

fspai-1.1 – Factorized Sparse Approximate Inverse Preconditioner

Thomas Huckle
Matous Sedlacek

2011 – 09 – 10

Technische Universität München
Research Unit Computer Science V
Scientific Computing in Computer Science

Contents

1	Abstract	2
2	fspai-1.1 Usage	2
2.1	Installation	2
2.2	How to run fspai-1.1	2
2.3	Options	3
2.4	Usage examples	7
3	FAQ	8
3.1	What does FSPAI stand for?	8
3.2	Why does fspai-1.1 taking forever to compute the preconditioner?	8
3.3	Which parameter setting is a good trade-off between performance and accuracy of the approximation?	8
3.4	Does FSPAI always work?	8
3.5	When I reduce ϵ the preconditioner does not really improve, why?	8
3.6	Caching does not improve the computation time, why?	8
3.7	Which dimension do the system matrices may have?	8
3.8	Which input matrices fspai-1.1 was tested with?	9
3.9	On which architectures fspai-1.1 was tested on and how many processors have been used so far?	9

4	Links	9
4.1	Contacts	9
4.2	Used tools	9
4.3	References	10

1 Abstract

Given a sparse symmetric positive definite matrix A , the FSPAI preconditioner computes a sparse approximate inverse lower triangular matrix $L_M \approx L_A^{-1}$, where L_A is the Cholesky factor of $A = L_A^T L_A$. The matrix $M = L_M L_M^T$ can be used as an approximate inverse preconditioner for iterative solvers, like, for example, PCG.

Based on an initial a priori chosen sparsity pattern, FSPAI automatically captures promising indices and updates the sparsity pattern of the preconditioner according to [4, 5]. The algorithm proceeds until $\|L_A L_M - I\|_2 < \epsilon$. By varying ϵ the user may control the quality and the cost of the preconditioner. A very sparse preconditioner is cheap to compute, but may not lead to significant improvements, whereas L_M becomes too expensive to compute if it becomes too dense.

The sequential version, which can be compiled without MPI, uses an own (P)CG implementation if a solver process is requested. The parallel version uses the highly scalable (P)CG solver from the HYPRE [6] package.

2 fspai-1.1 Usage

2.1 Installation

In its current version of fspai-1.1, there is no common configure script and no common installation package at the moment. Adapt the `Makefile` in the `src` directory and start the compilation process by typing `make` in the `src` directory to start compiling the sourcecode.

For compiling the sequential version fspai-1.1 requires Lapack [7], BLAS [1], CXSparse [3], and UFconfig [11]. The parallel version additionally requires HYPRE [6] and MPI [9]. To compile for a parallel environment use the compiler flag `ENV_MPI`, i.e., compile with `-DENV_MPI`.

2.2 How to run fspai-1.1

- **Sequential version:**

```
./fspai-1.1 <matrix.mtx> <pattern.mtx|-diag|-L> [options]
```

- **Parallel (MPI) version:**

```
mpirun -np <n> ./fspai-1.1 <matrix.mtx> <pattern.mtx|-diag|-L> [options]
```

The **mandatory parameters** are:

<code>matrix.mtx</code>	The SPD system matrix in <i>Matrix Market</i> [8] format.
<code>pattern.mtx -diag -L</code>	Path to start pattern file in <i>Matrix Market</i> format or <code>-diag</code> for generating a diagonal start pattern automatically without any file or <code>-L</code> for generating a start pattern which prescribes the lower triangular pattern of A itself for the preconditioner. No file required in these two cases.
<code>n</code>	Number of processors to use in MPI environment for a parallel computation of L_M .

For getting help on the shell use the `-h` option, e.g., `./fspai-1.1 -h`.

2.3 Options

Optional parameters can be used after the mandatory input in any order, always `-optional_parameter <number|path>`.

Optional parameters are:

<code>-h(elp)</code>	<p>Parameter descriptions to shell <i>Usage:</i> <code>./fspai-1.1 -h(elp)</code> <i>Default:</i> Don't display help Print the fspai-1.1 usage and optional parameters to shell.</p>
<code>-ep</code>	<p>ϵ tolerance <i>Default:</i> 0.4 <i>Choice:</i> <code>ep</code> ≥ 0.0 <code>ep</code> controls the quality of the approximation of L_M to the inverse of L_A. Lower values of <code>ep</code> lead to more computational work, more fill-in and usually to better preconditioners.</p>
<code>-ns</code>	<p>Max. number of improvement steps per column <i>Default:</i> 5 <i>Choice:</i> <code>ns</code> ≥ 0 Every column of the inverse L_A^{-1} is approximated in several improvement steps (pattern updates). The quality of the approximation is determined by ϵ. If the approximation could not reach the accuracy of ϵ after <code>ns</code> steps, fspai-1.1 will use the best approximation so far.</p>

