High Performance Factored Sparse Approximate Inverse Preconditioning with FSAIPACK

Roberta Baggio
University of Padova

UNIVERSITÀ DEGLI STUDI DI PADOVA

MATHEMATICAL METHODS and MODELS for ENGINEERING

Objectives

Engineering problems solved by computational methods involve the solution of large linear systems of equations with millions of unknowns. The FSAIPACK package is an advanced software tool that offers an innovative solution technique for Symmetric Positive Definite (SPD) matrices in a parallel computer. In this work an analysis of FSAIPACK performance in real-world engineering problems has been carried out.

The Need for Preconditioning

Nowadays Iterative, or Conjugate Gradient-like solvers, are regarded as the most efficient option for solving large and sparse systems, especially on parallel computers. However these methods prove competitive with direct solvers only if the original system matrix is properly modified using a preconditioner. A preconditioner is a matrix used to transform a system in an equivalent one that is expected to have a better performance with a particular solver. In some sense, the preconditioner M^{-1} should be as similar as possible to A^{-1} though preserving a high sparsity.

$$Ax = b \quad \Rightarrow \quad M^{-1}Ax = M^{-1}b$$

FSAI
A High Performance Preconditioner for SPD Systems

Among the preconditioners developed for a parallel computer FSAI (Factorized Sparse Approximate Inverse) seems to be quite effective and offers promising performances. Its use is particularly recommended for SPD matrices. Indeed, the FSAI algorithm has the property of preserving the positive defiteness property of the preconditioned matrix, thus allowing the use of the Preconditioned Conjugate Gradient (PCG) solver.
This preconditioner is expressed in the factored form: $M^{-1} = G^T G$

Where the factor G is lower triangular and is obtained minimizing the **Frobenius norm**:

$$\|I - GL\|_F \rightarrow min$$

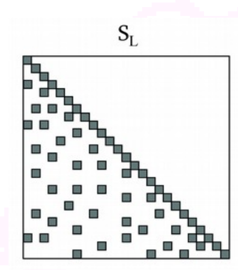

with L the exact Cholesky factor of A.
The sparsity pattern S_L of G is defined a priori by the user. The Frobenius norm minimization leads to the solution of a number of dense linear systems, one for every row of G. Each system can be assigned separately to a different processor, thus making this algorithm particularly suitable for parallel computation.

FSAIPACK: A Software Package for High Performance FSAI Preconditioning

The effectiveness of a FSAI preconditioner is strongly dependent on the selected sparsity pattern. FSAIPACK is a software package that collects several different strategies available in order to compute a high quality sparsity pattern for a FSAI preconditioner, helping the user to improve significantly the computational performance. Basically a sparsity pattern can be both computed statically or dinamically. In the former case the pattern is chosen a priori, while in the latter case it is progressively updated during the preconditioner computation. Every strategy is set selecting some user-specified parameters. Obviously the optimal pattern is strictly problem dependent. For such a reason the parameters have to be defined with reference to the specific problem at hand. Generally speaking, it is necessary to find a good balance between the cost for computing and applying the FSAI preconditioner and the expected reduction of the number of PCG iterations.

