	12	15	13	12	12	12	12
-1.50e+00	55	15	19	16	15	15	15	15
-9.19e-01	69	18	24	20	19	18	18	18
-5.62e-01	85	22	30	24	23	22	22	22
-3.44e-01	105	27	36	29	28	27	27	27
-2.10e-01	129	33	45	36	34	33	33	33
-1.28e-01	161	39	56	44	41	40	40	39
-7.86e-02	195	47	67	54	50	48	47	47
-4.80e-02	233	56	80	65	60	57	57	56
-2.94e-02	280	66	96	79	72	69	67	67
-1.80e-02	330	78	113	93	83	81	80	80
-1.10e-02	391	89	134	104	100	97	94	92
-6.72e-03	440	109	151	130	119	108	107	107
-4.11e-03	528	128	180	147	138	133	131	130
-2.51e-03	630	150	215	178	162	156	154	153
-1.54e-03	707	166	241	199	181	173	170	169
-9.40e-04	796	190	271	226	206	198	195	193
-5.77e-04	903	216	305	252	230	224	221	220
-3.55e-04	1000+	244	349	290	264	254	250	248
-2.21e-04	1000+	256	381	315	283	270	264	261
-1.42e-04	1000+	279	438	363	329	301	292	289
-9.85e-05	1000+	313	451	378	341	327	322	320
-7.90e-05	1000+	319	464	382	345	334	329	326
\sum_{σ_i} :	11269+	2930	4231	3500	3193	3055	3004	2982
relative to IC:	3.85+	1.00	1.44	1.19	1.09	1.04	1.03	1.02

Fig. 4 PCG iteration counts for DISC_5 using an exact IC factorization or up to 3 sweeps of the iterative IC algorithm starting with the standard initial guess



the iterative IC algorithm. The most important observation is that very few sweeps are needed to produce preconditioners that yield PCG iteration counts close to those yielded by exact IC factorization.

To understand the behavior of the iterative IC algorithm, Fig. 5 plots the relative nonlinear residual norm,

$$\left\| \left(A - \left(U^{(k)} \right)^T U^{(k)} \right)_S \right\|_F / \left\| \left(A - \left(U^{(0)} \right)^T U^{(0)} \right)_S \right\|_F,$$

and the relative IC residual norm,

$$\left\| A - \left(U^{(k)} \right)^T U^{(k)} \right\|_F / \left\| A - \left(U^{(0)} \right)^T U^{(0)} \right\|_F,$$

corresponding to Fig. 4, where $U^{(k)}$ denotes the IC approximation after k sweeps and $U^{(0)}$ is the initial value for this approximation. Note that the Frobenius norm in the relative nonlinear residual norm is only over the sparsity pattern S ; values outside this sparsity pattern are considered to be zero for the Frobenius norm. Note also that the relative IC residual norm is generally nonzero for an incomplete factorization (it is zero for a full factorization).

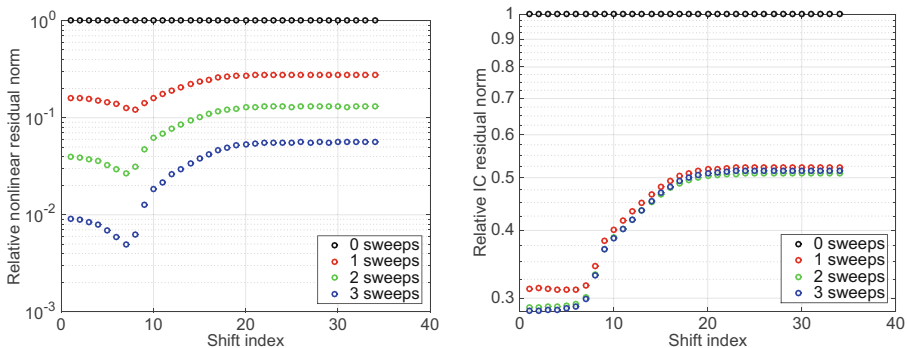


Fig. 5 Relative nonlinear residual norm (left) and the relative IC residual norm (right) for the DISC_5 problems for different shifts σ_i after applying 1, 2, and 3 sweeps of the iterative IC algorithm

Figure 5 (left) shows that after three sweeps the nonlinear residual norm has only been reduced 1–2 orders of magnitude. Since we saw in Fig. 4 that using three sweeps was sufficient for good preconditioning, we conclude that it is not necessary to converge the nonlinear residual norm to a high tolerance in order to produce useful preconditioners. Figure 5 (right) shows that the IC residual norm, however, converges after a very small number of sweeps.

