UNIVERSITÀ DEGLI STUDI DI PADOVA

DMMMSA

Department of Mathematical Methods and Models for Scientific Applications

Block FSAI preconditioning for the parallel solution to large linear systems

Carlo Janna, Massimiliano Ferronato and Giuseppe Gambolati

UNIVERSITÀ DEGLI STUDI DI GENOVA

Due Giorni di Algebra Lineare Numerica

Genova, 16-17 Febbraio 2012

Outline

- Introduction: preconditioning techiques for high performance computing
- Approximate inverse preconditioning: the Block FSAI approach
- Adaptive pattern research for Block FSAI preconditioning
- Numerical results:
  - solution to SPD linear systems by the Preconditioned Conjugate Gradient
- Conclusions
- Work in progress

- Preconditioning is *"the art of transforming a problem that appears intractable into another whose solution can be approximated rapidly"* [Trefethen and Bau, 1997]
- The use of an effective preconditioner is mandatory to achieve convergence with any system or eigenvalue solver used on matrices arising from real-world applications
- Convergence of iterative solvers is accelerated if the preconditioner M^{-1} resembles, in some way, A^{-1}
- At the same time, <u>M^{-1} must be sparse</u>, so as to keep the cost for the preconditioner computation, storage and application to a vector as low as possible
- <u>No rules: even naïve ideas can work surprisingly well!</u>

Introduction

- Algebraic preconditioners: robust tools which can be used knowing the coefficient matrix only, independently of the specific problem addressed
- Incomplete LU factorizations:
  - Incomplete Cholesky with zero fill-in
  - Partial fill-in and threshold value
  - Stabilization techniques

  } **Sequential Computations!**
- Approximate inverses:
  - Frobenius norm minimization
  - Bi-orthogonalization procedure
  - Approximate triangular factor inverse

  } **Parallel Computations!**

→ In real-world problems arising from the discretization of PDEs Stabilized Incomplete LU factorizations are often much more efficient than Approximate Inverses!

The Block FSAI approach

FSAI definition

- Factorized Sparse Approximate Inverse (FSAI): an almost perfectly parallel factored preconditioner [Kolotilina and Yeremin, 1993]

  $$M^{-1} = G^T G$$

  with G a lower triangular matrix such that:

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

  over the set of matrices with a prescribed lower triangular sparsity pattern S_L, e.g. the pattern of A or A^2, where L is the exact Cholesky factor of A

- Computed via the solution of n independent small dense systems and applied via matrix-vector products
- Nice features: (1) ideally perfect parallel construction of the preconditioner; (2) preservation of the positive definiteness of the native matrix

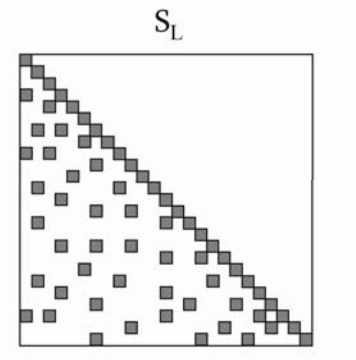

The Block FSAI approach

Block FSAI definition

- The Block FSAI (BF) preconditioner of a Symmetric Positive Definite matrix *A* is a generalization of the FSAI concept:

$$M^{-1} = F^T F$$

with *F* a block lower triangular matrix such that:

$$\| D - FL \|_F \rightarrow \min$$

over the set of matrices with a prescribed lower block triangular sparsity pattern S_{BL}, with *D* an arbitrary block diagonal matrix

- Minimization of the Frobenius norm yields:

$$[FA]_{ij} = \left[DL^T\right]_{ij} \qquad \forall (i, j) \in S_{BL}$$

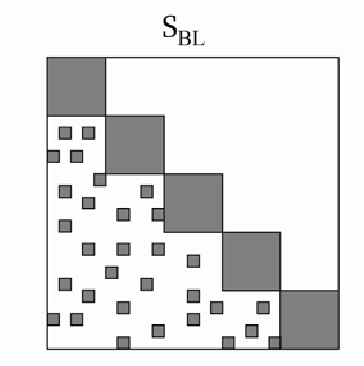

- As *D* is arbitrary, the coefficients of *F* lying in the diagonal blocks can be set arbitrarily, e.g. the diagonal blocks of *F* equate the identity

- In this case F can be computed by solving n independent linear systems with size equal to the number of non-zeroes in each row:

$$A[P_r, P_r]\mathbf{f}_r = -A[r, P_r] \qquad r = 1, \ldots, n$$

with P_r the set of integer numbers:

$$P_r = \{ j : (r, j) \in S_{BL} \}$$

- If A is SPD, existence and uniqueness of the solution for each linear system is guaranteed independently of the set S_{BL}
- Solution to each system is efficiently obtained by a dense factorization routine

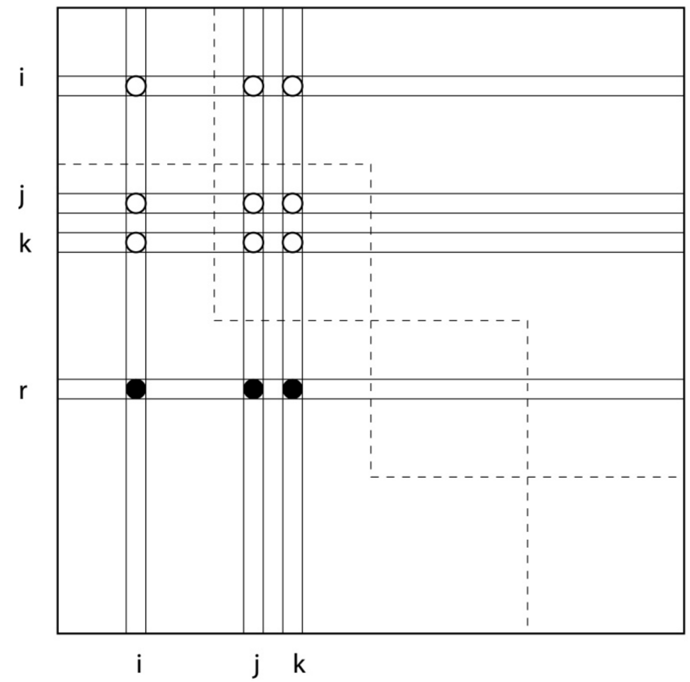

The Block FSAI approach

Block FSAI definition

- In practice, F is such that the largest entries of the preconditioned matrix FAF^T are concentrated in n_b diagonal blocks

