



New robust limited-memory incomplete Cholesky preconditioners

Jennifer Scott

STFC Rutherford Appleton Laboratory

Miroslav Tůma

Institute of Computer Science
Academy of Sciences of the Czech Republic

Overview

- ▶ Introduction : background and history of *IC* factorizations
- ▶ Semi-definite modification schemes : prevent breakdown
- ▶ Limited-memory approach
- ▶ Numerical experiments on practical problems
- ▶ The **indefinite case** signed *IC* factorizations for KKT systems

Introduction

We want to solve

$$Ax = b$$

where $A \in \mathbb{R}^{n \times n}$ is large and sparse.

Direct Methods: Factorize $A = LU$, solve $Ly = b$, $Ux = y$.

Black-box, robust.

Memory-hungry \Rightarrow unsuitable for very large problems.

Iterative Methods: CG, GMRES, BiCGStab, etc.

Matrix-free. Fast? Efficient? .

Non-robust, performance depends on preconditioner.

Introduction

An ideal preconditioner should be:

- ▶ cheap to compute
- ▶ sparse and fast to apply
- ▶ provide sufficient approximation of the algebraic problem
- ▶ result in rapidly converging preconditioned iterative method

Key target for library software is **robustness**

Introduction

Incomplete Cholesky (alert/IC) factorization

$$A \simeq LL^T$$

Some entries that occur in complete factorization are ignored.

- ▶ Long history of incomplete factorizations.
- ▶ Early days (late 1950s and 1960s) motivated by finite differences for PDEs. Often for specific problems.
- ▶ Real revolution in practical use and growth in popularity came in late 1970s.
- ▶ In particular, Meijerink and van der Vorst '77 recognised potential of incomplete factorizations as preconditioners for use with CG and proved existence for M -matrices (later extended to H -matrices).

Introduction

Different variants of incomplete factorizations:

- ▶ $IC(\tau)$: Dropping by value (Tuff and Jennings '73)
- ▶ $IC(\ell)$: originally exploited finite difference-based structure (small number of sub-diagonals). Generalised to level-based approach to preserve structure (Watts '81)
- ▶ $IC(p)$: Limited/prescribed memory: Axelsson, Munksgaard '83; Jones, Plassman '95; Saad '94.

Lots of variations/hybrids that combine approaches.

Introduction: problem of breakdown

- ▶ Kershaw '78 **locally perturbed** zero or negative diagonal entries to prevent breakdown so method more widely applicable. Straightforward but can give large growth and unstable preconditioner.
- ▶ Manteuffel '80 proposed **global diagonal shift** so that $A + \alpha I$ factorized for some $\alpha > 0$. Shift α chosen by trial-and-error but can be effective.
- ▶ Alternative approach: **positive semi-definite modifications**.

Our goals

- ▶ Study two positive semi-definite modification schemes:
 - ▶ Jennings and Malik '77,'78 (and Ajiz and Jennings '84)
 - ▶ Tismenetsky '91 (and Kaporin '98)

Our goals

- ▶ Study two positive semi-definite modification schemes:
 - ▶ Jennings and Malik '77,'78 (and Ajiz and Jennings '84)
 - ▶ Tismenetsky '91 (and Kaporin '98)
- ▶ Propose **memory-efficient** variant of Tismenetsky approach, optionally combined with Jennings and Malik modifications or diagonal shifts.

Our goals

- ▶ Study two positive semi-definite modification schemes:
 - ▶ Jennings and Malik '77,'78 (and Ajiz and Jennings '84)
 - ▶ Tismenetsky '91 (and Kaporin '98)
- ▶ Propose **memory-efficient** variant of Tismenetsky approach, optionally combined with Jennings and Malik modifications or diagonal shifts.
- ▶ Present comprehensive numerical results.