Referring back to Fig. 4, it can be observed the factorizations converge faster or slower for different values of the shift. More precisely, fewer sweeps of the iterative IC algorithm are needed to match the exact IC results for low shift indices than for high shift indices. Generally, matrices earlier in the sequence are easier to solve than matrices later in the sequence. This roughly corresponds to the condition numbers of the matrices shown in Fig. 3 (right). Although the convergence of the iterative IC algorithm is related to the norm of the Jacobian of the fixed-point iteration matrix, these Jacobians have smaller norms if the system matrix have large diagonal elements relative to off-diagonal elements [21], which in turn could be related to the condition number. In any case, connections between matrix properties and the convergence of the iterative IC algorithm are important to make.

Although we have only shown results DISC_5, the results are qualitatively similar for other discretizations. Figure 6 shows the convergence of the relative nonlinear residual norm and the relative IC residual norm for one system for each problem size. Each system corresponds to the last shift of each sequence, which corresponds to the hardest problem of that sequence. The results show only negligible differences in the convergence of the two different residual norms for different problem sizes.

The iterative IC algorithm is much faster than the conventional IC algorithm on highly parallel architectures such as GPUs. Table 3 lists the runtimes for the NVIDIA cuSPARSE [37] implementation of IC, and for 1, 3, and 5 sweeps of the iterative IC algorithm, for different sizes of our MOR problem (the shift used is immaterial here). The results show very significant speedups over the cuSPARSE factorization code, even when the iterative IC algorithm uses 5 sweeps. Note that 3 sweeps typically generate a preconditioner of comparable quality as the exact IC factorization,




Fig. 6 Convergence for the last shifted matrix in each sequence (different mesh sizes) of the relative nonlinear residual norm (left) and the relative IC residual norm (right). The relative IC residual norms of the exact IC factorizations are shown in the right margin of the graph

Table 3 Time [ms] for the exact incomplete Cholesky factorization using NVIDIA’s cuSPARSE library [38] and 1, 3, and 5 sweeps of the iterative counterpart – and the respective speedups

	IC	Time [ms]			Speedup		
		Number of sweeps			Number of sweeps		
		1	3	5	1	3	5
DISC_1	4.88	0.03	0.06	0.08	167.70	87.14	57.48
DISC_2	9.01	0.03	0.06	0.08	309.62	158.07	107.39
DISC_3	15.60	0.04	0.08	0.12	421.62	200.00	134.48
DISC_4	32.00	0.08	0.19	0.30	405.58	167.54	105.26
DISC_5	66.40	0.23	0.62	1.01	286.21	106.58	65.74
DISC_6	142.00	0.78	2.25	3.73	182.75	63.11	38.07
DISC_7	323.00	2.94	8.77	14.50	109.86	36.83	22.28

as shown earlier. We also note that for the iterative IC kernel for smaller problems, the GPU kernel launch overhead is a significant fraction of the total time.

4.3 Results on updating the factorization in the MOR problem sequence

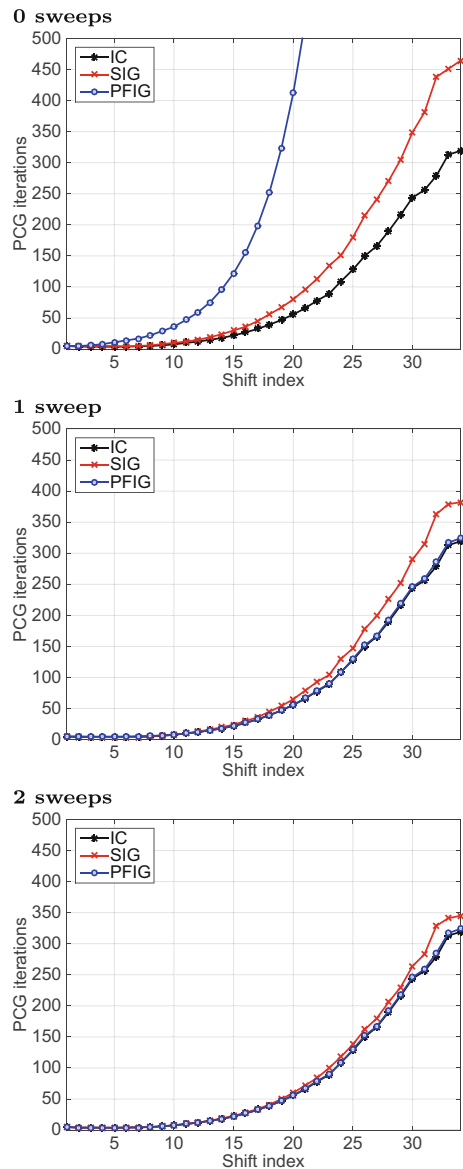
The tests shown so far have solved each shifted system independently, using the standard initial guess (which we now abbreviate as “SIG”) for the iterative IC factorization. We now address solving each shifted system in sequence. For simplicity, we use the same set of shifted systems for the DISC_5 problem presented in Section 4.1. As shown in Fig. 3, the systems generally increase in difficulty for increasing shift index, but the shifts also decrease in size, making the systems with high shift index more similar to each other than the systems with low shift index.