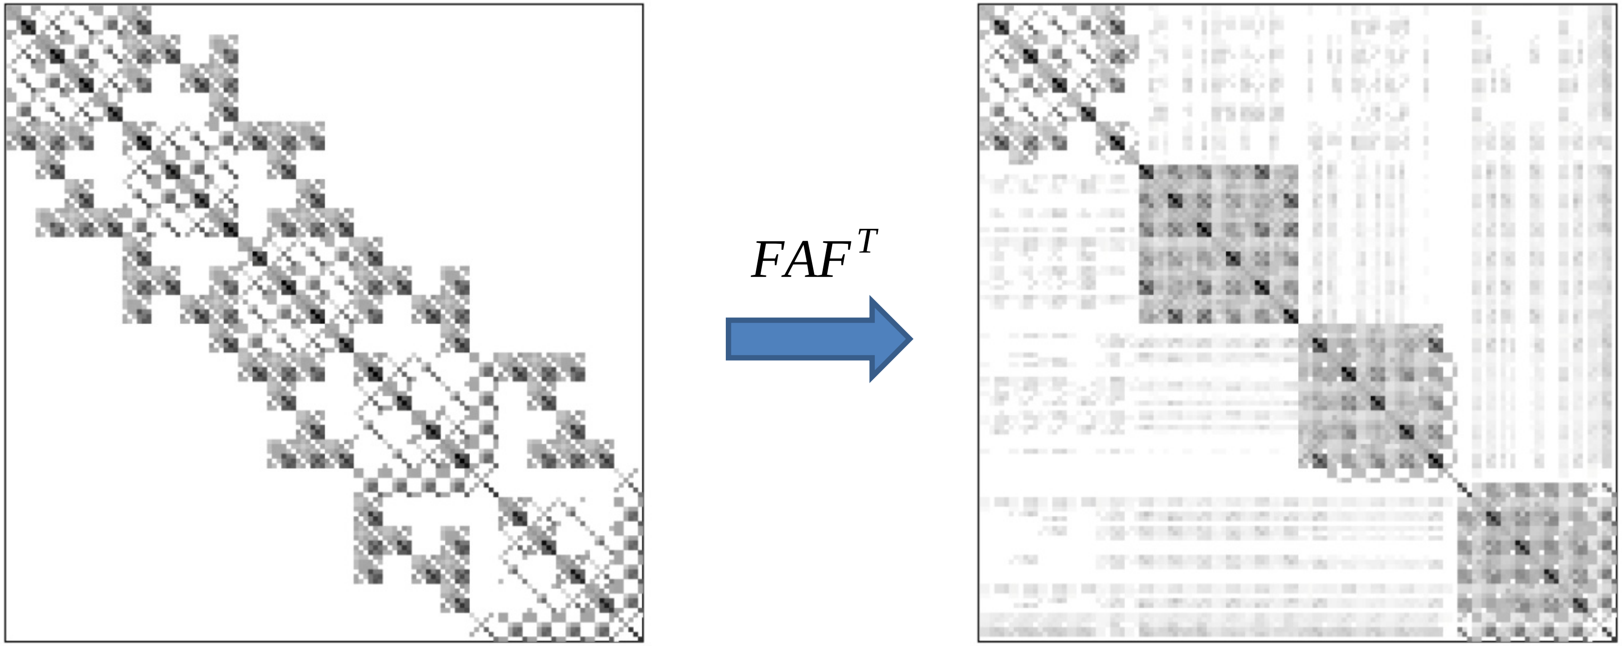

The Block FSAI approach

The BF-IC preconditioner

- As D is arbitrary, FAF^T is not necessarily better than A in an iterative solution method
- To accelerate convergence, FAF^T can be preconditioned again using a block diagonal matrix, e.g. an Incomplete Cholesky (IC) decomposition for each diagonal block B_i of FAF^T:

$$J = J_L J_L^T = \begin{bmatrix} L_1 & 0 & \cdots & 0 \\ 0 & L_2 & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & L_{nb} \end{bmatrix} \begin{bmatrix} L_1^T & 0 & \cdots & 0 \\ 0 & L_2^T & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & L_{nb}^T \end{bmatrix}$$

- The final preconditioned matrix is:

$$J_L^{-1} FAF^T J_L^{-T} = WAW^T$$

where the BF-IC preconditioner reads: $M^{-1} = W^T W = F^T J_L^{-T} J_L^{-1} F$

The Block FSAI approach

Adaptive pattern search

- One of the main difficulties stems from the selection of S_{BL} as an a priori sparsity pattern for F
- Using small powers of A is a popular choice, but for difficult problems high powers may be needed and the preconditioner construction can become quite heavy
- A most efficient option relies on selecting the pattern dynamically by an adaptive procedure which uses somewhat the "best" available positions for the non-zero coefficients
- The Kaporin conditioning number κ of an SPD matrix is defined as:

$$\kappa(A) = \frac{\mathrm{tr}(A)}{n \det(A)^{1/n}}$$

where:

$\kappa(A) \geq 1$ and $\kappa(A) = 1$ iff $\lambda_1 = \lambda_2 = \ldots = \lambda_n$

The Block FSAI approach

Adaptive pattern search

- It can be shown that the Kaporin conditioning number of the BF-IC preconditioned matrix satisfies the following inequality:

$$1 \le \kappa\left(WAW^T\right) \le C \cdot \psi(F)$$

where C is a constant depending on A and $\psi(F)$ is a scalar function depending on the F entries only:

$$\psi(F) = \left(\prod_{i=1}^{n} \left[FAF^T \right]_{ii} \right)^{1/n} = \left(\prod_{i=1}^{n} \left\{ \mathbf{f}_i^T A[P_i, P_i] \mathbf{f}_i + 2\mathbf{f}_i^T A[P_i, i] + [A]_{ii} \right\} \right)^{1/n}$$

- THEOREM. The Block FSAI factor F minimizes $\psi(F)$ for any sparsity pattern S_{BL}.
- <u>Idea:</u> select the non-zero positions in each row of F which provide the largest decrease in the $\psi(F)$ value!

