# Data-flow Parallel Implementation with OmpSs of ILUPACK

#### J.I. Aliaga<sup>1</sup> R.M. Badía<sup>2</sup> M. Barreda<sup>1</sup> M. Bollhöfer<sup>3</sup> E.S. Quintana-Ortí<sup>1</sup>

<sup>1</sup>Deparment of Computer Science and Engineering, Univ. Jaume I (Spain) {aliaga, mvaya, quintana}@icc.uji.es

> <sup>2</sup>Barcelona Supercomputing Center (BSC-CNS) (Spain). rosa.m.badia@bsc.es

<sup>3</sup>Institute of Computational Mathematics, TU-Braunschweig (Germany) m.bollhoefer@tu-bs.de

#### July, 2014

ILUPACK with OMPSS

パ

• We consider the efficient solution of linear systems

Ax = b

- $A \in \mathbb{R}^{n,n}$  is a **large** and **sparse** SPD coefficient matrix
- $x \in \mathbb{R}^n$  is the sought-after solution
- $b \in \mathbb{R}^n$  is a given r.h.s. vector
- The solution of this problem
  - arises in very many large-scale application problems
  - is frequently the most time-consuming part of the computation

j'i

Two main approaches to "attack" the problem:

- Direct methods (sparse Cholesky, sparse LU, etc.) may fail due to excessive fill-in of the factors
- Iterative methods (CG, GMRES, etc.) may present slow convergence or even divergence

Convergence of iterative methods can be accelerated via preconditioning:

- Incomplete LU factorization
- Sparse approximate inverse
- ...

パ

Two main approaches to "attack" the problem:

- Direct methods (sparse Cholesky, sparse LU, etc.) may fail due to excessive fill-in of the factors
- Iterative methods (CG, GMRES, etc.) may present slow convergence or even divergence

Convergence of iterative methods can be accelerated via preconditioning:

- Incomplete LU factorization
- Sparse approximate inverse
- ...



#### We propose to solve Ax = b iteratively using

- Preconditioned Conjugate Gradient (PCG)
  - Among the best iterative approaches to solve SPD linear systems
- ILUPACK preconditioning techniques
  - Incomplete Cholesky (IC) factorization-based preconditioners
     → Multilevel IC (MIC) preconditioners: Inverse-based approach
  - Successful for a wide range of large-scale application problems of scientific interest with up to million of equations



#### We propose to solve Ax = b iteratively using

- Preconditioned Conjugate Gradient (PCG)
  - Among the best iterative approaches to solve SPD linear systems
- ILUPACK preconditioning techniques
  - Incomplete Cholesky (IC) factorization-based preconditioners → Multilevel IC (MIC) preconditioners: Inverse-based approach
  - Successful for a wide range of large-scale application problems of scientific interest with up to million of equations



- The large scale of the problems faced by ILUPACK makes it evident the need to reduce the time-to-solution via parallel computing techniques
- Development of parallel ILUPACK-based solvers
  - Computation of adjacency graph G<sub>A</sub>
  - Parallelism revealed by nested dissection on G<sub>A</sub>
    - $\rightarrow$  recursive application of vertex separator partitioning
  - Dynamic scheduling to improve load-balancing during execution → run-time tuned for a specific platform/application problem
  - OmpSs is a data-flow programming model
    - ightarrow Exploit the data-dependencies between the tasks
    - $\rightarrow$  Useful to combine different platforms in heterogenous systems



- The large scale of the problems faced by ILUPACK makes it evident the need to reduce the time-to-solution via parallel computing techniques
- Development of parallel ILUPACK-based solvers
  - Computation of adjacency graph G<sub>A</sub>
  - Parallelism revealed by nested dissection on G<sub>A</sub>
    - $\rightarrow$  recursive application of vertex separator partitioning
  - Dynamic scheduling to improve load-balancing during execution
    - $\rightarrow$  run-time tuned for a specific platform/application problem
  - OmpSs is a data-flow programming model
    - ightarrow Exploit the data-dependencies between the tasks
    - $\rightarrow$  Useful to combine different platforms in heterogenous systems



- The large scale of the problems faced by ILUPACK makes it evident the need to reduce the time-to-solution via parallel computing techniques
- Development of parallel ILUPACK-based solvers
  - Computation of adjacency graph G<sub>A</sub>
  - Parallelism revealed by nested dissection on G<sub>A</sub>
    - $\rightarrow$  recursive application of vertex separator partitioning
  - Dynamic scheduling to improve load-balancing during execution
    - $\rightarrow$  run-time tuned for a specific platform/application problem
  - OmpSs is a data-flow programming model
    - $\rightarrow$  Exploit the data-dependencies between the tasks
    - $\rightarrow$  Useful to combine different platforms in heterogenous systems

