ALPINES

Algorithms and parallel tools for integrated numerical simulations

INRIA Rocquencourt, LJLL, UPMC

Members

- X. Claeys (MdC), L. Grigori (DR), F. Hecht (PR), F. Nataf (DR)

Associated members

- J. Burman (MdC Paris-Sud), J. Beauquier (PR Paris-Sud)

Currently 6+3 Phd students, +2 postdoctoral researcher, 1 software engineer

Research group

- Lab. J.-L. Lions, UPMC
  - Mesh generation techniques, FEM

    F. Hecht (Professor)
  - Domain decomposition, PDEs

    F. Nataf (Director of Research CNRS)
  - Boundary element methods, modelling of electromagnetic wave propopagation

    X. Claeys (MdC)

- INRIA
  - Linear algebra and high performance algorithms

    L. Grigori (Director of Research)

- University Paris 11 (associate members)
  - Fault tolerance

    J. Burman (MdC), J Beauquier (Professor)

- Currently 9 Phd students, 2 postdoctoral researchers, 1 software engineer

Methodology

- Mesh generation for parallel computation
  - Exploit the formalism of FreeFEM, a language dedicated to finite elements
  - F. Hecht, F. Nataf

- Solvers for numerical linear algebra
  - Design of domain decomposition and multilevel direction preserving preconditioners
  - High performance computing for boundary element methods
  - X. Claeys, L. Grigori, F. Hecht, F. Nataf

- Computational kernels for numerical linear algebra
  - Design of novel numerical algorithms that minimize communication
  - J. Beauquier, J. Burman, L. Grigori, F. Hecht

- Integration and validation in numerical simulations
  - All members

CA-ILU0
Block Filtering Decomposition (BFD), Nested Filtering Factorization (NFF)

R. Fezzani, P. Kumar, L. Grigori,
R. Lacroix, S. Moufawad, F. Nataf, L. Qu, K. Wang
INRIA, LJLL, UPMC

ANR Petal and Petalh projects
http://petal.saclay.inria.fr/

ILU0 with nested dissection and ghosting

Let α_0 be the set of equations to be solved by one processor
For $j = 1$ to s do
    Find β_j = *ReachableVertices* $(G(U), \alpha_{j-1})$
    Find γ_j = *ReachableVertices* $(G(L), \beta_j)$
    Find δ_j = *Adj* $(G(A), \gamma_j)$
    Set $\alpha_j = \delta_j$
end

Ghost data required:
$x(\delta)$, $A(\gamma,\delta)$,
$L(\gamma,\gamma)$, $U(\beta, \beta)$