The Block FSAI approach

Adaptive pattern search

- Compute the gradient of each factor of $\psi(F)$:

$$\mathbf{g} = \nabla_{\mathbf{f}_i} \left[F A F^T \right]_{ii}$$

and add to the pattern of the *i*-th row the position corresponding to the largest component of **g**

- Update the row $\mathbf{f}_i$ solving the related dense system
- Stop the selection of new positions when either a maximum number of entries are added to $\mathbf{f}_i$ or the relative decrease of $\psi(F)$ after *k* steps:

$$\Delta_k = \frac{[\psi(F)]_k - [\psi(F)]_{k-1}}{[\psi(F)]_{k-1}}$$

is smaller than a prescribed tolerance

- This gives rise to the Adaptive Block FSAI – Incomplete Cholesky (ABF-IC) preconditioner

Numerical results

ABF-IC preconditioner analysis

- BF-IC with 32 blocks, i.e. 32 processors

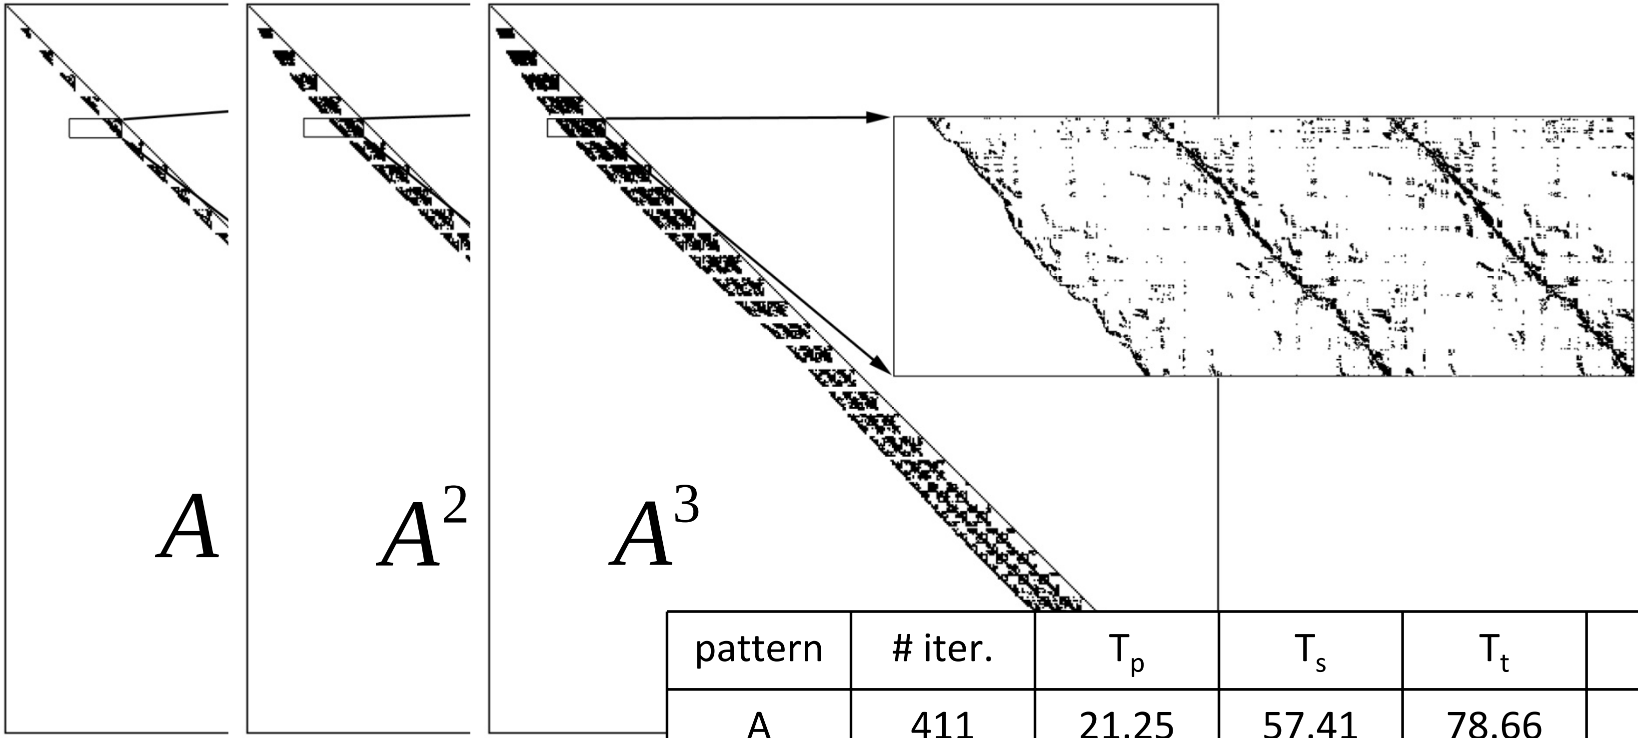

pattern	# iter.	T_p	T_s	T_t	μ_F
A	411	21.25	57.41	78.66	0.19
A^2	241	55.86	57.00	112.86	1.04
A^3	176	319.63	82.51	402.14	**3.13**