## Outline



- Motivation and Introduction
- ILUPACK inverse-based MIC preconditioners
- 3 Parallel inverse-based MIC factorization preconditioners
- 4 Experimental Results
- Parallel Preconditioned Conjugate Gradient
- 6 Conclusions and Future Work

## Outline



- ILUPACK inverse-based MIC preconditioners
- 3 Parallel inverse-based MIC factorization preconditioners
- 4 Experimental Results
- 5 Parallel Preconditioned Conjugate Gradient
- 6 Conclusions and Future Work



• The (root-free) sparse Cholesky decomp. computes  $L \in \mathbb{R}^{n,n}$ , sparse unit lower triangular, and  $D \in \mathbb{R}^{n,n}$ , diagonal with positive entries s.t.

$$A = LDL^T$$

• In the Incomplete Cholesky (IC) factorization A is factored as

$$A = \tilde{L}\tilde{D}\tilde{L}^{T} + E$$

where  $E \in \mathbb{R}^{n,n}$  is a "small" perturbation matrix consisting of those entries dropped during factorization

M = LDL<sup>T</sup> is applied to the original system Ax = b in order to accelerate the convergence of the PCG solver

$$(\tilde{D}^{-1/2}\tilde{L}^{-1})A(\tilde{L}^{-T}\tilde{D}^{-1/2})$$

- Inverse-based IC factorizations construct  $M = \tilde{L}\tilde{D}\tilde{L}^T$  s.t.  $\|\tilde{L}^{-1}\| \leq \kappa$ 
  - $\kappa = 5$  or  $\kappa = 10$  are good choices for many problems
  - Pivoting and **multilevel** methods are required in general



• The (root-free) sparse Cholesky decomp. computes  $L \in \mathbb{R}^{n,n}$ , sparse unit lower triangular, and  $D \in \mathbb{R}^{n,n}$ , diagonal with positive entries s.t.

$$A = LDL^T$$

• In the Incomplete Cholesky (IC) factorization A is factored as

$$A = \tilde{L}\tilde{D}\tilde{L}^{T} + E$$

where  $E \in \mathbb{R}^{n,n}$  is a "small" perturbation matrix consisting of those entries dropped during factorization

$$(\tilde{D}^{-1/2}\tilde{L}^{-1})A(\tilde{L}^{-T}\tilde{D}^{-1/2})$$

- Inverse-based IC factorizations construct  $M = \tilde{L}\tilde{D}\tilde{L}^T$  s.t.  $\|\tilde{L}^{-1}\| \leq \kappa$ 
  - $\kappa = 5$  or  $\kappa = 10$  are good choices for many problems
  - Pivoting and **multilevel** methods are required in general



- The (root-free) sparse Cholesky decomp. computes *L* ∈ ℝ<sup>n,n</sup>, sparse unit lower triangular, and *D* ∈ ℝ<sup>n,n</sup>, diagonal with positive entries s.t.
- In the Incomplete Cholesky (IC) factorization A is factored as

$$A = \tilde{L}\tilde{D}\tilde{L}^{T} + E$$

where  $E \in \mathbb{R}^{n,n}$  is a "small" perturbation matrix consisting of those entries dropped during factorization

M = L D L<sup>T</sup> is applied to the original system Ax = b in order to accelerate the convergence of the PCG solver

$$(\tilde{D}^{-1/2}\tilde{L}^{-1})A(\tilde{L}^{-T}\tilde{D}^{-1/2})$$

- Inverse-based IC factorizations construct  $M = \tilde{L}\tilde{D}\tilde{L}^T$  s.t.  $\|\tilde{L}^{-1}\| \leq \kappa$ 
  - $\kappa = 5$  or  $\kappa = 10$  are good choices for many problems
  - Pivoting and **multilevel** methods are required in general



- The (root-free) sparse Cholesky decomp. computes *L* ∈ ℝ<sup>n,n</sup>, sparse unit lower triangular, and *D* ∈ ℝ<sup>n,n</sup>, diagonal with positive entries s.t.
- In the Incomplete Cholesky (IC) factorization A is factored as

$$A = \tilde{L}\tilde{D}\tilde{L}^{T} + E$$

where  $E \in \mathbb{R}^{n,n}$  is a "small" perturbation matrix consisting of those entries dropped during factorization