Positive semi-definite modifications I

- ▶ Diagonal modification scheme first introduced by Jennings and Malik '77, '78 (also Jennings and Ajiz '84).
- ▶ Every time off-diagonal entry discarded, corresponding diagonal entries modified by adding **SPSD** matrix

$$\begin{array}{c}
 \\
 i \\
 j
 \end{array}
 \begin{pmatrix}
 & & i & & j & & \\
 & \ddots & & & & & \\
 & & |a_{ij}| & & -|a_{ij}| & & \\
 & & & \ddots & & & \\
 & & -|a_{ij}| & & |a_{ij}| & & \\
 & & & & & \ddots & \\
 & & & & & & \ddots
 \end{pmatrix}$$

Jennings-Malik approach

- ▶ **Breakdown-free** factorization that can be expressed as

$$A = LL^T - E$$

where error matrix E is sum of SPSD matrices.

- ▶ **But** modifications to A can be significant.
- ▶ Popular in some engineering applications.

Positive semi-definite modifications II

- ▶ More sophisticated modification scheme due to Tismenetsky '91 (and Kaporin '98).
- ▶ Introduces use of **intermediate memory** that is employed during construction of L but then discarded.
- ▶ Shown to be very **robust** but it “has unfortunately attracted surprisingly little attention” (Benzi '02).
- ▶ One possible reason for this is it suffers from a serious drawback: **memory requirements can be prohibitively high**.

We aim to address memory problem, while retaining robustness.

Tismenetsky approach

Based on matrix decomposition of form

$$A = (L + R)(L + R)^T - E$$

- ▶ L is lower triangular with positive diagonal entries used for preconditioning,
- ▶ R is strictly lower triangular with small entries that is used to stabilise the factorization process, and
- ▶ E has the structure

$$E = RR^T.$$

Tismenetsky approach

- ▶ On j -th step, decompose col. 1 of Schur complement S into

$$l_j + r_j \quad \text{with} \quad |l_j|^T |r_j| = 0,$$

where entries of l_j are retained in incomplete factorization and those in r_j are discarded.

- ▶ On next step, S updated by subtracting

$$(l_j + r_j)(l_j + r_j)^T.$$

- ▶ Tismenetsky omits the term

$$E_j = r_j r_j^T. \tag{1}$$

- ▶ Thus, SPSD matrix implicitly added to A .

Kaporin's use of drop tolerances

- ▶ Obvious choice for r_j are smallest off-diagonal entries in col j .
- ▶ Controls size of L but **not** memory required to compute it.
- ▶ Kaporin '98: entries of magnitude at least τ_1 kept in L and those smaller than τ_2 are dropped from R .
- ▶ Now E has structure

$$E = RR^T + F + F^T,$$

F strictly lower triangular matrix that is **not computed**;
 R used in computation of L but **discarded**.

Problem of unrestricted L and R

- ▶ With no restriction on size of L and R , can achieve **high quality** preconditioner but memory demands **high**.
- ▶ Also can be very **expensive** to compute making approach impractical for the very large problems iterative methods designed for.

Remedy: impose memory limit on L and R .

What about breakdown?

- ▶ If we impose memory limit and/or drop small entries, Tismenetsky approach **not** guaranteed breakdown free.
- ▶ Use **global diagonal shift**? (Manteuffel)
Note: **multiple restarts** may be required so potentially **expensive**.
- ▶ Or **combine** with Jennings-Malik compensation?

How to combine approaches?

There are a number of possibilities:

- ▶ Compensate for **all** entries not retained in L or R .
- ▶ Allow entries in RR^T that do not lead to any further fill-in and compensate for all remaining entries of RR^T .

Test environment

- ▶ Problems from University of Florida Collection.
- ▶ Selected all non-diagonal SPD matrices with $n > 1000$.
- ▶ Removed those with duplicate sparsity patterns.
- ▶ All problems prescaled (this is **important**).
- ▶ Following initial experiments, 8 problems discarded as unable to achieve convergence without large amount of fill.
- ▶ **Test set of 145 problems.**

Test environment (continued)

- ▶ CG used with $x_0 = 0$, b computed so that $x = 1$, and stopping criteria

$$\|Ax_k - b\| \leq 10^{-10} \|b\|$$

with limit of 2000 iterations.