Università degli Studi di Padova

Numerical results

ABF-IC preconditioner analysis

- ABF-IC 32 blocks, i.e. 32 processors

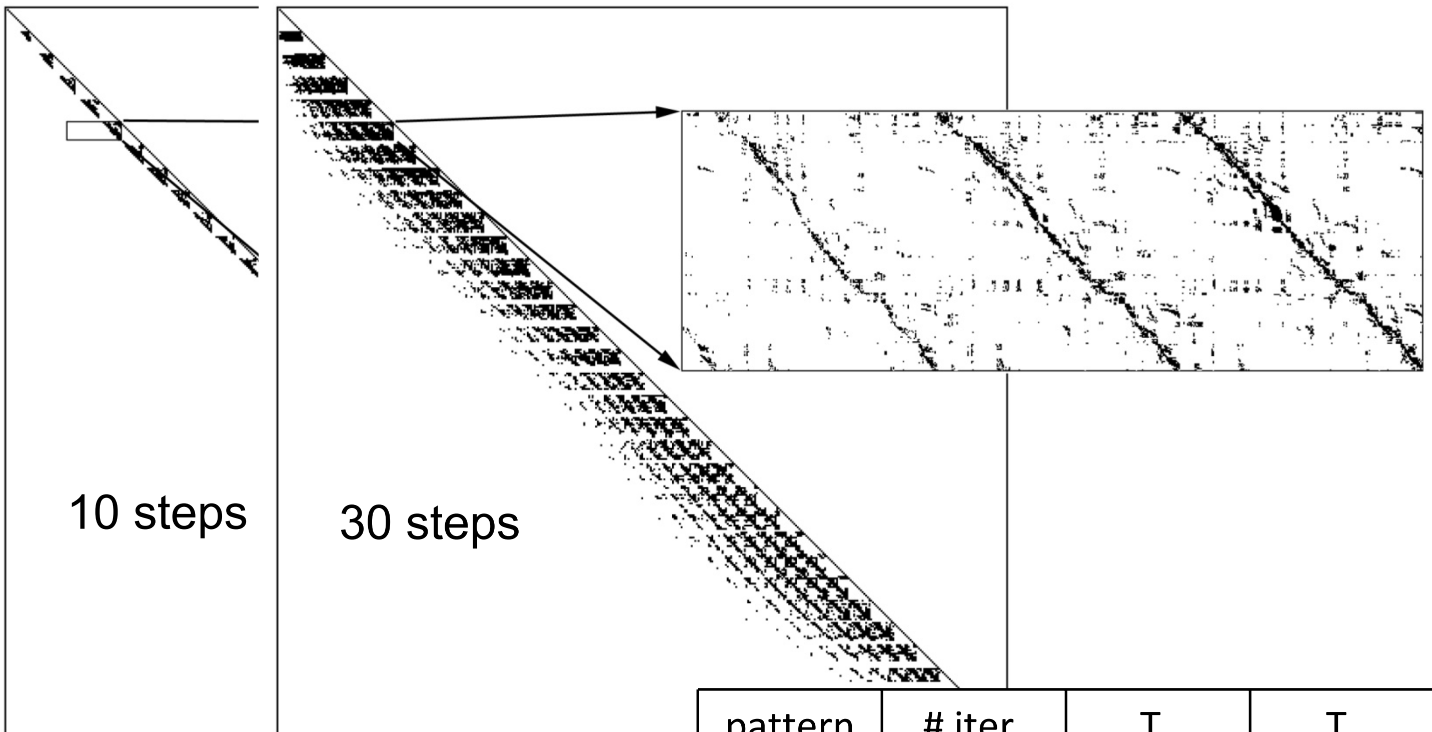

pattern	# iter.	T_p	T_s	T_t	μ_F
10 steps	233	29.87	32.96	62.83	0.17
30 steps	209	100.91	38.27	139.18	**0.46**

Numerical results

Test problems

	Size	# non-zeroes
Fault-639	638,802	28,614,564
StocF-1465	1,465,137	21,005,389
Geo-923	923,136	41,005,206
Mech-1103	1,102,614	48,987,558

Geo-923

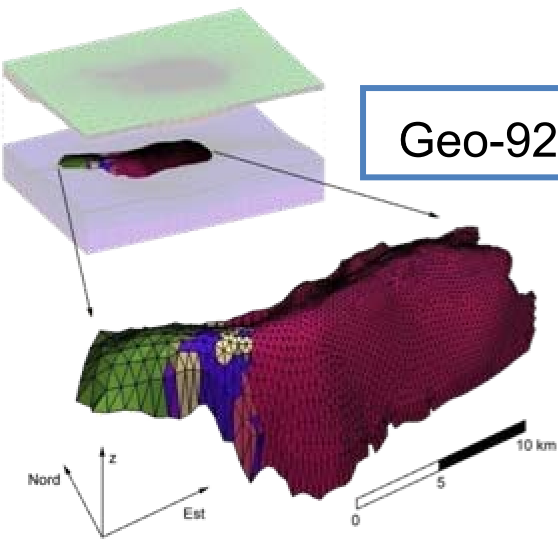

Fault-639

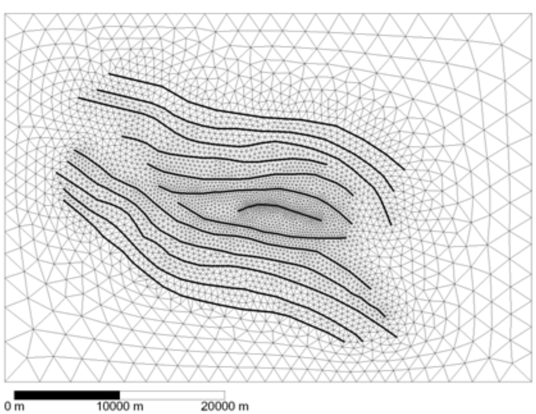

Mech-1103

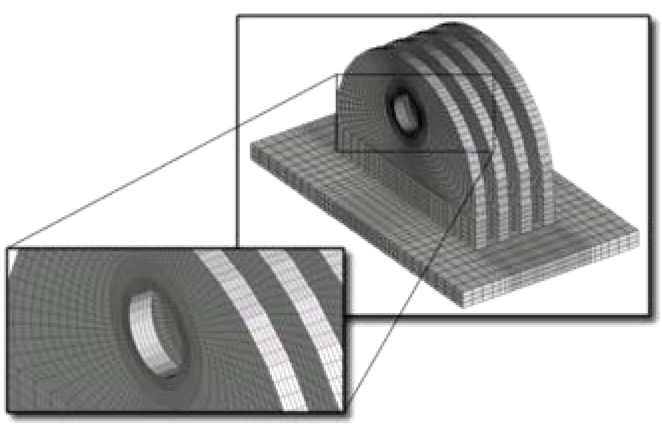

StochF-1465

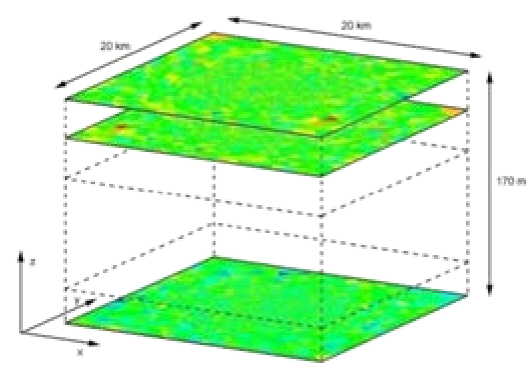

UNIVERSITÀ
DEGLI STUDI
DI PADOVA

Numerical results

Linear system solution with a parallel PCG algorithm

Geo-923

Total wall-clock time [s]