To solve the shifted matrices in sequence, we begin by assuming that we have a factorization of the first matrix. We use an exact IC factorization for this, although an inexact factorization computed using the iterative IC procedure can also be used. Then, each subsequent matrix is factorized by updating the factorization of the previous matrix using some number of sweeps.

Figure 7 shows the PCG iteration counts when the iterative IC factorization was computed using this procedure (which we call using the previous factorization initial guess, or “PFIG”). The results using SIG and exact IC are also shown for comparison. In the Figure, each graph shows the results using 0, 1, or 2 sweeps. PFIG using 0 sweeps corresponds to using the exact IC factorization for the first system to precondition the entire sequence of problems. As can be observed, these PCG iteration counts depart rapidly from those for exact IC.

PFIG with 1 sweep gives results very close to those of exact IC computed for every shift. SIG with 1 sweep does not match exact IC for the harder problems with high shift indices. Assuming 1 sweep, the total cost for using PFIG is one exact factorization for the first matrix, and one sweep per each additional matrix. In summary, PFIG

Fig. 7 PCG iteration counts using a preconditioner generated after 0, 1, and 2 sweeps of the iterative IC algorithm. The ‘IC’ plot uses the exact IC factorization for every system of the sequence. The ‘SIG’ plot computes the factorization using a given number of sweeps on the standard initial guess. The ‘PFIG’ plot computes the factorization using a given number of sweeps on the previous factorization. In this case, the IC factorization of the first matrix of the sequence was computed exactly





using 1 sweep appears sufficient for good performance when an update is performed for every shift. SIG using even 2 sweeps does not match this level of performance.

4.4 Reusing the factorization for many systems

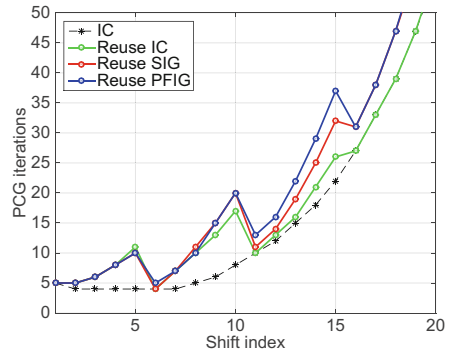

An incomplete factorization preconditioner may be reused for more than one system in a sequence, which can save computations. In this section, we show results for

DISC_5 from computing an incomplete factorization for every 3rd, 5th, and 7th system, i.e., the factorization is reused for 3, 5, or 7 consecutive systems. Below, we use ℓ to denote the recomputation interval. In practice, degradation in the PCG iteration count would be used for triggering recomputation of the factorization.

Reuse 3 times



Reuse 5 times



Reuse 7 times





Fig. 8 PCG iteration counts when each factorization is computed every ℓ systems and reused for ℓ consecutive systems, for $\ell = 3, 5, 7$. Right side graphs zoom in on the left side graphs. A single sweep was used in the SIG and PFIG cases. “IC” denotes the case where the exact IC factorization is used for every system

Exact and iterative IC factorizations were computed every ℓ systems, and were used to precondition ℓ consecutive systems. The iterative IC factorizations were computed using a single sweep applied to either the standard or previous factorization initial guesses (SIG or PFIG). In the case of PFIG, larger ℓ means that the factorization used as initial guess for the current factorization is more “stale” or numerically farther. It could be expected that for larger ℓ , the PFIG updating strategy may be worse than the SIG strategy without updating.