-mn	<p>Max. number of new nz candidates per step <i>Default:</i> 5 <i>Choice:</i> $mn \geq 0$</p> <p>For each improvement step there will be chosen <code>mn</code> new indices for the current pattern to compute a better approximation in the next step. If no new candidates can be found, the iteration aborts and the approximation will finish.</p>
-hs	<p>Hashtable size <i>Default:</i> 6 <i>Choice:</i> $hs \in \{0, 1, 2, 3, 4, 5, 6\}$</p> <p>fspai-1.1 uses a hashtable to cache remote messages and avoid redundant process communication. It is recommended to use larger hashtables to avoid the linear rehashing mechanism. <code>-hs 0</code> forces to switch the hashtable off – no communication will be cached locally.</p> <p><code>-hs 0</code> Don't use any hashtable. <code>-hs 1</code> Hashtable has size 101. <code>-hs 2</code> Hashtable has size 503. <code>-hs 3</code> Hashtable has size 2503. <code>-hs 4</code> Hashtable has size 12503. <code>-hs 5</code> Hashtable has size 62501. <code>-hs 6</code> Hashtable has size 104743.</p> <p>This parameter is not available in the sequential version.</p>
-wp	<p>Write preconditioner L_M to file <i>Default:</i> 1 (true) <i>Choice:</i> $wp \in \{0, 1\}$</p> <p>Whether to write the computed preconditioner into a file or not. The preconditioner will be written in <i>Matrix Market</i> format into the specified output file within the current working directory. Default output file is <code>precond.mtx</code>.</p>
-um	<p>Use mean value <i>Default:</i> 1 (true) <i>Choice:</i> $um \in \{0, 1\}$</p> <p>Whether to use a mean bound value for augmenting indices during an improvement step or not. During each improvement step new nz candidates will be computed for a better approximation. With this mean value the number of indices will be reduced to get only few candidates with best improvement. Using this recommended option will reduce the computational costs for constructing the FSPAI.</p>

-out	<p>Specify output file <i>Default:</i> <code>precond.mtx</code> <i>Choice:</i> <code>out</code> be any string containing alphanumeric characters Where to write the computed preconditioner. The preconditioner will be written in <i>Matrix Market</i> format to a file named by the specified output string in the current working directory.</p>
-sol	<p>Solver to use <i>Default:</i> <code>1</code> <i>Choice:</i> <code>sol</code> \in $\{0, 1, 2, 3\}$ Which solver is to be used. In a sequential environment only <code>sol</code> \in $\{0, 1, 2\}$ is possible. Note, that a right-hand side has to be defined using the <code>rhs</code> option. -<code>sol</code> <code>0</code> Don't invoke (P)CG to solve system. -<code>sol</code> <code>1</code> Use PCG solver with computed FSPAI preconditioner. -<code>sol</code> <code>2</code> Perform CG without preconditioner and PCG using computed FSPAI. -<code>sol</code> <code>3</code> Perform PCG with computed FSPAI and PCG using ParaSails preconditioner. The ParaSails preconditioner will be constructed via the HYPRE package automatically.</p>
-sol_tol	<p>Solver tolerance <i>Default:</i> 10^{-6} <i>Choice:</i> <code>sol_tol</code> ≥ 0.0 <code>sol_tol</code> controls the quality/accuracy of the solution computed by the (P)CG solver.</p>
-sol_maxit	<p>Max. number of solver iterations <i>Default:</i> <code>1000</code> <i>Choice:</i> <code>sol_maxit</code> ≥ 0 <code>sol_maxit</code> defines the maximum number of iterations (P)CG will perform to compute the solution.</p>
-sol_out	<p>Output file for (P)CG solution <i>Default:</i> <code>solution.mtx</code> <i>Choice:</i> <code>sol_out</code> be any string containing alphanumeric characters Where to write the computed (P)CG solution. The solution will be written in <i>Matrix Market</i> format to a file named by the specified output string in the current working directory.</p>
-rhs	<p>Right-hand side for solver <i>Default:</i> <code>1</code> <i>Choice:</i> <code>rhs</code> \in $\{0, 1\}$</p>