$\Rightarrow$ Half of the work performed on one processor

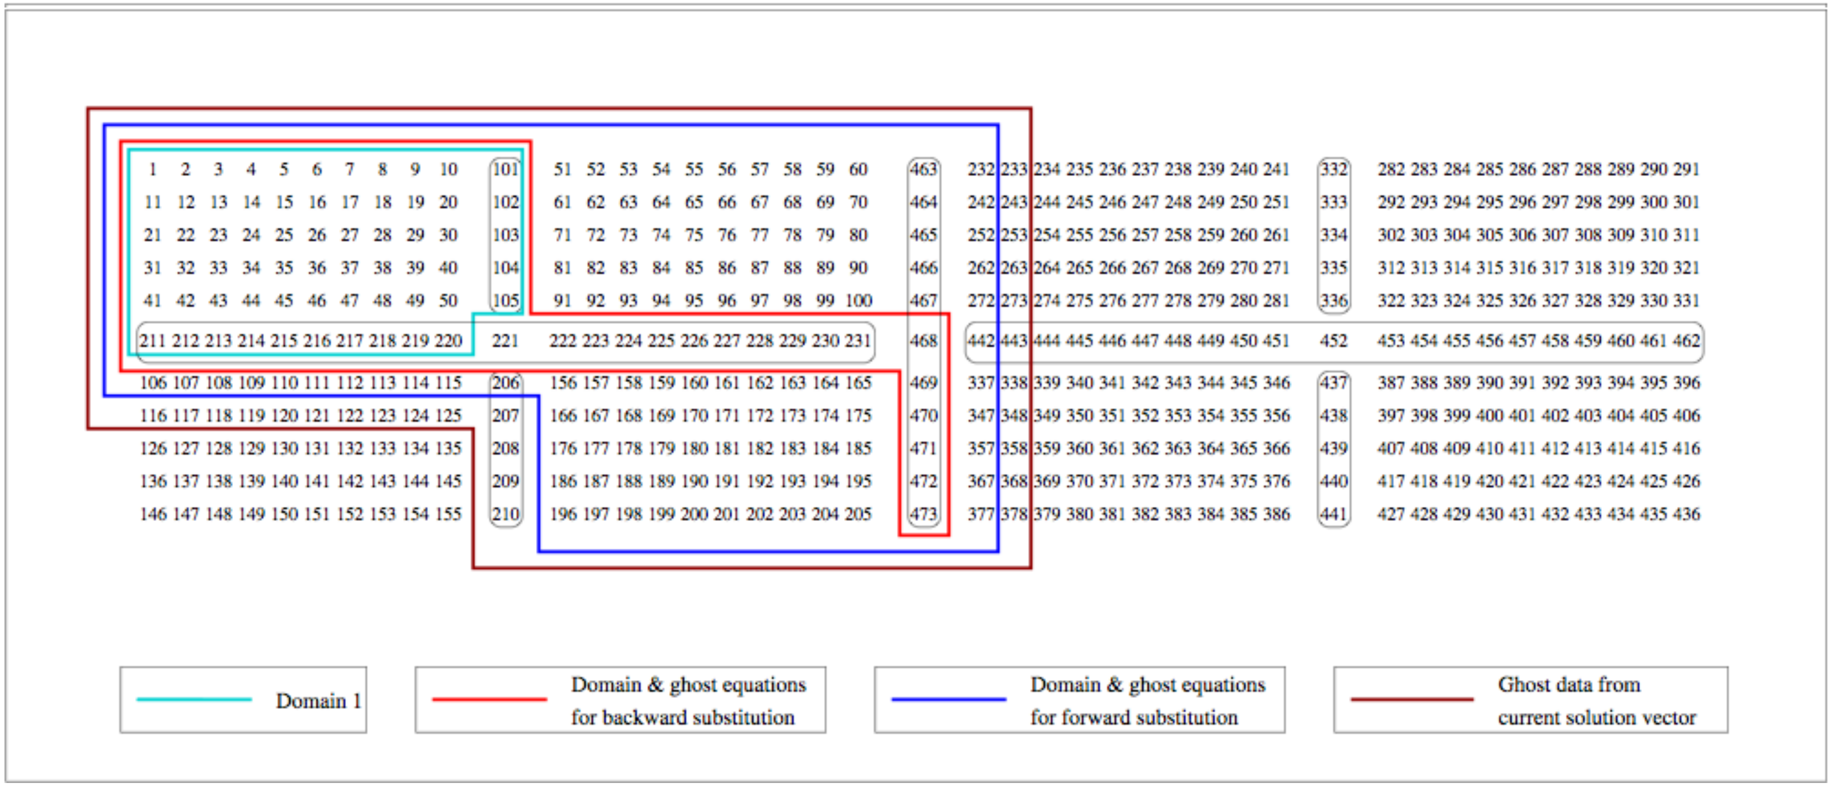

5 point stencil on a 2D grid

CA-ILU0 with AMML reordering and ghosting

- Reduce volume of ghost data by reordering the vertices using Alternating Min-Max Layers (AMML) reordering:
  - First number the vertices at odd distance from the separators
  - Then number the vertices at even distance from the separators
- CA-ILU0 computes a standard ILU0 factorization