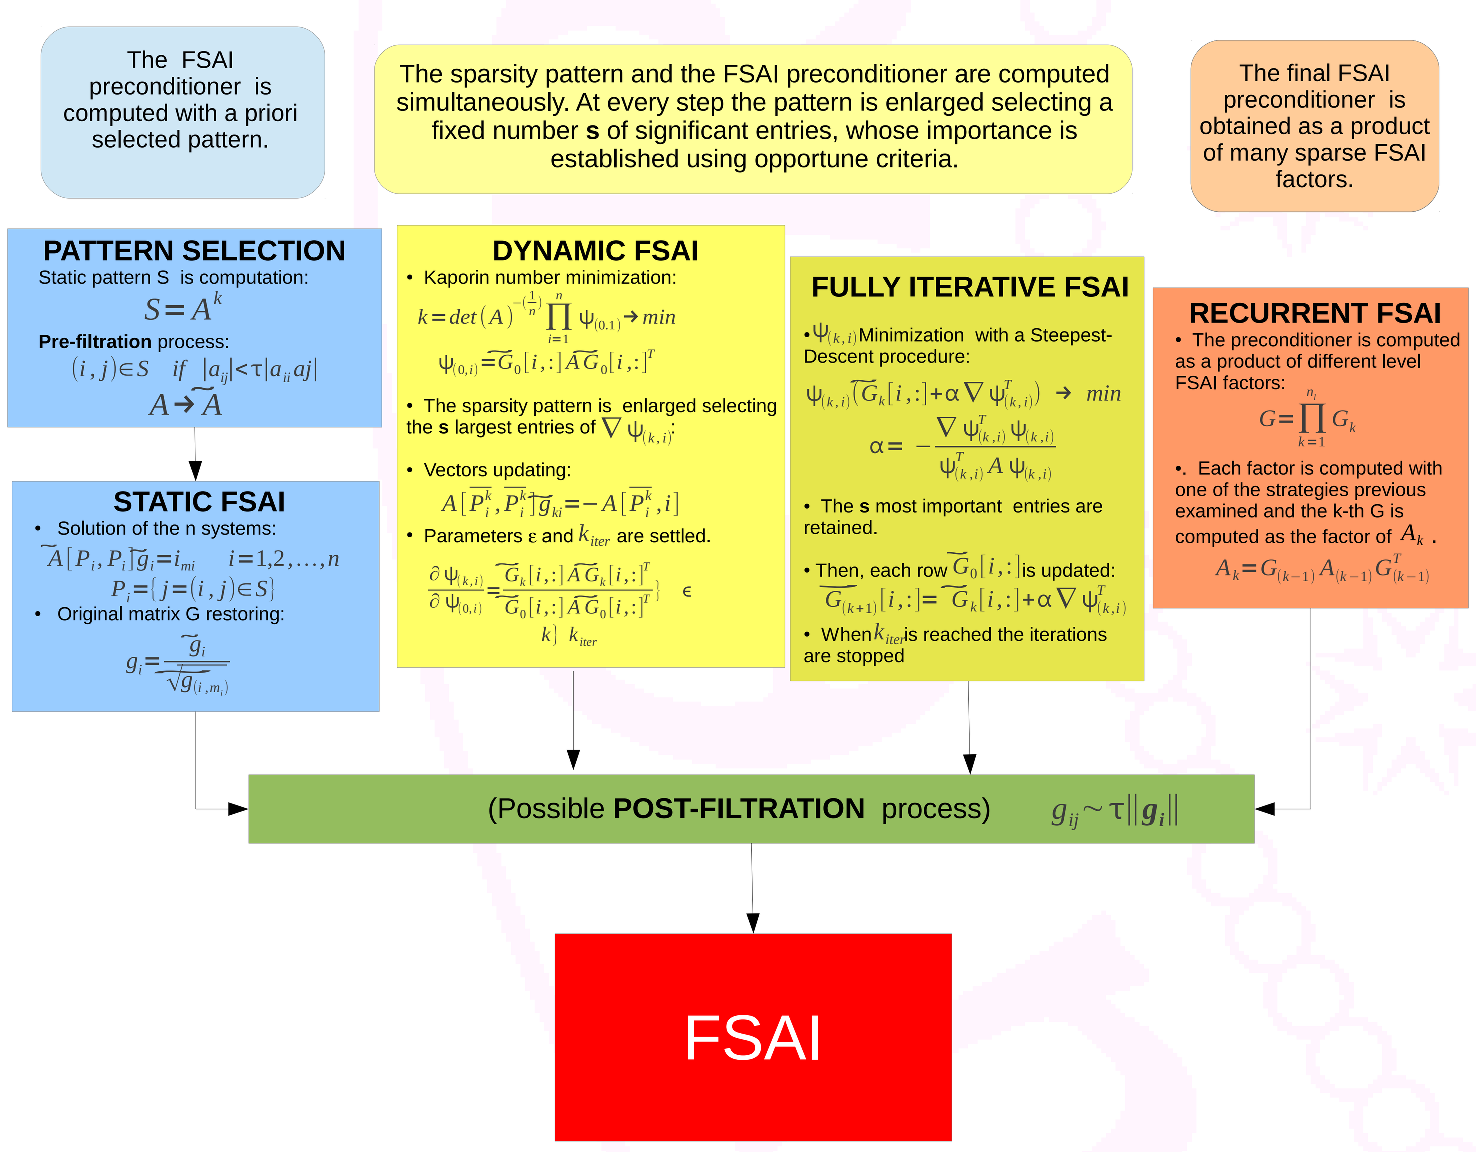

The FSAIPACK software package allows for combining the different algorithms in order to obtain multiple strategies. This approach provides a great flexibility in FSAI computation, making easier to design an optimal preconditioner for the particular problem at hand.
The FSAIPACK package is freely available at: http: //www.dmsa.unipd.it/~janna/software.html.

FSAIPACK Perfect Scalability

A number of tests has been performed on an high performance computer in order to show the FSAIPACK optimal parallelization degree. The machine is FERMI, the most powerful system provided by CINECA. FERMI is an IBM BG/Q computer provided with IBM PowerA2, 1.6 GHz processors, each equipped with 16 cores.
The tests have been performed by doubling each time the number of processors, starting from one up to 16.