	<p>Specifying the right-hand side for linear system which is to be solved with (P)CG.</p> <p>-rhs 0 Use random vector as right-hand side. If multiple solution processes are requested, i.e., <code>sol > 1</code> is used, it is guaranteed that the same rhs is used in all solver requests.</p> <p>-rhs 1 Use ones vector as right-hand side.</p>
-para_levels	<p>Number of levels for ParaSails preconditioner</p> <p><i>Default:</i> 1</p> <p><i>Choice:</i> <code>para_levels</code> \geq 0</p> <p>When constructing ParaSails preconditioner with the HYPRE package, this option specifies the number of levels to be used. See [2] for more details on this parameter.</p>
-para_thresh	<p>Threshold parameter for ParaSails preconditioner</p> <p><i>Default:</i> 0.1</p> <p><i>Choice:</i> <code>para_thresh</code> \in [0.0; 1.0]</p> <p>When constructing ParaSails preconditioner with the HYPRE package, this option specifies the threshold to be used. See [2] for more details on this parameter.</p>
-para_filter	<p>Filter parameter for ParaSails preconditioner</p> <p><i>Default:</i> 0.05</p> <p><i>Choice:</i> <code>para_filter</code> \in [0.0; 1.0]</p> <p>When constructing ParaSails preconditioner with the HYPRE package, this option specifies the post-thresholding filter parameter to be used. See [2] for more details on this parameter.</p>
-alg	<p>Fspai algorithm to use for computing preconditioner</p> <p><i>Default:</i> 0</p> <p><i>Choice:</i> <code>alg</code> \in {0, 1, 2, 3}</p> <p>It is possible to invoke different algorithmic branches in fspai-1.1 . For some special matrices <code>alg > 0</code> may speedup the computation.</p> <p><code>alg 0</code> use LAPACK to solve subsystems</p> <p><code>alg 1</code> use CXSparse to solve subsystems</p> <p><code>alg 2</code> use Fspai Caching and LAPACK</p> <p><code>alg 3</code> use Fspai Hashing and LAPACK</p> <p>Fspai Caching uses a buffer of fixed size to store computed Cholesky decompositions. Fspai Hashing uses a hash table. Note that this hash table has nothing in common with the option <code>-hs</code>.</p>
-cs	<p>Cache size for Fspai Caching</p> <p><i>Default:</i> 0</p>

Choice: $cs \geq 0$

If Fspai Caching is requested via `-alg 2`, a positive cache size greater than zero is necessary.

2.4 Usage examples

- Starting sequential computation of a FSPAI for the matrix *nos3* [8]. Based on a diagonal start pattern the system matrix will be approximated inversely. With an ϵ tolerance of 10^{-3} there will be a maximum of 10 improvement steps, within each a maximum of 3 new candidates will be added to the current pattern without using a mean value bound. Only the preconditioner is to be computed and no solver is to be invoked.

```
./fspai-1.1 /Matrices/nos3.mtx -diag -ep 0.001 -ns 10 -mn 3 -um 0 -sol 0
```

- Starting parallel computation of a FSPAI on three processors to compute the preconditioner for the matrix *nos3*. Start pattern is the lower triangular sparsity pattern of the system matrix. With an ϵ tolerance of 10^{-2} and a maximum of 5 improvement steps, within each a maximum of 5 new candidates using the mean value bound will be added to the current pattern, the system matrix will be approximated inversely. Both an unpreconditioned CG and a PCG using the computed FSPAI is requested to solve the linear system. The solver tolerance is set to 10^{-7} with a maximum number of 800 iterations. The right-hand side for the linear system is the default one, i.e., a vector of all ones. The Fspai Caching algorithm is to be used with a fixed cache size of 50.

```
mpirun -np 3 opt01 opt02 opt03 ./fspai-1.1 /Matrices/nos3.mtx -L -ep 0.01  
-ns 5 -mn 5 -sol 2 -sol_tol 0.0000001 -sol_maxit 800 -alg 2 -cs 50
```

- Starting parallel computation on five processors to compute the FSPAI for the matrix *bcsstk16.mtx* [8]. Based on a specific start pattern from file, the system matrix will be approximated inversely. A static FSPAI without any improvement steps is invoked and both the PCG using FSPAI and PCG using a ParaSails preconditioner are requested to solve the linear system containing a right-hand side of all zeros. For the construction of the ParaSails preconditioner the number of levels is set to 2 and the threshold parameter is switched off with 0.0.

```
mpirun -np 5 opt01 opt02 opt03 opt04 opt05 ./fspai-1.1  
/Matrices/bcsstk16.mtx /Matrices/startpattern.mtx -ns 0 -sol 3 -para_levels  
2 -para_thresh 0.0 -rhs 0
```

3 FAQ

3.1 What does FSPAI stand for?

FSPAI stands for Factorized SParse Approximate Inverse. The FSPAI algorithm is the factorized analogy to the well-known SPAI algorithm. [4, 5]

3.2 Why does fspai-1.1 taking forever to compute the preconditioner?

Make sure your ϵ tolerance is not too small. You could start with a larger value (e.g., `-ep 0.3`) and reduce it progressively until total execution time starts to increase again.