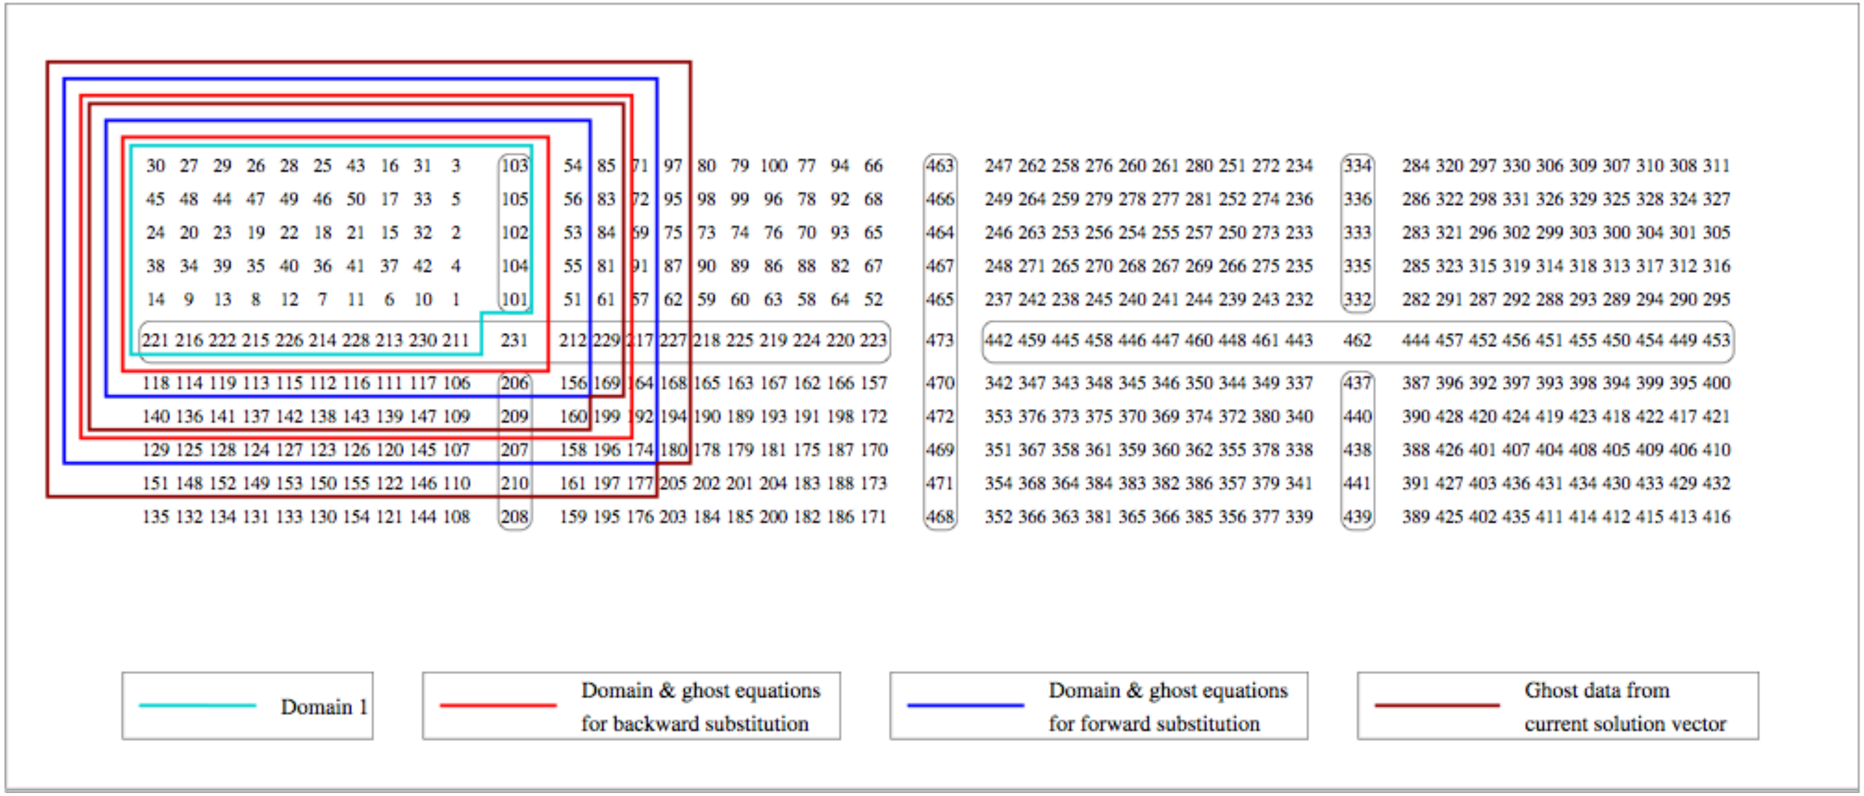

5 point stencil on a 2D grid

Comparison with Block Jacobi

- Block Jacobi is another preconditioner which does not require communication
- Tests for a boundary value problem (Achdou, Nataf), 40x40x40 grid

$$-div(\kappa(x)\nabla u) = f \quad in\,\Omega$$

$$u = 0 \quad on\,\partial\Omega_D$$

$$\frac{\partial u}{\partial n} = 0 \quad on\,\partial\Omega_N$$

$$\Omega = [0,1]^3, \partial\Omega_N = \partial\Omega \setminus \partial\Omega_D$$

κ jumps from 1 to 10^3

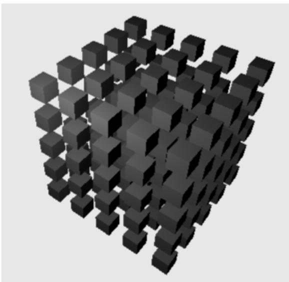

Methods tested:

- Natural ordering NO+ILU0
- CA-ILU0 - kway+AMML(1)+ILU0
- Block Jacobi using LU - BJ+ILU0
- Block Jacobi using ILU0 - BJ-ILU0

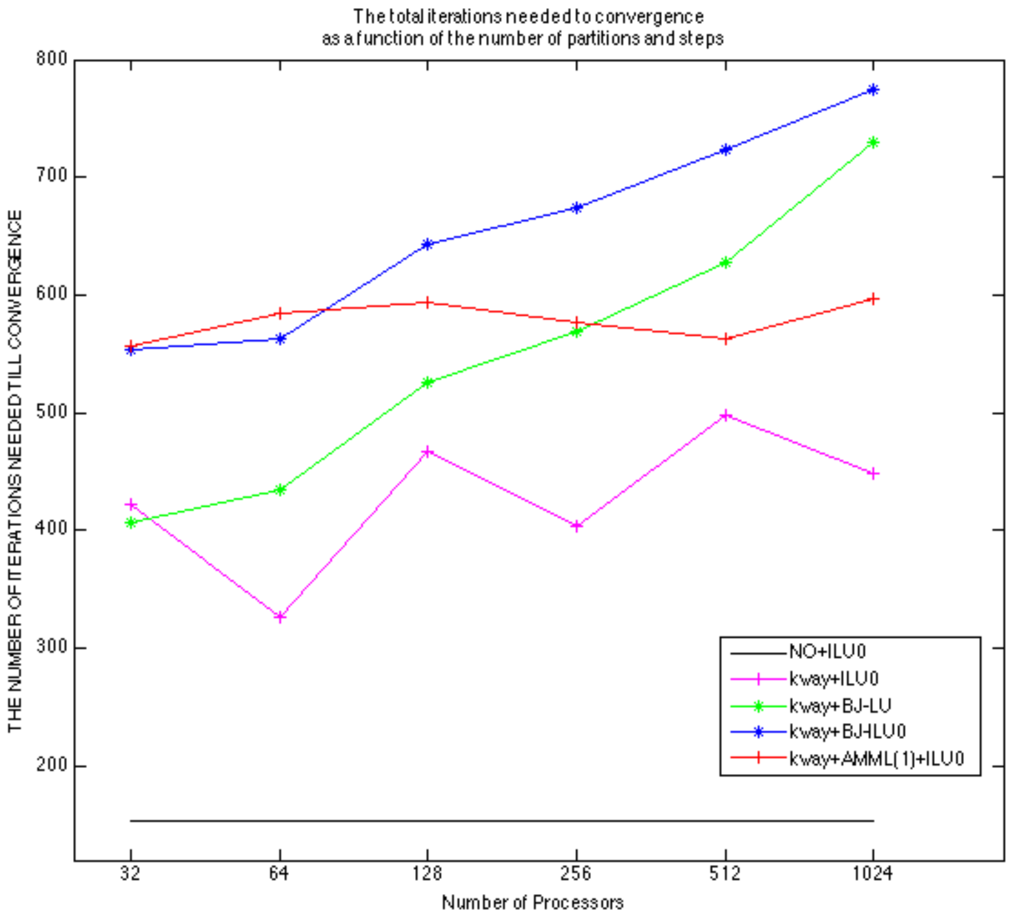

Motivation

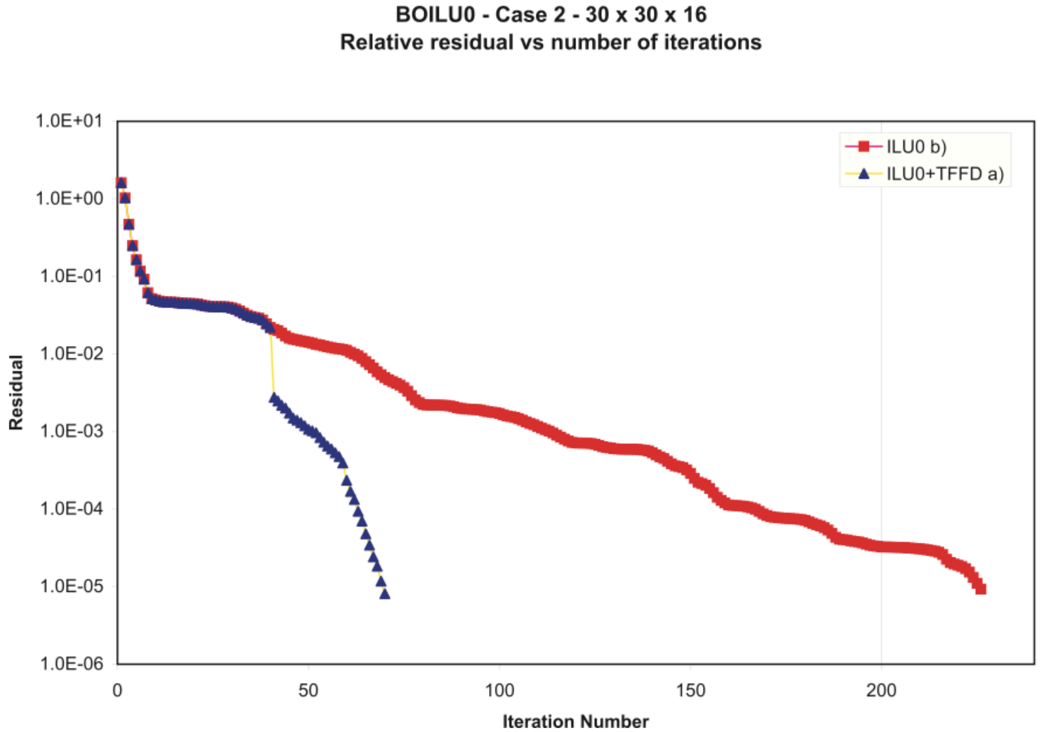

Source: Y. Achdou, F. Nataf

Solve

$$\mathbf{M^{-1} A x = M^{-1} b}$$

Incomplete LU has plateaus in the convergence

Often due to the presence of few low eigenvalues

Filtering factorization

- Preconditioner **M** satisfies a filtering property for input **A** and set of vectors **T**

  $\mathbf{MT = AT}$ or $\mathbf{T^T M = T^T A}$
- Filtering vectors **T** are chosen to improve the convergence
- Complementary with incomplete LU factorization

Preserving directions of interest

- Pointwise approximate factorization satisfying a row-sum criteria, Dupont, Kendall, and Rachford (1968), Gustafsson (1978)
  - Improves the condition number of the preconditioned matrix for matrices arising from finite difference approximation of second order elliptic equations
- Nested factorization, Appleyard, Cheshire (1983)
  - If $t^T r_0 = 0$, then at any iteration $t^T r_k = 0$, this ensures a mass conservation property
- Filtering factorization, Wagner, Wittum (1997), Achdou, Nataf (2001)

- Direction preserving semiseparable approximation of SPD matrices, Gu, Li, Vassilevski (2010)