The total time spent is nearly halved whenever the number of processors is doubled, thus showing FSAIPACK high scalability. For such a reason this software seems to be a very powerful option on parallel computers.

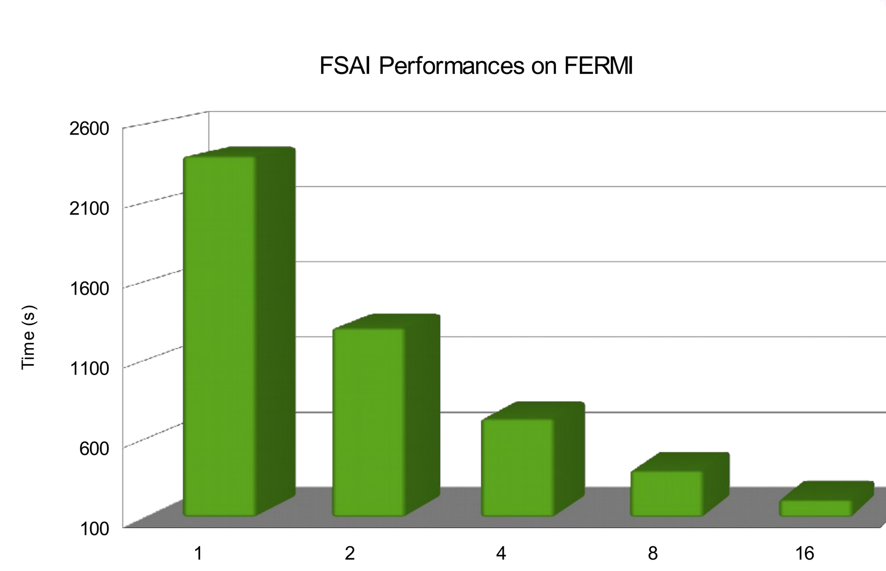

Numerical Results

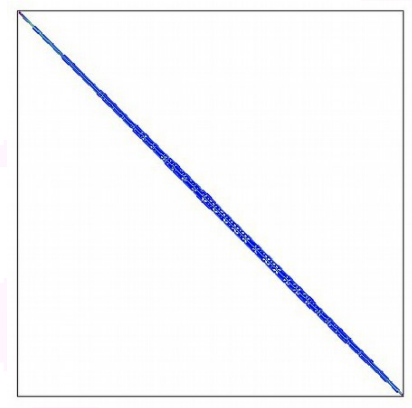

StochF_2030
Number of unknowns=2030336
Number of non-zeroes =30088954

Flow in porous media

$$\frac{\partial}{\partial x}\left(T_x \frac{\partial h}{\partial x}\right) + \frac{\partial}{\partial y}\left(T_y \frac{\partial h}{\partial y}\right) + \frac{\partial}{\partial z}\left(T_z \frac{\partial h}{\partial z}\right) = F(x,y,z)$$

$$\nabla \cdot (\boldsymbol{T} \nabla h) = F(x,y,z)$$

Matrix "StochF_2030" arises from from a tetrahedral Finite Element discretization of an underground aquifer.