3.3 Which parameter setting is a good trade-off between performance and accuracy of the approximation?

In general it depends on the system/problem to be preconditioned/solved. However, our experiments revealed that augmenting with more indices in several update steps is more efficient. Depending on sparsity, structure, size, and the initial pattern, update settings can be for instance `ns = 10` with `mn = 3` or `ns = 5` with `mn = 5`. A tolerance of `ep = 0.3` or `ep = 0.4` can suffice to get a satisfying quality.

3.4 Does FSPAI always work?

In principle, yes. Unlike many other preconditioners, FSPAI cannot break down if the matrix A is non-singular. Moreover, if one keeps reducing ϵ and use `-mn 1`, fspai-1.1 will eventually compute the exact inverse L_A^{-1} . This may take a very long time.

3.5 When I reduce ϵ the preconditioner does not really improve, why?

Try to increase the number of improvement steps `-ns` and/or the number of indices to be added during one step `-mn`. If that does not help, try to switch off the mean value bound, i.e., use `-um 0`.

3.6 Caching does not improve the computation time, why?

It is advised to use the caching algorithm only when the system matrix is very structured (band, block, etc...). Try to make the cache larger (e.g. 500) when A is supposed to have many different submatrices. Otherwise a small cache of size 50 should be sufficient.

3.7 Which dimension do the system matrices may have?

There is no restriction to the number of columns of A . fspai-1.1 was tested with lots of matrices of different dimension and density. The largest matrices had approximately $1.6 \cdot 10^7$ columns and rows. Note that the system matrices must be square.

3.8 Which input matrices fspai-1.1 was tested with?

Besides matrices generated on our own, like, for example Laplace matrices, fspai-1.1 was tested for various matrices from Matrix Market [8] and from the University of Florida Sparse Matrix Collection [10].

3.9 On which architectures fspai-1.1 was tested on and how many processors have been used so far?

See the website for the environments fspai-1.1 was tested on. Depending on the cluster, fspai-1.1 was used to solve preconditioned systems with more than 8000 processors. We observed high scalability for all our runtime results.

4 Links

4.1 Contacts

- Thomas Huckle

Department of Computer Science - Research Unit V
Technische Universität München
Boltzmannstr. 3
85748 Garching bei München
Germany

☎ +49-(0)89/289 18 609

FAX +49-(0)89/289 18 607

✉ huckle@informatik.tu-muenchen.de

www http://www5.in.tum.de/wiki/index.php/Univ.-Prof._Dr._Thomas_Huckle

- Matous Sedlacek

Department of Computer Science - Research Unit V
Technische Universität München
Boltzmannstr. 3
85748 Garching bei München
Germany

☎ +49-(0)89/289 18 613

FAX +49-(0)89/289 18 607

✉ sedlacek@in.tum.de

www http://www5.in.tum.de/wiki/index.php/Matous_Sedlacek

4.2 Used tools

The sequential compilation uses:

- **BLAS**: Basic Linear Algebra subroutines [1].

- **LAPACK**: Linear Algebra Package [7].
- **CXSpase**: Extended version of CSparse – Concise Sparse Matrix package [3].
- **UFconfig**: Required by CXSpase [11].

The parallel compilation additionally uses:

- **HYPRE**: High Performance Preconditioners [6].
- **MPI**: Message Passing Interface [9].

4.3 References

References

- [1] ATLAS - AUTOMATICALLY TUNED LINEAR ALGEBRA SOFTWARE. <http://math-atlas.sourceforge.net/>.
- [2] E. CHOW, *A priori sparsity patterns for parallel sparse approximate inverse preconditioners*, SIAM J. Sci. Comput., 21 (2000), pp. 1804–1822.
- [3] CXSPARSE - EXTENDED CONCISE SPARSE PACKAGE. <http://www.cise.ufl.edu/research/sparse/CXSpase/>.
- [4] T. HUCKLE, *Factorized sparse approximate inverses for preconditioning and smoothing*, Selçuk J. of Appl. Math., 1(1) of Sonderheft for 60th Birthday of Prof. Dr. Christoph Zenger (2000), pp. 63–70.
- [5] ———, *Factorized sparse approximate inverses for preconditioning*, The Journal of Supercomputing, 25 (2003), pp. 109–117.
- [6] HYPRE - HIGH PERFORMANCE PRECONDITIONERS. <http://acts.nersc.gov/hypre>.
- [7] LAPACK - LINEAR ALGEBRA PACKAGE. <http://www.netlib.org/lapack/index.html>.
- [8] MATRIX MARKET. <http://math.nist.gov/MatrixMarket/>.
- [9] MPI - MESSAGE PASSING INTERFACE. <http://www-unix.mcs.anl.gov/mpi/>.
- [10] THE UNIVERSITY OF FLORIDA SPARSE MATRIX COLLECTION. <http://www.cise.ufl.edu/research/sparse/matrices/>.
- [11] UFCONFIG. <http://www.cise.ufl.edu/research/sparse/UFconfig/>.