Figure 8 shows the PCG iteration counts due to these preconditioners, for $\ell = 3, 5, 7$. The right side graphs zoom in on the left side graphs. A typical sawtooth pattern can be observed, as the PCG iteration count degrades as a factorization is reused, and then improves again when it is recomputed.

From the graphs on the left side of Fig. 8, fewer total number of PCG iterations are needed when factorizations are updated using PFIG, compared to when factorizations are computed from scratch using SIG. When the recomputation interval is larger, e.g., $\ell = 7$, the overall cost decreases (due to requiring fewer total number of sweeps) but the benefit of PFIG also decreases compared to SIG. However, even for $\ell = 7$, the PFIG strategy is better than the SIG strategy.

It can also be observed that PFIG is better when the difference between shifts is small (high shift indices), since previous factorizations are better guesses when the difference is small, but SIG is *slightly* better when the difference between shifts is large (low shift indices; see zoomed graphs on right side). The cutoff between SIG being better and PFIG being better is about 15 for $\ell = 3$, about 20 for $\ell = 5$, and about 28 for $\ell = 7$. The cutoff increases for larger ℓ which could be expected.

A subtle but interesting feature that can also be noticed is that the degradation in PCG iteration counts when the iterative IC preconditioners are reused is slightly worse than when exact IC preconditioners are reused. This is a very small effect, but can be observed, for example, in the $\ell = 7$ case between indices 8 and 14.

4.5 Approximate sparse triangular solves with block right-hand sides

Although not the main focus of this paper, in this section we address the problem of applying the preconditioner in parallel, i.e., the problem of parallel sparse triangular solves. The MOR application requires solving systems with a block of 7 right-hand sides. Significant research has focused on level scheduling and other parallelization approaches for single right-hand sides [2, 34, 40, 44] that could easily be extended to block problems. On GPUs, it is natural to use NVIDIA’s cuSPARSE function for level-scheduled triangular solves for a block of vectors.

However, following [5], we replace exact sparse triangular solves used as the preconditioning operation, by approximate triangular solves using a small number of Jacobi iterations. For our purpose, we extended the Jacobi method to a block version that is able to iterate multiple vectors simultaneously. In Table 4, we sum the number of PCG iterations needed to solve the sequence of linear systems for each discretization. We compare the use of approximate triangular solves with the cuSPARSE

Table 4 Comparison of iteration counts (first line) and timings [s] (second line) of PCG solving the complete sequence of problems and using different methods of solving with the IC factors: exact triangular solves with cuSPARSE or approximate triangular solves using a fixed number of Jacobi steps

	Number of Jacobi iterations for approximate triangular solves						
	1	2	3	4	5	6	7
Exact							
solves							
DISC.1	\sum_{σ_i} iters: 216	404	239	227	198	200	218
	\sum_{σ_i} time [s]: 0.72	0.59	0.39	0.39	0.38	0.40	0.42
DISC.2	\sum_{σ_i} iters: 405	844	479	440	361	409	372
	\sum_{σ_i} time [s]: 2.27	1.28	0.81	0.76	0.75	0.77	0.80
DISC.3	\sum_{σ_i} iters: 751	1688	954	842	794	742	765
	\sum_{σ_i} time [s]: 8.13	2.72	1.72	1.59	1.60	1.62	1.68
DISC.4	\sum_{σ_i} iters: 1323	3058	1836	1536	1356	1409	1398
	\sum_{σ_i} time [s]: 30.85	6.28	4.79	4.48	4.50	4.84	5.23
DISC.5	\sum_{σ_i} iters: 2935	7169	4162	3612	3483	3200	3226
	\sum_{σ_i} time [s]: 158.70	27.97	24.09	24.41	26.87	27.49	30.73
DISC.6	\sum_{σ_i} iters: 5904	15005	9016	7469	7139	7457	6282
	\sum_{σ_i} time [s]: 785.55	156.36	158.56	158.48	177.82	211.90	201.32
DISC.7	\sum_{σ_i} iters: 10501	26803	16523	14030	12358	12132	12381
	\sum_{σ_i} time [s]: 3207.04	982.42	1074.86	1116.08	1149.81	1306.34	1502.27

block triangular solves mentioned above. The factorizations were constructed using 1 update sweep of the iterative IC algorithm started with the previous factorization initial guess, where the updates are performed for every shift.