  - If the near null-space of the original fine grid matrix is preserved, then view the preconditioner as a coarse discretization matrix
  - Conditioning analysis performed by Napov, components dropped are orthogonal to components preserved
- Multigrid methods
  - Bootstrup AMG (Karsten Kahl)

Arbitrary matrices

- Let A be partitioned into a block matrix of size N x N
- The square diagonal blocks are not necessarily of a same size

$$A = \begin{pmatrix} A_{11} & \cdots & A_{1N} \\ \vdots & \ddots & \vdots \\ A_{N1} & \cdots & A_{NN} \end{pmatrix}$$

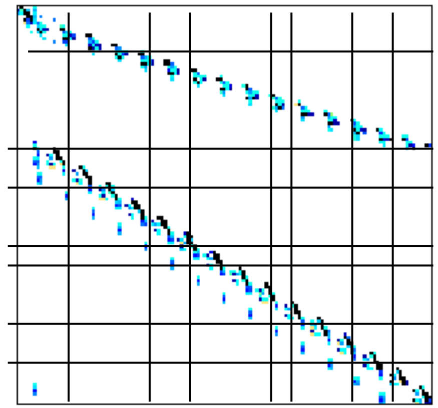

- The generalization of filtering preconditioner to arbitrary matrices is a step forward towards parallel computation

Exact factorization of arbitrary matrices

- An exact block LDU factorization of A is:

$$A = (L + D)D^{-1}(D + U)$$

$$= \begin{pmatrix} D_{11} & & & \\ L_{21} & D_{22} & & \\ \vdots & \ddots & \ddots & \\ L_{N1} & \cdots & L_{N,N-1} & D_{NN} \end{pmatrix} \cdot \begin{pmatrix} D_{11}^{-1} & & & \\ & D_{22}^{-1} & & \\ & & \ddots & \\ & & & D_{NN}^{-1} \end{pmatrix} \cdot \begin{pmatrix} D_{11} & U_{12} & \cdots & U_{1N} \\ & D_{22} & \ddots & \vdots \\ & & \ddots & U_{N-1,N} \\ & & & D_{NN} \end{pmatrix}$$