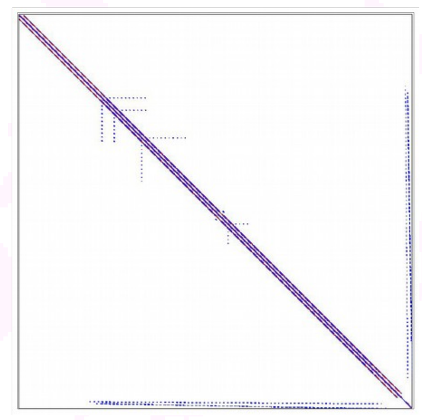

Geomech_2911
Number of unknowns=2911419
Number of non-zeroes = 127729899

Geomechanical problem:

$$\frac{\partial \sigma_x}{\partial x} + \frac{\partial \tau_{xy}}{\partial y} + \frac{\partial \tau_{xz}}{\partial z} = F_x$$
$$\frac{\partial \tau_{yx}}{\partial x} + \frac{\partial \sigma_y}{\partial y} + \frac{\partial \tau_{yz}}{\partial z} = F_y$$
$$\frac{\partial \tau_{zx}}{\partial x} + \frac{\partial \tau_{zy}}{\partial y} + \frac{\partial \sigma_z}{\partial z} = F_z$$

Matrix "Geomech_2911" is related to the geomechanics of a gas reservoir embedded in a sedimentary basin

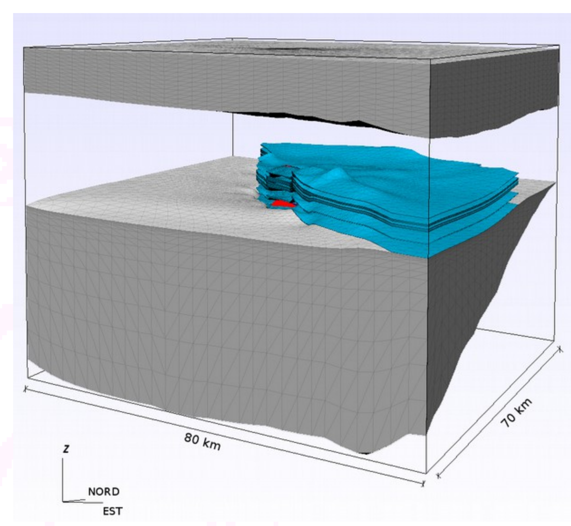

Examples of geomechanical problems discretizations

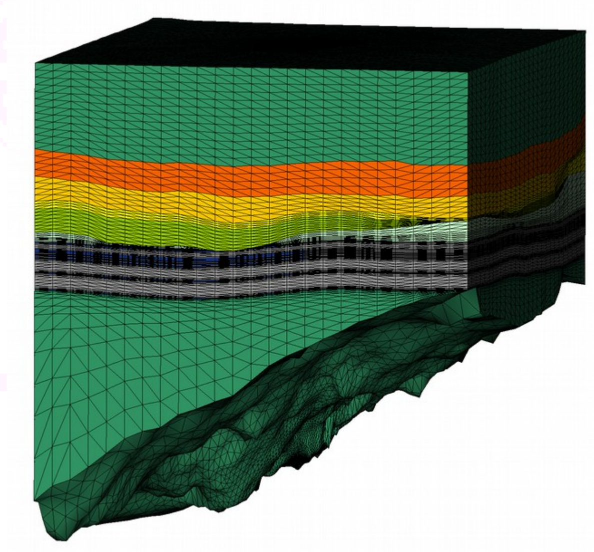

All the following tests have been performed on a on a computer equipped with two quad-core processors Intel Xeon E5-2643@ 3.30GHz, with a 256 GB RAM.

FSAI preconditioned PCG vs Direct solvers

A Comparison has been carried out between parallel PCG preconditioned with Static FSAI and **PARDISO** (PARallel DIrect SOlver), one of the most popular direct solvers for parallel computing.
The tests suggest that direct solvers are not competitive with FSAI-preconditioned PCG on large sparse matrices.