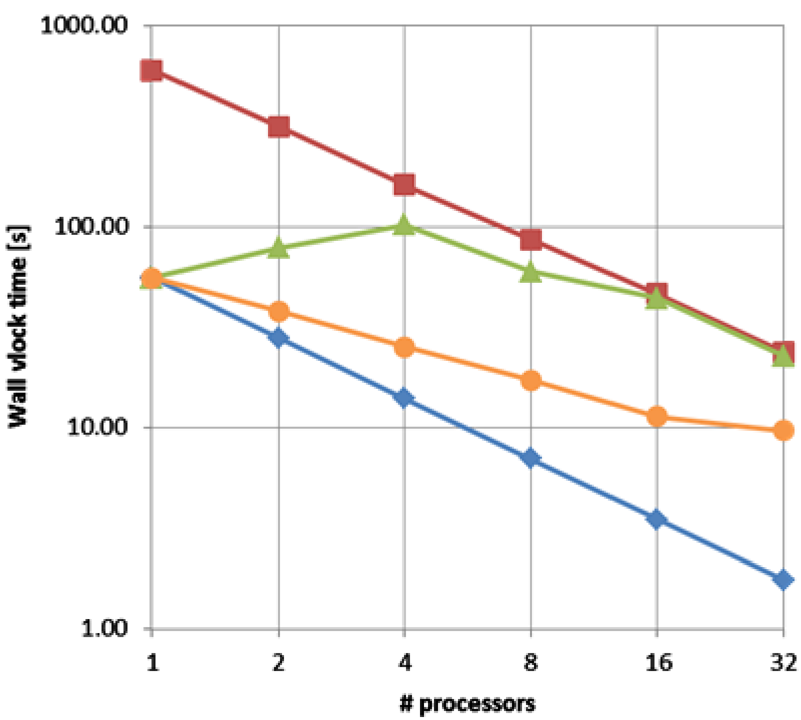

Ratio with Ideal IC

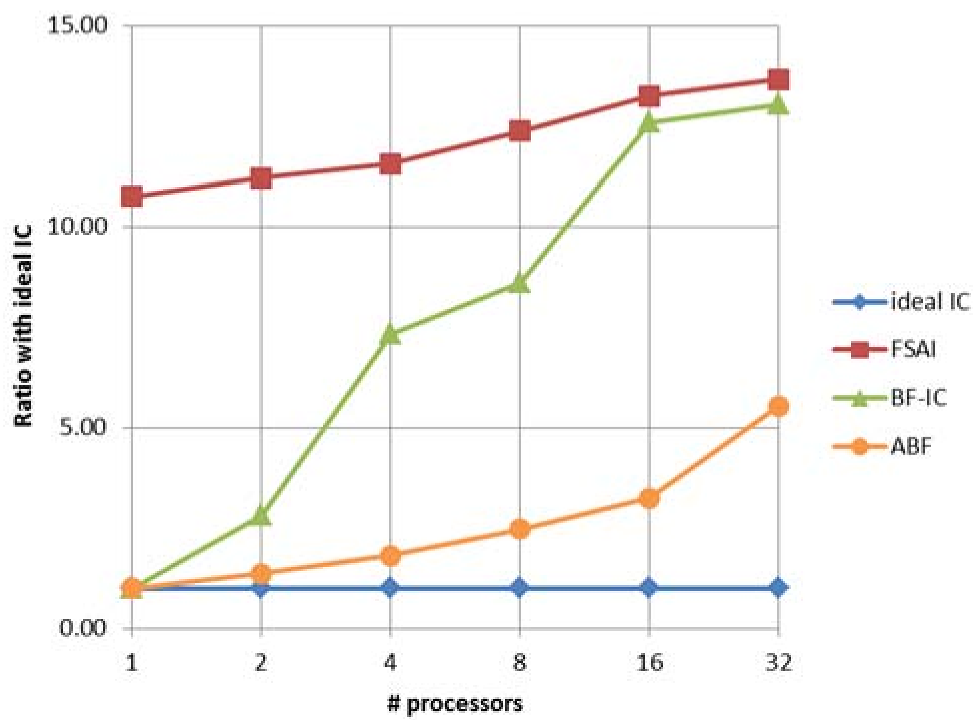

Numerical results

Linear system solution with a parallel PCG algorithm

Fault-639

Total Wall-clock time [s]