`

- Let $C=L+D+U$. Each block of L, D, U is computed as:

$$C_{ij} = \begin{cases} A_{ij}, i = 1, j = 1 \\ A_{ij} - \sum_{k=1, L_{ik} \neq 0, U_{kj} \neq 0}^{\min(i,j)-1} L_{ik} D_{kk}^{-1} U_{kj}, i > 1, or j > 1 \end{cases}$$

Block Filtering Decomposition (BFD)

- Let t be a filtering vector. A BFD preconditioner M is written as:

$$M = (\overline{L} + \overline{D})\overline{D}^{-1}(\overline{D} + \overline{U})$$

- Let $\overline{C} = \overline{L} + \overline{D} + \overline{U}$. The blocks of M are computed with the following formula, where $i,j = 1 \ldots N$

$$\overline{C}_{ij} = \begin{cases} A_{ij}, & i = 1 \quad or \quad j = 1 \\ A_{ij} - \displaystyle\sum_{k=1, L_{ik} \neq 0, U_{kj} \neq 0}^{\min(i,j)-1} \overline{L}_{ik} F_{kj} \overline{U}_{kj}, & i > 1 \quad or \quad j > 1 \end{cases}$$

where F_{kj} is a sparse approximation such that

$$\overline{L}_{ik} F_{kj} \overline{U}_{kj} t_j = \overline{L}_{ik} \overline{D}_{kk}^{-1} \overline{U}_{kj} t_j$$

Suitability for parallel computation

- Partition the matrix using nested dissection, thus enabling parallelism

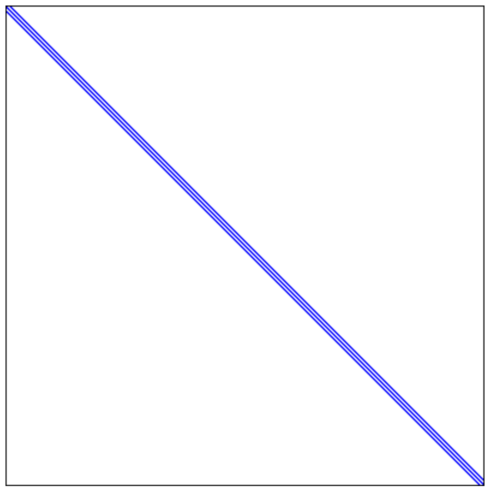

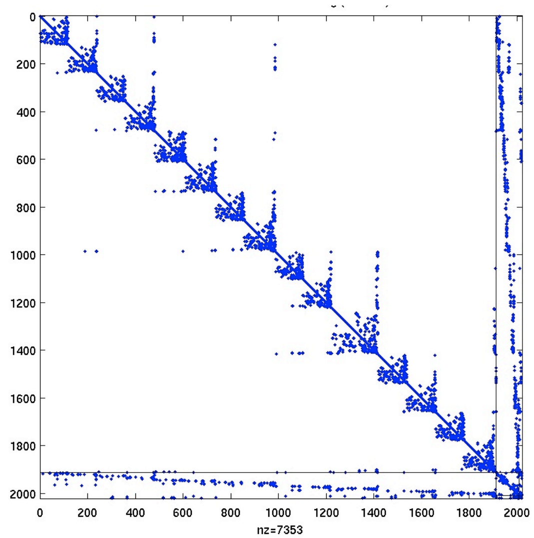

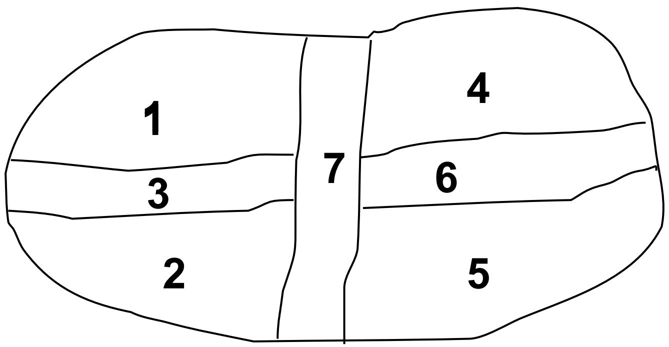

$$
\begin{pmatrix}
A_{11} & & A_{13} & & & & A_{17} \\
 & A_{22} & A_{23} & & & & A_{27} \\
A_{31} & A_{32} & A_{33} & & & & A_{37} \\
 & & & A_{44} & & A_{46} & A_{47} \\
 & & & & A_{55} & A_{56} & A_{57} \\
 & & & A_{64} & A_{65} & A_{66} & A_{67} \\
A_{71} & A_{72} & A_{73} & A_{74} & A_{75} & A_{76} & A_{77}
\end{pmatrix}
$$