M = L D L<sup>T</sup> is applied to the original system Ax = b in order to accelerate the convergence of the PCG solver

$$(\tilde{D}^{-1/2}\tilde{L}^{-1})A(\tilde{L}^{-T}\tilde{D}^{-1/2})$$

- Inverse-based IC factorizations construct  $M = \tilde{L}\tilde{D}\tilde{L}^T$  s.t.  $\|\tilde{L}^{-1}\| \leq \kappa$ 
  - $\kappa = 5$  or  $\kappa = 10$  are good choices for many problems
  - Pivoting and **multilevel** methods are required in general

By using the Crout variant of the IC factorization:



- The IC decomp. employs Inverse-based pivoting
  - $\rightarrow$  Row/Col. k of  $\hat{L} / \hat{L}^T$  obtained at step k along with estimation  $t_k \approx ||e_k^T \hat{L}^{-1}|$  $\rightarrow$  Row/Col. s.t.  $t_k > \kappa$  are moved to the bottom/right-end of the matrix
- The IC factorization stops when the trailing block only contains rejected pivots  $\rightarrow$  The computation of the approximate Schur complement  $\tilde{S}_c$  completes . . .

By using the Crout variant of the IC factorization:



The IC decomp. employs Inverse-based pivoting

 $\rightarrow$  Row/Col. *k* of  $\tilde{L} / \tilde{L}^{\tilde{T}}$  obtained at step *k* along with estimation  $t_k \approx ||e_k^T \tilde{L}^{-1}||$  $\rightarrow$  Row/Col. s.t.  $t_k > \kappa$  are moved to the bottom/right-end of the matrix

• The IC factorization stops when the trailing block only contains rejected pivots  $\rightarrow$  The computation of the approximate Schur complement  $\tilde{S}_c$  completes . . .

By using the Crout variant of the IC factorization:



- The IC decomp. employs Inverse-based pivoting
  - $\rightarrow$  Row/Col. *k* of  $\tilde{L} / \tilde{L}^{T}$  obtained at step *k* along with estimation  $t_{k} \approx ||e_{k}^{T}\tilde{L}^{-1}||$
  - $\rightarrow$  Row/Col. s.t.  $\mathit{t_k} > \kappa$  are moved to the bottom/right-end of the matrix
- The IC factorization **stops** when the trailing block only contains rejected pivots  $\rightarrow$  The computation of the approximate Schur complement  $\tilde{S}_c$  completes ...

By using the Crout variant of the IC factorization:



• ... a **partial** IC decomp. of a permuted system with  $\tilde{L}_B^{-1}$  satisfying  $\|\tilde{L}_B^{-1}\| \leq \kappa$ 

$$\hat{P}^{T}\hat{A}\hat{P} \equiv \left(\begin{array}{cc} B & F^{T} \\ F & C \end{array}\right) = \left(\begin{array}{cc} \tilde{L}_{B} & 0 \\ \tilde{L}_{F} & I \end{array}\right) \left(\begin{array}{cc} \tilde{D}_{B} & 0 \\ 0 & \tilde{S}_{C} \end{array}\right) \left(\begin{array}{cc} \tilde{L}_{B}^{T} & \tilde{L}_{F}^{T} \\ 0 & I \end{array}\right) + E$$

• The whole method is restarted on  $\tilde{S}_c$  leading to a **multilevel** approach

MIC factorization of five-point matrix arising from Laplace PDE discretization



 $\kappa=5$  and  $\tau=10^{-3}$  leads to a MIC factorization with 5 levels

## Outline



- Motivation and Introduction
- ILUPACK inverse-based MIC preconditioners
- 3 Parallel inverse-based MIC factorization preconditioners
- 4 Experimental Results
- 6 Parallel Preconditioned Conjugate Gradient
- 6 Conclusions and Future Work

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• Analyzing the sparsity pattern of A we can permute it,

$$\Pi^{T} A \Pi = \begin{pmatrix} A_{11} & 0 & A_{31}^{T} \\ 0 & A_{22} & A_{32}^{T} \\ \hline A_{31} & A_{32} & A_{33} \end{pmatrix}$$

Â <sub>11</sub> 0 Â <sub>31</sub>				* *	$\begin{bmatrix} A_{11} & 0 & 0 & 0 \\ 0 & A_{22} & 0 & 0 \\ 0 & 0 & A_{33} & 0 \end{bmatrix}$	$\begin{bmatrix} A_{51}^T & 0 \\ A_{52}^T & 0 \\ 0 & A_{63}^T \end{bmatrix}$	$\begin{vmatrix} A_{71}^T \\ A_{72}^T \\ A_{73}^T \end{vmatrix}$
				*	0 0 0 A <sub>44</sub>		$A_{74}^{I}$
				2 *			
					0 0 A <sub>63</sub> A <sub>64</sub>		
*	244	*	* * *				