PROBLEM	PARDISO Total time (s)	Static FSAI Total Time (s)
StochF_2030	300.04	21.42
Geomech_2911	1757.52	149.08

Best performances achieved with FSAIPACK

Single-Strategy approach

StochF_2030

Strategy	Preconditioner Time (s)	System Time (s)	**Total Time (s)**	Iteration Number	Preconditioner Density
Static	0.79	20.62	**21.41**	767	0.411
Dynamic	2.04	19.21	**21.25**	552	1.042
Iterative	3.11	18.64	**21.75**	603	0.742

Geomech_2911

Strategy	Preconditioner Time (s)	System Time (s)	**Total Time (s)**	Iteration Number	Preconditioner Density
Static	8.09	140.99	**149.08**	1472	0.343
Dynamic	12,58	64.26	**76.84**	766	0.222
Iterative	9.24	116.82	**126.06**	2017	0.09

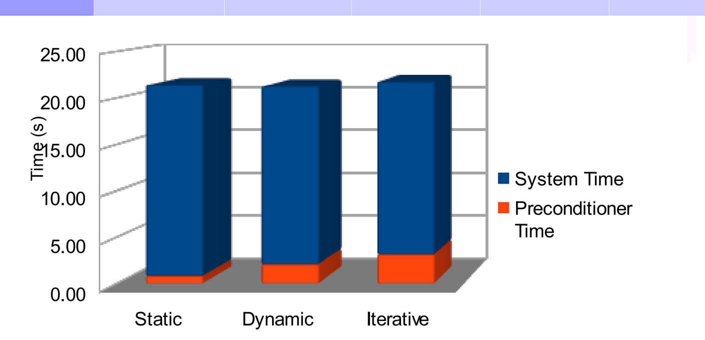

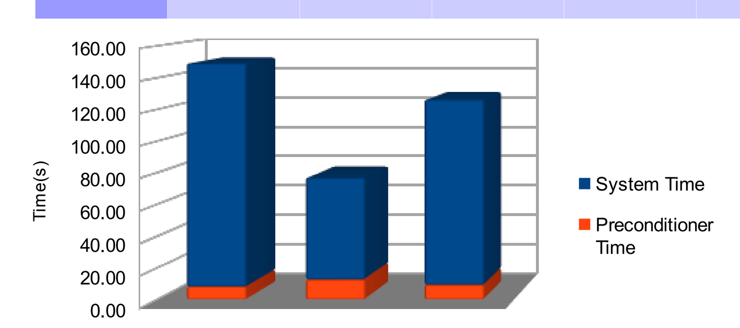

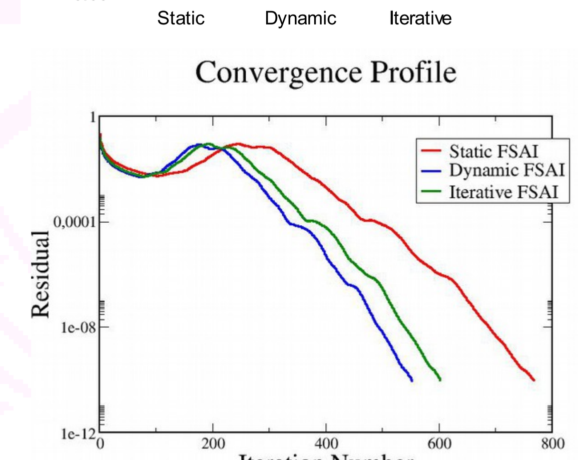

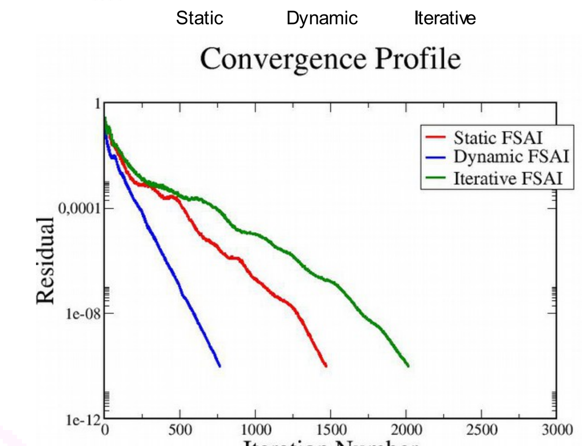

Multiple-Strategy approach on matrix "Geomech_2911"