Results for a boundary value problem

- SKY (provided by Achdou, Nataf), discretized on a 400x400x400 grid (64 millions unknowns, 447 millions nonzeros)

$$-div(\kappa(x)\nabla u) = f \quad in\,\Omega$$
$$u = 0 \quad on\,\partial\Omega_D$$
$$\frac{\partial u}{\partial n} = 0 \quad on\,\partial\Omega_N$$
$$\Omega = [0,1]^3, \partial\Omega_N = \partial\Omega \setminus \partial\Omega_D$$

κ jumps from 1 to 10^3

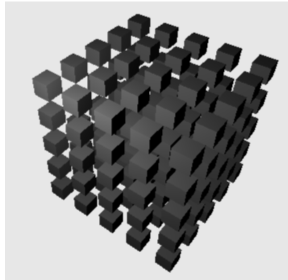

- Tests use GMRES (PETSc), tolerance = 10^{-8}

SKY 15x15x15

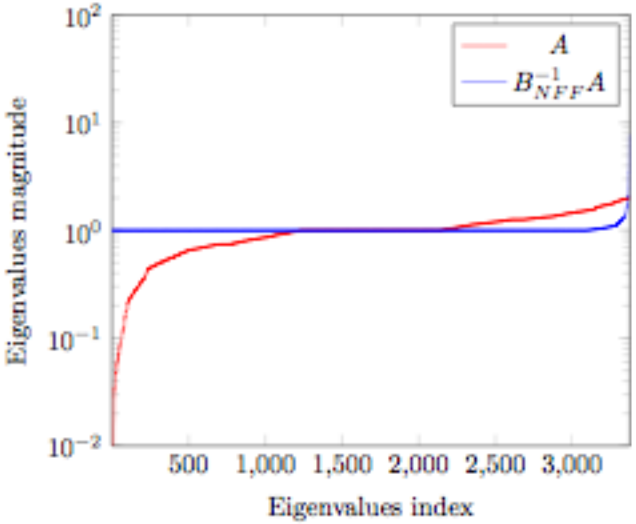

NFF, SKY 225x225x225

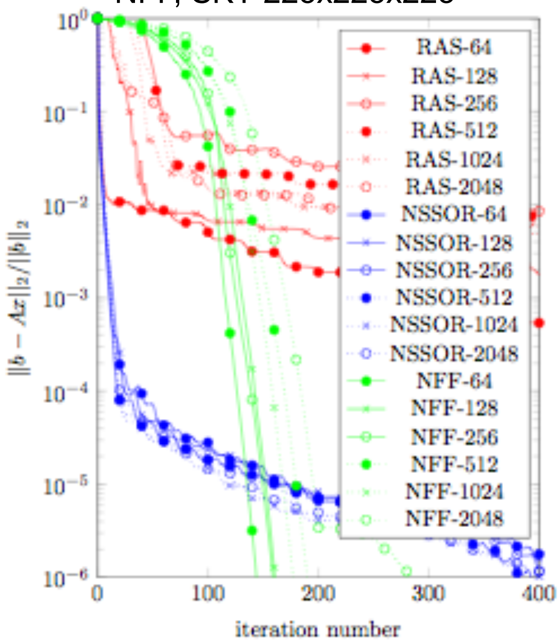

NFF, SKY 400x400x400

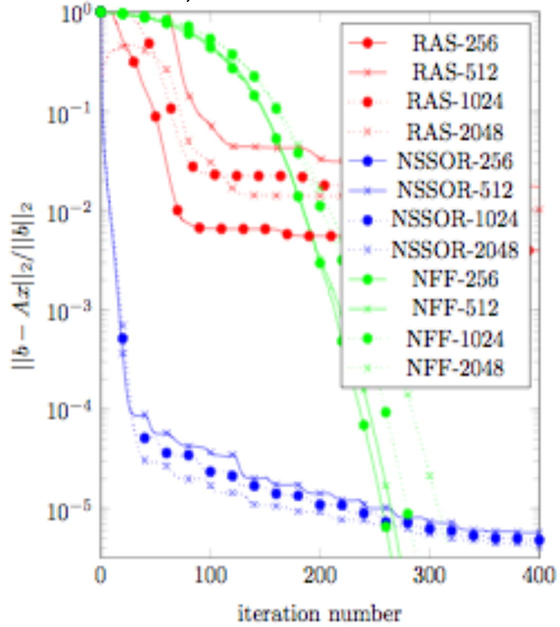

Comparison with Restricted Additive Schwarz (RAS)