We can disasemble the permuted matrix into 4 submatrices

 $\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}^{T}} & \underline{A_{71}^{T}} \\ \underline{A_{51}} & \underline{A_{75}^{T}} & \underline{A_{75}^{T}} \\ A_{71} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{22}} & \underline{A_{52}^{T}} & \underline{A_{72}^{T}} \\ \underline{A_{52}} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \\ A_{72} & \underline{A_{75}^{T}} & \underline{A_{77}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{33}} & \underline{A_{63}^{T}} & \underline{A_{73}^{T}} \\ \underline{A_{63}} & \underline{A_{66}^{3}} & \underline{A_{76}^{3}} \\ A_{73} & \underline{A_{76}^{3}} & \underline{A_{77}^{3}} \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{64}^{T}} & \underline{A_{74}^{T}} \\ \underline{A_{64}} & \underline{A_{66}^{A}} & \underline{A_{76}^{A}} \\ A_{74} & \underline{A_{76}^{A}} & \underline{A_{77}^{A}} \end{bmatrix}$ 

where

ILUPACK with OMPSS



The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• Analyzing the sparsity pattern of A we can permute it,

$\begin{bmatrix} \hat{A}_{11} & 0 & \hat{A}_{31} \end{bmatrix}$	0 0 0	*	1 [		
0 $\hat{A}_{22} \hat{A}_{32}^{T}$	0 0 0	*			
Â <sub>31</sub> Â <sub>32</sub> Â <sub>33</sub>	0 0 0	*		0 0 A <sub>33</sub> 0	$A_{\rm Z3}^T$
0 0 0	$\bar{A}_{11} \ 0 \ \bar{A}_{31}^T$	*	$\rightarrow$	0 0 0 A <sub>44</sub>	$A_{74}^{T}$
0 0 0	0 $\bar{A}_{22} \bar{A}_{32}^T$	*			
0 0 0	$\bar{A}_{31} \bar{A}_{32} \bar{A}_{33}$	*		0 0 A <sub>63</sub> A <sub>64</sub>	
* * *	* * *	A <sub>33</sub>			

We can disasemble the permuted matrix into 4 submatrices

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}^{T}} & \underline{A_{71}^{T}} \\ \underline{A_{51}} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \\ A_{71} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{22}} & \underline{A_{52}^{T}} & A_{72}^{T} \\ \underline{A_{52}} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \\ A_{72} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{54}^{T}} & \underline{A_{74}^{T}} \\ \underline{A_{56}} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \\ A_{72} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{56}^{T}} & A_{74}^{T} \\ \underline{A_{66}} & \underline{A_{76}^{T}} \\ A_{73} & \underline{A_{76}^{T}} & A_{77}^{T} \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{66}^{T}} & A_{76}^{T} \\ \underline{A_{76}} & \underline{A_{76}^{T}} \\ A_{74} & \underline{A_{76}^{T}} & A_{77}^{T} \end{bmatrix}$$

$$\begin{split} A_{55} &= A_{55}^1 + A_{55}^2 \ , \ A_{75} &= A_{75}^1 + A_{75}^2 \ | \ A_{66} &= A_{66}^3 + A_{66}^4 \ , \ A_{76} &= A_{76}^3 + A_{76}^4 \\ A_{77} &= A_{77}^1 + A_{77}^2 \ + A_{77}^3 + A_{77}^4 \end{split}$$

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

Analyzing the sparsity pattern of A we can permute it,

$\begin{bmatrix} \hat{A}_{11} & 0 & \hat{A}_{31}^T \\ 0 & \hat{A}_{22} & \hat{A}_{32}^T \\ \hat{A}_{31} & \hat{A}_{32} & \hat{A}_{33} \end{bmatrix}$	0 0 0 0 0 0 0 0 0	*	$\begin{bmatrix} A_{11} & 0 & 0 & 0 \\ 0 & A_{22} & 0 & 0 \\ 0 & 0 & A_{33} & 0 \end{bmatrix}$	$egin{array}{ccc} A_{51}^T & 0 \ A_{52}^T & 0 \ 0 & A_{63}^T \end{array}$	$\begin{bmatrix} A_{71}^T \\ A_{72}^T \\ A_{73}^T \end{bmatrix}$
0 0 0	$\bar{A}_{11}$ 0 $\bar{A}_{31}^T$	* →	0 0 0 A <sub>44</sub>	$0 A_{64}^{T}$	$A_{74}^T$
0 0 0	$0 \bar{A}_{22} \bar{A}_{32}^{T}$	*	A <sub>51</sub> A <sub>52</sub> 0 0	A <sub>55</sub> 0	$A_{75}^T$
0 0 0	Ā <sub>31</sub> Ā <sub>32</sub> Ā <sub>33</sub>	*	0 0 A <sub>63</sub> A <sub>64</sub>	0 A <sub>66</sub>	$A_{76}^T$
* * *	* * *	A <sub>33</sub>	A <sub>71</sub> A <sub>72</sub> A <sub>73</sub> A <sub>74</sub>	A <sub>75</sub> A <sub>76</sub>	A <sub>77</sub>