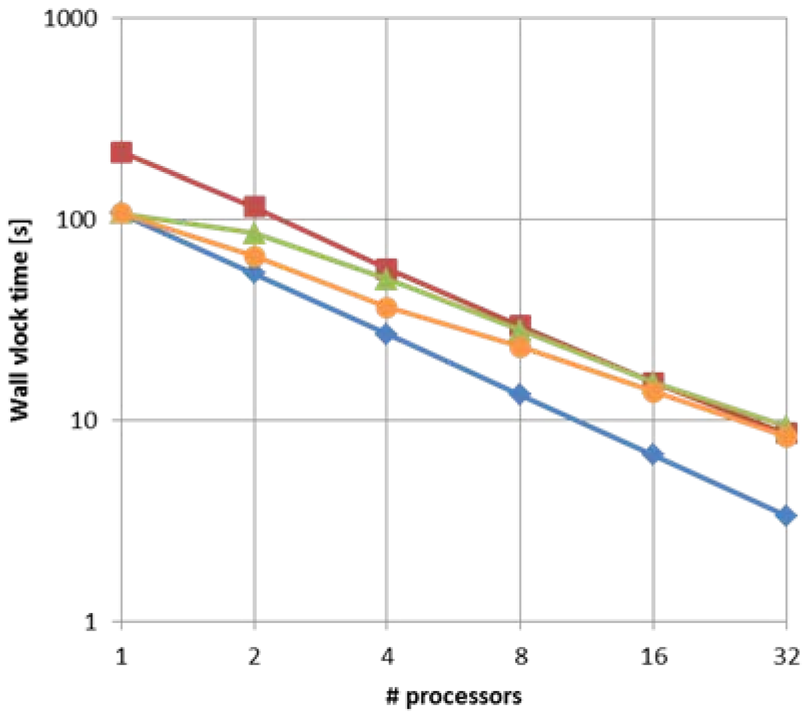

Ratio with Ideal IC

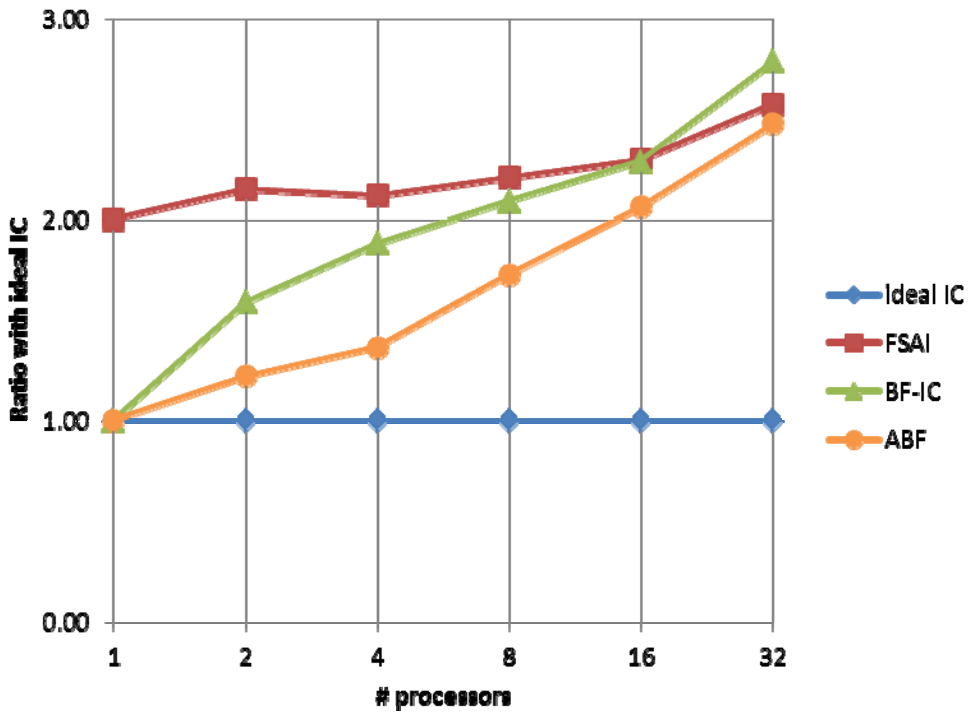

StochF-1465

Total Wall-clock time [s]