- ▶ All software written in Fortran.


Test environment (continued)

- ▶ What to measure? iteration counts? timings? sparsity of L ?
- ▶ We define the **efficiency** of preconditioner to be

$$iter \times nz(L)$$


- ▶ **Performance profiles** (Moré, Dolan '02) used to assess performance.
- ▶ In our tests, **lsize** is max. number of fill entries in each col. of L and **rsize** is max. number of entries in each col. of R .

Efficiency for $rsize=0$, no diagonal compensation



- ▶ These results are **without** diagonal compensation and **no dropping** of small entries equivalent to ICFS code of Lin and Moré '99.
- ▶ Rather insensitive to choice of `lsize`.

Efficiency for $rsize=0$, with/without SJM



- ▶ These results are **with** and **without** standard Jennings-Malik (SJM) diagonal compensation.
- ▶ Conclude that compensation not generally useful in this case.

Iterations and time for $rsize=0$, with/without SJM

Comparison of using global diagonal shifts (**GDS**) with the Jennings-Malik strategy (**SJM**) ($lsize = 10$).

Figures in parentheses are number of shifts and final shift; times are seconds

Problem	Iterations		Total time	
	GDS	SJM	GDS	SJM
HB/bcsstk28	232 (2, $4.0 * 10^{-3}$)	468	0.120	0.221
Cylshell/s3rmq4m1	648 (2, $4.0 * 10^{-3}$)	838	0.381	0.459
GHS_psdef/lldoor	437 (3, $8.0 * 10^{-3}$)	643	66.4	91.5
GHS_psdef/audikw_1	707 (2, $2.0 * 10^{-3}$)	1442	157	303

- Our experience: **generally better to use diagonal shift.**



Results for `rsize` varying

We now consider using **intermediate memory** (`rsize`>0).

We start by performing **no** diagonal compensation.

Results for `rsize` varying

Efficiency (left) and total time (right) (`lsize=5`)



- `rsize=-1` is unlimited memory for R (not practical).

Results with/without diagonal compensation

Recall:

Limited memory Tismenetsky approach based on decomposition

$$A = LL^T + LR^T + L^T R - E, \quad E = RR^T + F + F^T,$$

where F is not computed but R is.

Results with/without diagonal compensation


Consider three strategies for dealing with RR^T :

- ▶ $jm = 0$: allow entries of RR^T that cause no further fill in $LL^T + LR^T + L^T R$ and discard all other entries of RR^T .
- ▶ $jm = 1$: as above but use Jennings-Malik compensation for discarded entries of RR^T .
- ▶ $jm = 2$: discard all entries of RR^T .

We run these options **with (T)** and **without (F)** diagonal compensation for entries discarded from R .

Results with/without diagonal compensation

Efficiency (left) and total time (right) ($lsize=rsize=10$)





- ▶ Compensating for dropped entries of R generally not beneficial.
- ▶ Reliability slightly improved if entries of RR^T allowed ($jm=0$) but faster and better efficiency to ignore RR^T ($jm=2$).

New *IC* code

- ▶ Based on our findings, we have developed a new IC code called **HSL_MI28**.
- ▶ Can be used as a “black-box” to compute an efficient and robust *IC* preconditioner.
- ▶ But also **flexible**, allowing user to choose the scaling, ordering, diagonal shift, drop tolerances etc.
- ▶ Importantly, the amount of **memory** used (for both L and R) is under the user's control.



Comparison with level-based approach ($IC(3)$)

Efficiency (left) and iterations (right).



Comparison with direct solver HSL_MA97

Total time: all problems (left) and large problems (right).



HSL_MI28 can sometimes compete with direct solver (and succeed when HSL_MA97 runs out of memory).

Concluding remarks on positive-definite case

Concluding remarks on positive-definite case

- ▶ We have explored the use of diagonal compensation with a **limited memory Tismenetsky** approach.