Settings:

- Curie supercomputer based on Bullx system, nodes composed of two eight-core Intel Sandy Bridge.
- Subdomains solved using Pardiso, separators solved using MUMPS.
- GMRES and RAS from PETSc.

NFF vs RAS, SKY 400x400x400

Subdom	Iteration	Error	Iteration	Error
256	5489	5.9e-7	268	2.2e-6
512	6126	2.7e-6	273	3.2e-6
1024	7163	1.8e-6	289	2.6e-6
2048	10000	3.7e-6	317	3.8e-6

NFF performance, SKY 400x400x400

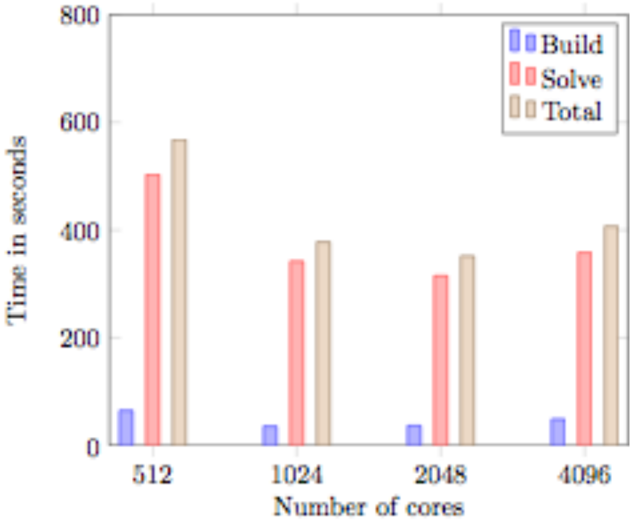

NFF vs RAS, SKY 225x225x225

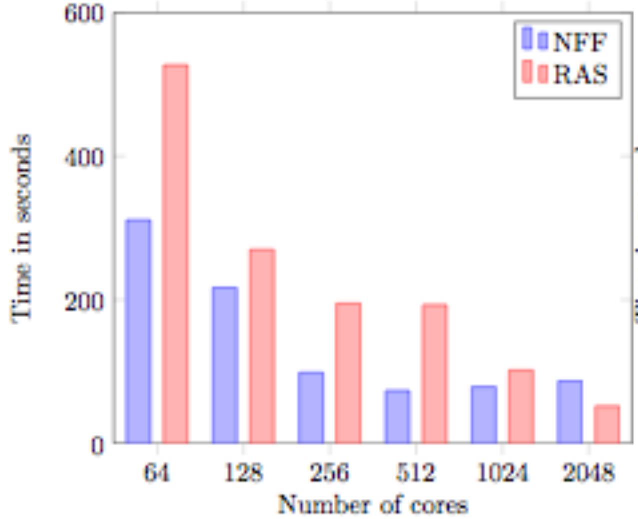

NFF vs RAS, SKY 400x400x400

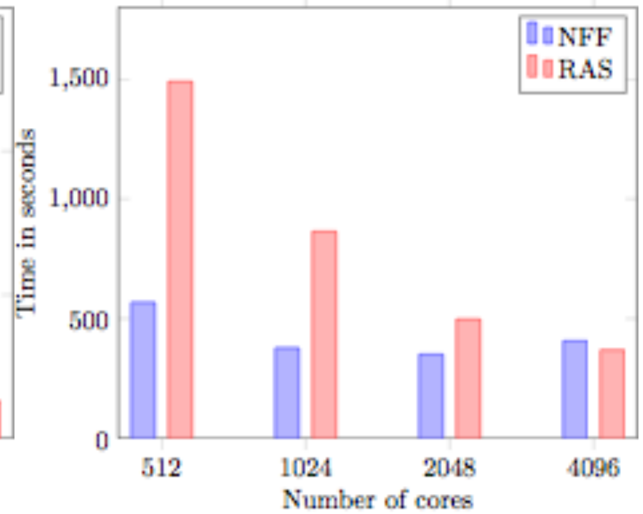

Best student paper finalist, Qu, LG, Nataf, SC'13