Using only a few Jacobi iterations results in a less accurate preconditioner, and generally requires a larger number of PCG iterations. However, the total time for the iterations is generally improved. The tradeoff between the speed and accuracy of these solves is problem-dependent in general. The iterative triangular solves, however, accelerate the solution process for all considered cases. The optimal choice of Jacobi iterations results in up to $7\times$ speedup. Speedup would be less in a method like GMRES compared to PCG since the former has more work per iteration outside the preconditioner.

5 Conclusions

This paper demonstrated the use of an iterative algorithm for computing incomplete factorizations as part of solving a sequence of linear systems from model order reduction. The iterative algorithm can exploit an initial guess for the factorization. We tested the use of a previous factorization as the initial guess. While this initial guess is generally effective, it is not always better than using a standard initial guess when the factors to be updated correspond to a matrix that is numerically far from the current matrix to be factored.

For the MOR benchmark problem, we have only addressed solving one set of linear systems. This reflects, e.g., the case of the LRADI iteration (Section 2) with a pre-computed set of shifts. In the LRADI with an adaptive shift procedure, or the IRKA approach (Section 2), a much smaller batch of shifts is available at once, but the values of the shifts change as they are optimized. In the case of IRKA that means that systems for different shifts can be solved simultaneously, while additional systems for other shifts only become available after solves are performed. For the LRADI the systems for one batch still need to be solved sequentially, but still preconditioners could be computed or updated as soon as the batch becomes available. In this bigger picture, different possibilities for parallelism exist, depending on the available parallel resources.

A general protocol for solving sequences of linear systems involves reuse of the preconditioner and occasional updates or recomputation, the schedule for which will depend on how quickly the matrices change from one to the next. Reuse of the preconditioner may also be combined with other techniques, such as recycling of Krylov subspaces. This paper shows that an iterative algorithm for computing incomplete factorizations can be an important part of solving sequences of linear systems on highly parallel computers.

Acknowledgments The authors thank Maximilian Behr for providing the FEniCS based matrices that served as the benchmark application. This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Numbers DE-SC-0012538 and DE-SC-0010042, the Ministry of Education and Science of the Russian Federation (Agreement N14-11-00190), and NVIDIA.