Concluding remarks on positive-definite case

- ▶ We have explored the use of diagonal compensation with a **limited memory Tismenetsky** approach.
- ▶ The proposed limited memory Tismenetsky approach has been shown to be **robust and efficient**.

Concluding remarks on positive-definite case

- ▶ We have explored the use of diagonal compensation with a **limited memory Tismenetsky** approach.
- ▶ The proposed limited memory Tismenetsky approach has been shown to be **robust and efficient**.
- ▶ Using **restricted intermediate memory** improves efficiency.

Concluding remarks on positive-definite case

- ▶ We have explored the use of diagonal compensation with a **limited memory Tismenetsky** approach.
- ▶ The proposed limited memory Tismenetsky approach has been shown to be **robust and efficient**.
- ▶ Using **restricted intermediate memory** improves efficiency.
- ▶ But diagonal compensation to prevent breakdown appears **less important** than generally supposed.

Concluding remarks on positive-definite case

- ▶ We have explored the use of diagonal compensation with a **limited memory Tismenetsky** approach.
- ▶ The proposed limited memory Tismenetsky approach has been shown to be **robust and efficient**.
- ▶ Using **restricted intermediate memory** improves efficiency.
- ▶ But diagonal compensation to prevent breakdown appears **less important** than generally supposed.
- ▶ Our extensive experiments favour use of **global diagonal shifts** (works well provided the problem is well scaled).

Concluding remarks on positive-definite case

- ▶ We have explored the use of diagonal compensation with a **limited memory Tismenetsky** approach.
- ▶ The proposed limited memory Tismenetsky approach has been shown to be **robust and efficient**.
- ▶ Using **restricted intermediate memory** improves efficiency.
- ▶ But diagonal compensation to prevent breakdown appears **less important** than generally supposed.
- ▶ Our extensive experiments favour use of **global diagonal shifts** (works well provided the problem is well scaled).
- ▶ New *IC* code **HSL_MI28**.

What about indefinite problems?

Saddle-point systems $Kx = b$ where

$$K = \begin{pmatrix} A & B^T \\ B & -C \end{pmatrix}$$

- ▶ A is $n \times n$ symmetric positive definite
- ▶ B is rectangular and of full rank
- ▶ C is $m \times m$ ($m \leq n$) symmetric positive semi-definite

Signed Cholesky factorization

$$K = \begin{pmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix} \begin{pmatrix} L_{11}^T & L_{21}^T \\ 0 & L_{22}^T \end{pmatrix},$$

$$A = L_{11}L_{11}^T$$

$$L_{21} = BL_{11}^{-T}$$

$$S = C + L_{21}L_{21}^T = L_{22}L_{22}^T$$

This factorization always exists, without pivoting (although numerical stability is not guaranteed).

Signed Cholesky factorization

Set

$$\mathcal{L} = \begin{pmatrix} L_{11} & 0 \\ L_{21} & L_{22} \end{pmatrix},$$

then

$$P = \mathcal{L}\mathcal{L}^T$$

can be used as a (split) preconditioner.

$\mathcal{L}^{-1}K\mathcal{L}^{-T}$ has two distinct eigenvalues ± 1 , so symmetric Krylov subspace method (MINRES or SYMMLQ) converges in at most two iterations.

Signed Cholesky factorization

In practice,

$$A \approx \tilde{L}_{11} \tilde{L}_{11}^T,$$

$$\tilde{L}_{11}^T \tilde{L}_{21} = B^T$$

$$\tilde{S} = C + \tilde{L}_{21} \tilde{L}_{21}^T \approx \tilde{L}_{22} \tilde{L}_{22}^T$$

- ▶ **Key question:** how to sparsify \tilde{L}_{21}^T ? (need \tilde{S} sparse)
- ▶ We take a different approach that does not limit working to the (1, 1) block and then the (2, 2) block.

Constrained orderings

Consider for now the **complete** factorization.

Constrained ordering: find permutation Q such that QKQ^T can be factorized stably **without** numerical pivoting and **without** modifying entries in K , while keeping sparsity.