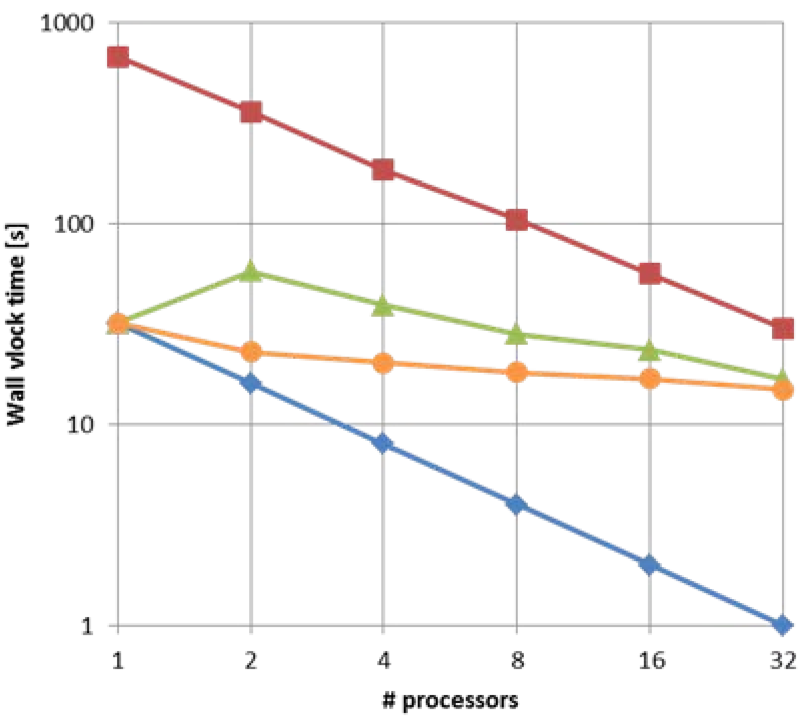

Ratio with Ideal IC

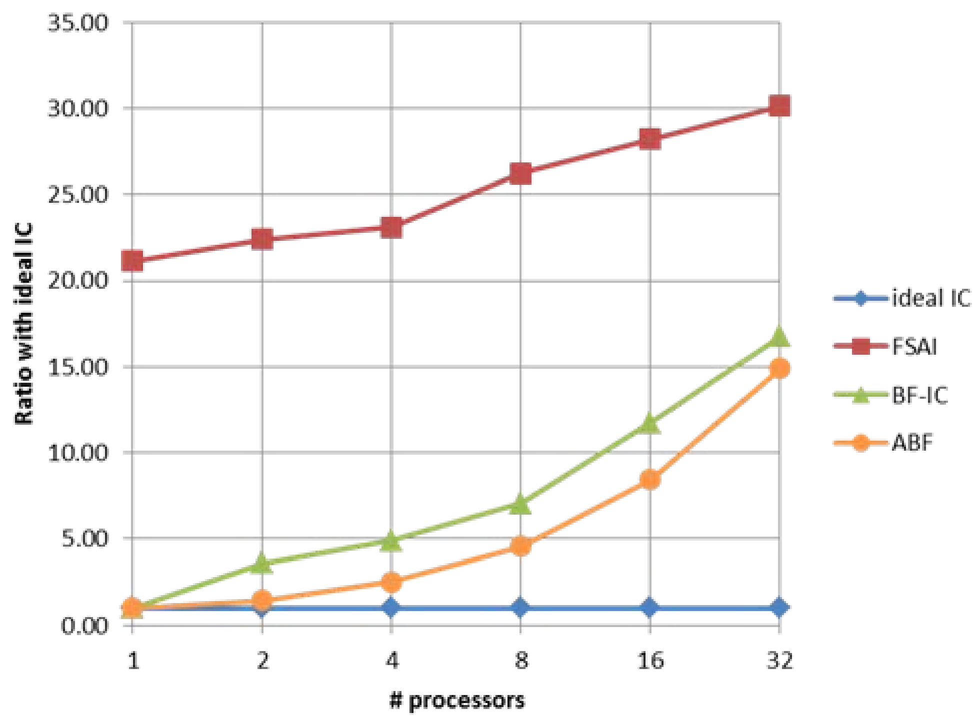

Numerical results

Linear system solution with a parallel PCG algorithm

Mech-1103

Total Wall-clock time [s]

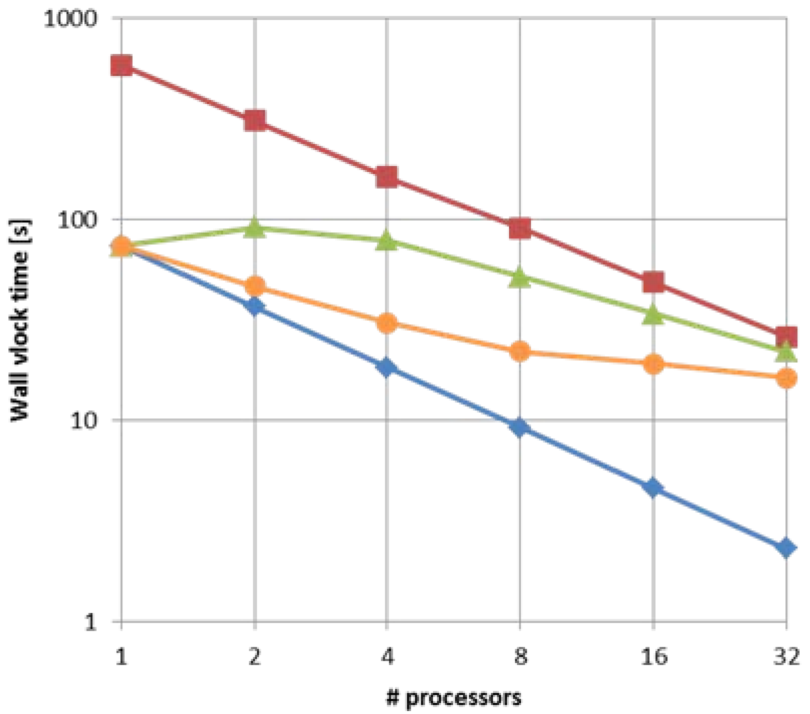

Ratio with Ideal IC

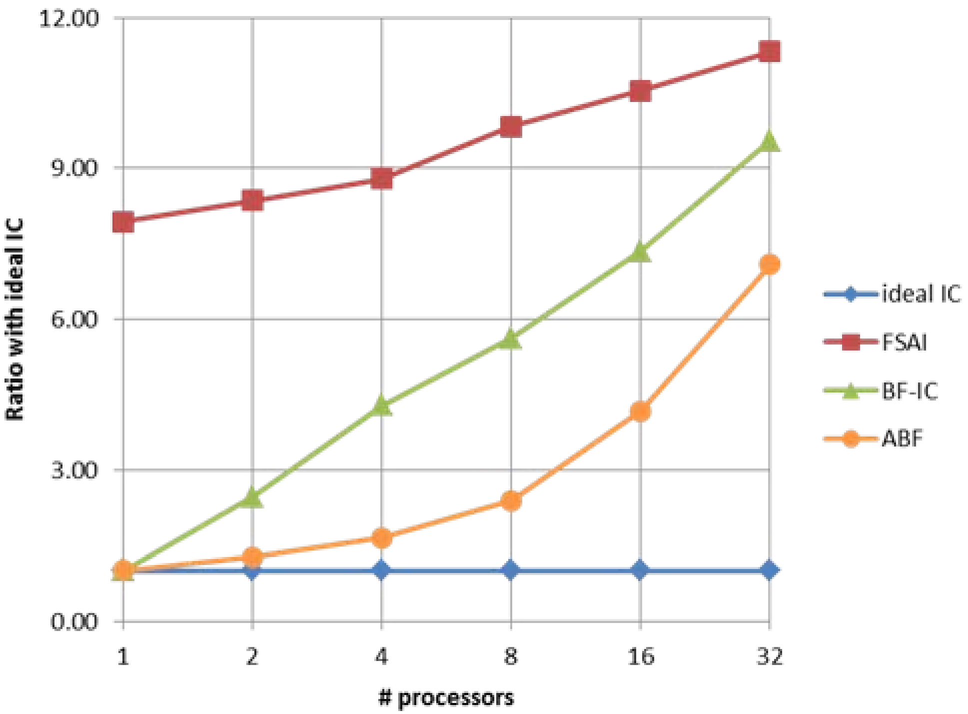

Conclusions

Results…

- The Adaptive Block FSAI – Incomplete Cholesky algorithm is a novel preconditioner coupling the attractive features of both approximate inverses and incomplete factorizations
- The adaptive pattern search can improve considerably the Block FSAI efficiency, especially in ill-conditioned problems
- The main quality of the proposed adaptive search is the capability of capturing the most significant terms belonging to high powers of A (even larger than 10) very efficiently
- ABF-IC has proven equally efficient for solving both SPD linear systems within the PCG algorithm and SPD eigenproblems within the Jacobi-Davidson algorithm
- ABF-IC turns out to be particularly attractive when a relatively small number of processors is used, e.g. with the increasingly popular multi-core processor technology

Conclusions

… and open issues

- Extension of the ABF-IC approach to non symmetric indefinite matrices (non symmetric FSAI is less robust than the SPD one, by contrast ABF-IC appears to be equally robust)
- Improvement of the preconditioner scalability on massively parallel computers coupling ABF-IC with **Domain Decomposition techniques**
- A free OpenMP Implementation of Block FSAI-IC is available online at:

http://www.dmsa.unipd.it/~ferronat/software.html

References:

C. Janna, M. Ferronato, G. Gambolati. *A Block FSAI-ILU parallel preconditioner for symmetric positive definite linear systems.* **SIAM Journal on Scientific Computing**, 32, pp. 2468-2484, 2010.