References

1. Ahmad, M.I., Szyld, D.B., van Gijzen, M.B.: Preconditioned multishift biCG for H2-optimal model reduction. Research Report 12-06-15. Department of Mathematics, Temple University (2012)
2. Alvarado, F.L., Schreiber, R.: Optimal parallel solution of sparse triangular systems. *SIAM J. Sci. Comput.* **14**, 446–460 (1993)
3. Antoulas, A.C.: Approximation of large-scale dynamical systems. SIAM publications, philadelphia (2005)
4. Anzt, H.: Asynchronous and multiprecision linear solvers - scalable and fault-tolerant numerics for energy efficient high performance computing. Ph.D. thesis, Karlsruhe Institute of Technology Institute for Applied and Numerical Mathematics (2012)
5. Anzt, H., Chow, E., Dongarra, J.: Iterative sparse triangular solves for preconditioning. In: Träff, J.L., Hunold, S., Versaci, F. (eds.) Euro-Par 2015: parallel processing, lecture notes in computer science, vol. 9233, pp. 650–661. Springer, Berlin (2015). doi:[10.1007/978-3-662-48096-0_50](https://doi.org/10.1007/978-3-662-48096-0_50)
6. Badía, J.M., Benner, P., Mayo, R., Quintana-Ortí, E.S.: Parallel algorithms for balanced truncation model reduction of sparse systems. In: Dongarra, J.J., Madsen, K., Wasniewski, J. (eds.) Applied parallel computing: 7th international conference, PARA 2004, pp. 267–275. Springer, Berlin (2006). no. 3732 in Lecture Notes in Comput. Sci.
7. Badía, J.M., Benner, P., Mayo, R., Quintana-Ortí, E.S., Quintana-Ortí, G., Remón, A.: Balanced truncation model reduction of large and sparse generalized linear systems. Technical report Chemnitz Scientific Computing Preprints 06-04. Fakultät für Mathematik, TU Chemnitz (2006)
8. Badía, J.M., Benner, P., Mayo, R., Quintana-Ortí, E.S., Quintana-Ortí, G., Saak, J.: Parallel order reduction via balanced truncation for optimal cooling of steel profiles. In: Cunha, J.C., Medeiros, P.D. (eds.) Euro-Par 2005 parallel processing, Lecture Notes in Comput. Sci., vol. 3648, pp. 857–866. Springer, Berlin (2005). doi:[10.1007/11549468_93](https://doi.org/10.1007/11549468_93)
9. Baumann, M., van Gijzen, M.B.: Nested Krylov methods for shifted linear systems. Report 14-01, Delft Institute of Applied Mathematics, Delft University of Technology (2014)
10. Beattie, C.A., Gugercin, S., Wyatt, S.: Inexact solves in interpolatory model reduction. *Linear Algebra Appl.* **436**(8), 2916–2943 (2012). doi:[10.1016/j.laa.2011.07.015](https://doi.org/10.1016/j.laa.2011.07.015)
11. Bellavia, S., De Simone, V., di Serafino, D., Morini, B.: Efficient preconditioner updates for shifted linear systems. *SIAM J. Sci. Comput.* **33**(4), 1785–1809 (2011). doi:[10.1137/100803419](https://doi.org/10.1137/100803419)
12. Benner, P., Ezzatti, P., Kressner, D., Quintana-Ortí, E.S., Remón, A.: Accelerating model reduction of large linear systems with graphics processors. In: Jónasson, K. (ed.) Applied Parallel and Scientific Computing, Lecture Notes in Comput. Sci., vol. 7134, pp. 88–97. Springer, Berlin (2012). doi:[10.1007/978-3-642-28145-7_9](https://doi.org/10.1007/978-3-642-28145-7_9)
13. Benner, P., Ezzatti, P., Quintana-Ortí, E.S., Remón, A.: Using hybrid CPU-GPU platforms to accelerate the computation of the matrix sign function. In: Lin, H.X., Alexander, M., Forsell, M., Knüpfer, A., Prodan, R., Sousa, L., Streit, A. (eds.) Euro-Par 2009 – parallel processing workshops, lecture notes in Comput. Sci., vol. 6043, pp. 132–139. Springer, Berlin (2010). doi:[10.1007/978-3-642-14122-5_17](https://doi.org/10.1007/978-3-642-14122-5_17)
14. Benner, P., Li, J.R., Penzl, T.: Numerical solution of large Lyapunov equations, Riccati equations, and linear-quadratic control problems. *Numer. Lin. Alg. Appl.* **15**(9), 755–777 (2008)
15. Benner, P., Saak, J.: A semi-discretized heat transfer model for optimal cooling of steel profiles. In: Benner, P., Mehrmann, V., Sorensen, D. (eds.) Dimension reduction of large-scale systems, lect. Notes Comput. Sci. Eng., vol. 45, pp. 353–356. Springer, Berlin (2005)
16. Benzi, M., Joubert, W., Mateescu, G.: Numerical experiments with parallel orderings for ILU preconditioners. *Electron. Trans. Numer. Anal.* **8**, 88–114 (1999)
17. Bergman, K., et al.: Exascale computing study: technology challenges in achieving exascale systems. DARPA IPTO ExaScale Computing Study (2008)
18. Bertaccini, D.: Efficient preconditioning for sequences of parametric complex symmetric linear systems. *Electron. Trans. Numer. Anal.* **18**, 49–64 (2004)
19. Calgari, C., Chehab, J.P., Saad, Y.: Incremental incomplete LU factorizations with applications. *Numerical Linear Algebra with Applications* **17**(5), 811–837 (2010). doi:[10.1002/nla.756](https://doi.org/10.1002/nla.756)
20. Chow, E., Anzt, H., Dongarra, J.: Asynchronous iterative algorithm for computing incomplete factorizations on GPUs. In: Lecture Notes in Comput. Sci., vol. 9137, pp. 1–16 (2015)
21. Chow, E., Patel, A.: Fine-grained parallel incomplete LU factorization. *SIAM J. Sci. Comput.* **37**, C169–C193 (2015)