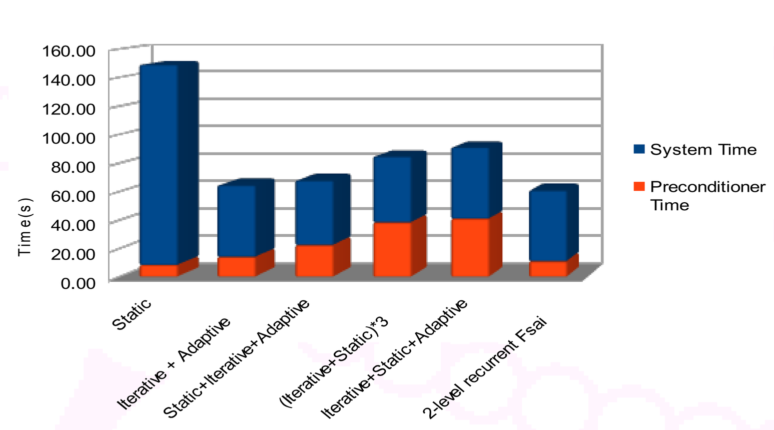

Strategy	Preconditioner Time (s)	System Time (s)	**Total Time (s)**	Iteration Number	Density	**Strategy Improvment**
STATIC+ADAPTIVE	13.87	50.59	**64.47**	570	0.274	**58%**
STATIC+ITERATIVE+ADAPTIVE	22.22	45.54	**67.76**	516	0.323	**54%**
(ITERATIVE+STATIC)*3	38.08	46.52	**84.61**	555	0.246	**38%**
ITERATIVE+STATIC+ADAPTIVE	40.67	50.32	**90.99**	419	0.893	**39%**
2-LEVEL RECURRENT FSAI	10.68	50.19	**60.87**	552	0.268	**59%**

Closing Remarks

- The FSAIPACK solver generally provides very promising results, especially if compared with direct parallel solvers as PARDISO.
- FSAIPACK allows for computing the FSAI preconditioner in a very flexible way, thus making it possible to improve the performance especially in case of ill-conditioned systems.
- The multiple-strategy approach is a powerful technique in order to find the optimal preconditioner for the particular problem at hand.

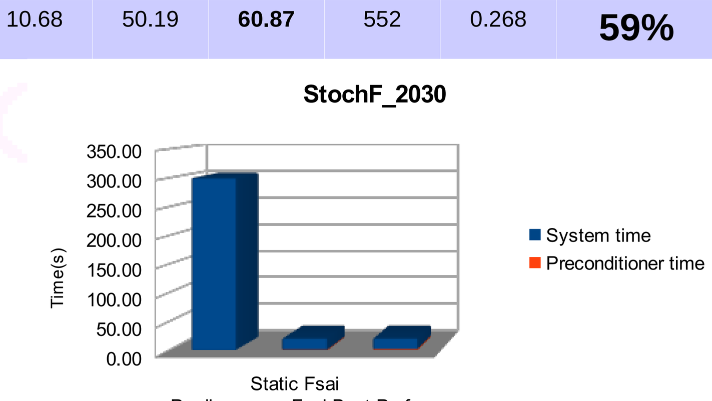

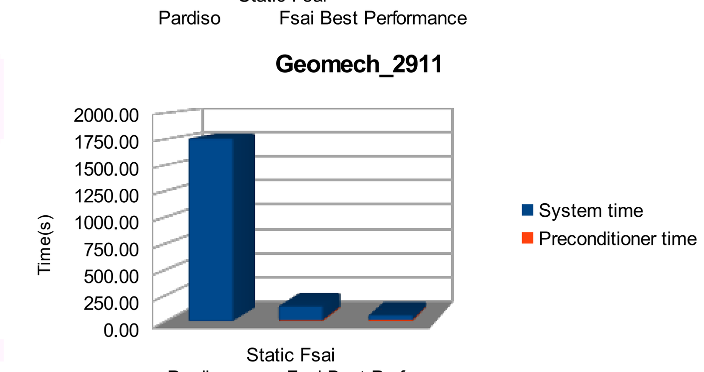

Acknowledgements: Mathematical Methods and Models for Engineering