C. Janna, M. Ferronato. *Adaptive pattern search for Block FSAI preconditioning.* **SIAM Journal on Scientific Computing**, to appear.

M. Ferronato, C. Janna, G. Pini. *Efficient parallel solution to large size sparse eigenproblems with Block FSAI preconditioning*. **Numerical Linear Algebra with Applications**, to appear.

Work in progress

- An analogy can be recognized between Block FSAI preconditioning and Domain Decomposition techniques
- It can be shown that, for a given block subdivision, preconditioning the Schur complement of a Domain Decomposition with Block FSAI is equivalent to apply Block FSAI to the whole system using high accuracy for the internal unknowns
- Following this observation, we separate the internal from the interface unknowns, preconditioning the former with an IC factorization and applying *F* on the latter only.
- This gives rise to a hybrid preconditioner mixing Domain Decomposition and Block FSAI → DD-ABF-IC

Work in progress

	Size	# non-zeroes
Dosso-2911	2,911,419	130,383,395

n_p	ABF-IC				DD-ABF-IC			
	Iter.	Tp	Ts	T_t	Iter.	Tp	Ts	T_t
16	166	13.06	26.49	**39.55**	164	10.94	20.12	**31.06**
32	183	10.31	16.22	**26.53**	205	9.5	12.12	**21.62**
64	204	8.04	9.43	**17.47**	201	6.81	5.99	**12.8**
128	273	5.95	6.9	**12.85**	200	5.01	3.19	**8.2**
256	265	4.64	3.85	**8.49**	223	3.79	2.05	**5.84**

Work in progress

Total Wall-clock time [s]

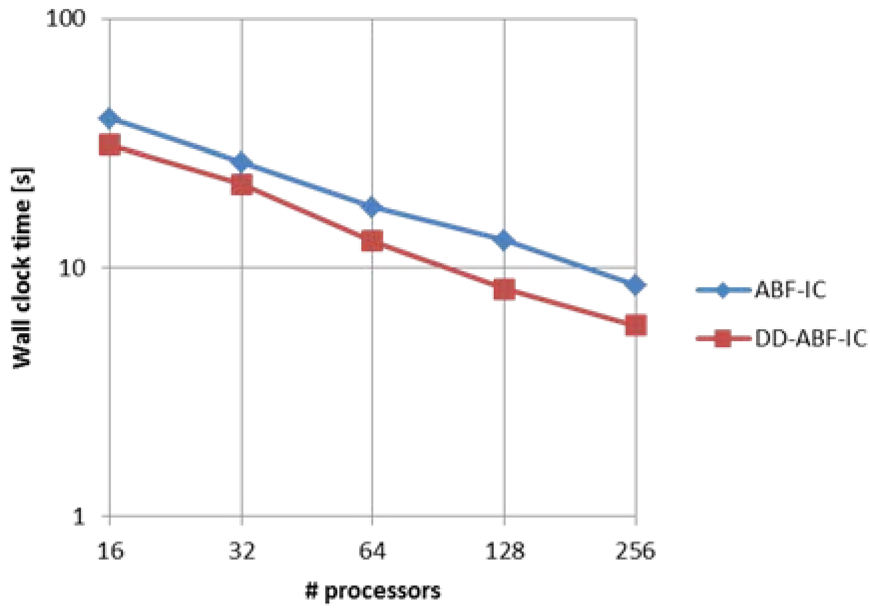

Ratio between DD-ABF-IC and ABF-IC

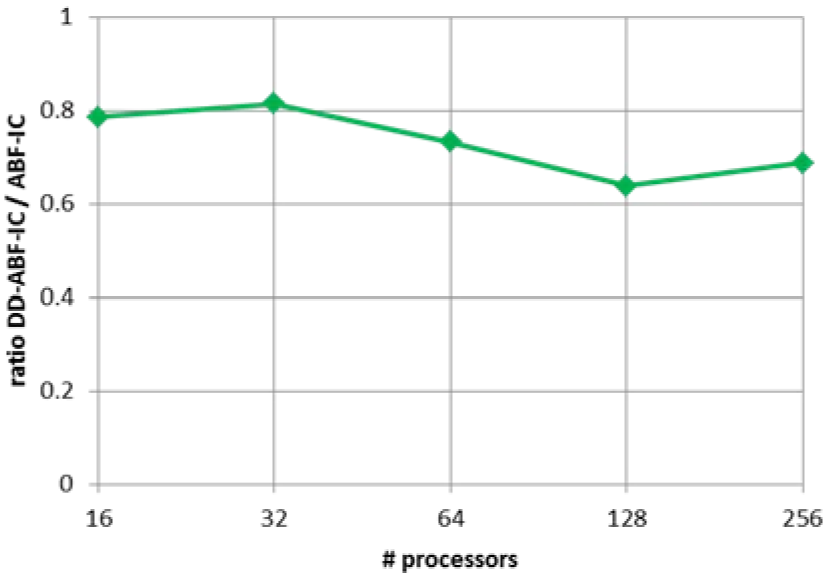

UNIVERSITÀ DEGLI STUDI DI PADOVA

DMMMSA

Department of Mathematical Methods and Models for Scientific Applications

Thank you for your attention

UNIVERSITÀ DEGLI STUDI DI GENOVA

Due Giorni di Algebra Lineare Numerica

Genova, 16-17 Febbraio 2012