22. Doi, S.: On parallelism and convergence of incomplete LU factorizations. *Appl. Numer. Math.* **7**(5), 417–436 (1991). doi:[10.1016/0168-9274\(91\)90011-N](https://doi.org/10.1016/0168-9274(91)90011-N)
23. Druskin, V., Knizhnerman, L., Simoncini, V.: Analysis of the rational Krylov subspace and ADI methods for solving the Lyapunov equation. *SIAM J. Numer. Anal.* **49**, 1875–1898 (2011)
24. Duff, I.S., Meurant, G.A.: The effect of ordering on preconditioned conjugate gradients. *BIT* **29**(4), 635–657 (1989)
25. Duintjer Tebbens, J., Tüma, M.: Efficient preconditioning of sequences of nonsymmetric linear systems. *SIAM J. Sci. Comput.* **29**(5), 1918–1941 (2007)
26. Frommer, A., Szyld, D.B.: On asynchronous iterations. *J. Comput. Appl. Math.* **123**, 201–216 (2000)
27. Gugercin, S., Antoulas, A.C., Beattie, C.: \mathcal{H}_2 model reduction for large-scale dynamical systems. *SIAM J. Matrix Anal. Appl.* **30**(2), 609–638 (2008)
28. Innovative Computing Lab: Software distribution of MAGMA version 1.6. <http://icl.cs.utk.edu/magma/> (2014)
29. Köhler, M., Saak, J.: Efficiency improving implementation techniques for large scale matrix equation solvers. Chemnitz scientific computing prep. CSC 09-10 TU chemnitz (2009)
30. Köhler, M., Saak, J.: A shared memory parallel implementation of the IRKA algorithm for \mathcal{H}_2 model order reduction. In: Manninen, P., Öster, P. (eds.) *Applied Parallel and Scientific Computing*, Lecture Notes in Comput. Sci., vol. 7782, pp. 541–544. Springer, Berlin (2013). doi:[10.1007/978-3-642-36803-5_42](https://doi.org/10.1007/978-3-642-36803-5_42)
31. Laub, A.J., Heath, M.T., Paige, C.C., Ward, R.C.: Computation of system balancing transformations and other applications of simultaneous diagonalization algorithms. *IEEE Trans. Automat. Control* **32**(2), 115–122 (1987)
32. Logg, A., Mardal, K.A., Wells, G.N., et al.: *Automated solution of differential equations by the finite element method*. Springer (2012)
33. Lukarski, D.: *Parallel sparse linear algebra for multi-core and many-core platforms - parallel solvers and preconditioners*. Ph.D. thesis, Karlsruhe Institute of Technology (KIT), Germany (2012)
34. Naumov, M.: *Parallel solution of sparse triangular linear systems in the preconditioned iterative methods on the GPU*. Tech. Rep. NVR-2011-001 NVIDIA (2011)
35. Corporation, N.V.I.D.I.A.: *NVIDIA CUDA Compute Unified Device Architecture Programming Guide 2.3.1 Edn* (2009)
36. NVIDIA Corporation: *NVIDIA CUDA TOOLKIT V6.5* (2014)
37. NVIDIA Corporation: *CUBLAS library user guide Du-06702-001.v6.5 edn* (2014)
38. NVIDIA Corporation: *CUSPARSE LIBRARY* (2014)
39. Poole, E.L., Ortega, J.M.: Multicolor ICCG methods for vector computers. *SIAM J. Numer. Anal.* **24**, 1394–1417 (1987)
40. Pothén, A., Alvarado, F.: A fast reordering algorithm for parallel sparse triangular solution. *SIAM J. Sci. Stat. Comput.* **13**(2), 645–653 (1992). doi:[10.1137/0913036](https://doi.org/10.1137/0913036)
41. Saad, Y.: *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, Philadelphia (2003)
42. Tombs, M.S., Postlethwaite, I.: Truncated balanced realization of a stable nonminimal state-space system. *Int. J. Control* **46**(4), 1319–1330 (1987)
43. Wachspress, E.L.: *The ADI model problem*. Springer, New York (2013)
44. Wolf, M., Heroux, M., Boman, E.: Factors impacting performance of multithreaded sparse triangular solve. In: Palma, J., Daydé, M., Marques, O., Lopes, J. (eds.) *High Performance Computing for Computational Science – VECPAR 2010*, Lecture Notes in Comput. Sci., vol. 6449, pp. 32–44. Springer, Berlin (2011). doi:[10.1007/978-3-642-19328-6_6](https://doi.org/10.1007/978-3-642-19328-6_6)