Bridson (2007):

- ▶ Divide nodes of adjacency graph of K into two disjoint sets: **A-nodes** that correspond to diagonal entries of A and remaining nodes are **C-nodes**.
- ▶ A **C-node** can only be ordered **after** all its **A-node neighbours** K have been ordered.

Constrained orderings

- ▶ Provided A is definite and B is of full row rank, with this ordering the LDL^T factorization exists, with L unit lower triangular and D diagonal.
- ▶ Moreover, the pivots associated with A -nodes are positive and those associated with C -nodes are negative.
- ▶ Rescaling, $L \leftarrow L|D|^{1/2}$ and $D \leftarrow \text{sign}(D) = \text{diag}(\pm 1)$

This gives a **signed Cholesky** factorization of K .

Is this good approach?

- ▶ **Advantage:** simple modification to Cholesky code, avoiding need for pivoting (static data structures set up during analyse phase).
- ▶ **Disadvantages:**
 - ▶ no pivoting so stability not guaranteed
 - ▶ constrained ordering may lead to more fill (more work and more memory)

Our experiments showed better to use a state-of-the-art **indefinite** direct solver that incorporates **pivoting**.

But can we use it for incomplete signed IC ?

- ▶ Want to adapt our IC code to compute **signed IC** factorization.
- ▶ Need to identify C -nodes and **post process** ordering so C -node ordered only after its A -node neighbours.
- ▶ Allow use of **two** shifts

$$\bar{K} = SQ \begin{pmatrix} A & B^T \\ B & -C \end{pmatrix} Q^T S + \begin{pmatrix} \alpha(1)I & 0 \\ 0 & -\alpha(2)I \end{pmatrix}$$

where S is diagonal scaling, Q permutation (cf. Saunders and Tomlin '96 and also Altman and Gondzio '99 regularization of symmetric indefinite systems).

Early results $C = 0$ GMRES(100) convergence results. $\alpha(1) = 0.0$.

Identifier	<i>lsize</i>	<i>rsize</i>	$\alpha(2)$	<i>fill</i>	efficiency	iters
boyd1	20	20	0.0	0.82	3.75×10^6	7
d_pretok	5	5	2.56×10^{-1}	1.96	4.85×10^7	28
ncvxqp9	10	10	4.19×10^3	3.98	2.51×10^5	2
sit100	20	20	0.0	3.61	1.72×10^6	14
stokes64	10	10	0.0	2.73	3.19×10^7	157
tuma1	20	20	0.0	5.79	5.86×10^6	20
turon_m	20	20	1.60×10^{-2}	4.36	2.23×10^8	56

Early results $C = 10^{-8}$ (interior-point matrices)GMRES(100) convergence results. $lsize = rsize = 10$.

Identifier	$\alpha(1)$	$\alpha(2)$	<i>fill</i>	efficiency	iters
c-55	0.01	0.64	2.08	5.30×10^7	117
c-68	0.01	2.56	2.31	3.21×10^7	44
c-69	0.01	0.64	2.35	5.68×10^7	70
c-70	0.01	0.64	2.32	6.09×10^7	72
c-71	0.01	0.04	2.17	8.42×10^7	83
c-big	0.01	0.64	2.63	3.85×10^8	109

Future directions

- ▶ Results suggest we can use a signed *IC* approach
but we want to do better.
- ▶ Developing limited memory **incomplete LDL^T** factorization, again using intermediate memory that is discarded.
- ▶ This will involve using **2×2 pivots** for stability (shifts will not be needed).
- ▶ **Note:** Recent work by Chen and Liu to develop **SYM-ILDL** code, based on the work of Li and Saad '05.
This is giving some good results.
- ▶ Also want to explore effects of **matching orderings** (used by Hagemann and Schenk '06 for incomplete factorizations).



Thank you!

HSL_MI28 is available as part of HSL 2013.

Technical Reports RAL-P-2013-004 and RAL-P-2013-005
(to appear in SISC and TOMS)

Supported by EPSRC grant EP/I013067/1

Grant Agency of the Czech Republic Project No. P201/13-06684