We can disasemble the permuted matrix into 4 submatrices

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}^{T}} & \underline{A_{71}^{T}} \\ \underline{A_{51}} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \\ A_{71} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{22}} & \underline{A_{52}^{T}} & A_{72}^{T} \\ \underline{A_{52}} & \underline{A_{52}^{T}} & \underline{A_{55}^{T}} \\ A_{72} & \underline{A_{55}^{T}} & \underline{A_{75}^{T}} \end{bmatrix}, \begin{bmatrix} \underline{A_{33}} & \underline{A_{53}^{T}} & A_{73}^{T} \\ \underline{A_{63}} & \underline{A_{66}^{T}} & \underline{A_{76}^{T}} \\ A_{73} & \underline{A_{76}^{T}} & A_{77}^{T} \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{64}^{T}} & A_{74}^{T} \\ \underline{A_{64}} & \underline{A_{66}^{T}} & A_{76}^{T} \\ A_{74} & \underline{A_{76}^{T}} & A_{76}^{T} \\ A_{76} & A_{76}^{T} \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{64}^{T}} & A_{74}^{T} \\ \underline{A_{64}} & \underline{A_{66}^{T}} & A_{76}^{T} \\ A_{76} & A_{76}^{T} \\ A_{76} & A_{76}^{T} \\ A_{76} & A_{76}^{T} \\ \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{66}} & A_{76}^{T} \\ \underline{A_{76}} & A_{76}^{T} \\ A_{76} & A_{76}^{T} \\ A_{77} & A_{76}^{T} \\ A_{77} & A_{77}^{T} \\ A_{77} & A_{77}^{T}$$

$$\begin{split} A_{55} &= A_{55}^1 + A_{55}^2 \ , \ A_{75} &= A_{75}^1 + A_{75}^2 \ | \ A_{66} &= A_{66}^3 + A_{66}^4 \ , \ A_{76} &= A_{76}^3 + A_{76}^4 \\ A_{77} &= A_{77}^1 + A_{77}^2 \ + A_{77}^3 + A_{77}^4 \end{split}$$

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

Analyzing the sparsity pattern of A we can permute it,

$\begin{bmatrix} \hat{A}_{11} & 0 & \hat{A}_{31}^T \\ 0 & \hat{A}_{22} & \hat{A}_{32}^T \\ \hat{A}_{31} & \hat{A}_{32} & \hat{A}_{33} \end{bmatrix}$	0 0 0 0 0 0 0 0 0	* * *	$\begin{bmatrix} A_{11} & 0 & 0 & 0 \\ 0 & A_{22} & 0 & 0 \\ 0 & 0 & A_{33} & 0 \end{bmatrix}$	$\begin{array}{ccc} A_{51}^T & 0 \\ A_{52}^T & 0 \\ 0 & A_{63}^T \end{array}$	$\begin{bmatrix} A_{71}^T \\ A_{72}^T \\ A_{73}^T \end{bmatrix}$
0 0 0	$\bar{A}_{11} \ 0 \ \bar{A}_{31}^T$	$* \rightarrow$	0 0 0 A <sub>44</sub>	0 $A_{64}^{T}$	$A_{74}^T$
0 0 0	$0 \bar{A}_{22} \bar{A}_{32}^{T}$	*	A <sub>51</sub> A <sub>52</sub> 0 0	A <sub>55</sub> 0	$A_{75}^T$
0 0 0	Ā <sub>31</sub> Ā <sub>32</sub> Ā <sub>33</sub>	*	0 0 A <sub>63</sub> A <sub>64</sub>	0 A <sub>66</sub>	$A_{76}^T$
* * *	* * *	A <sub>33</sub>	A71 A72 A73 A74	A <sub>75</sub> A <sub>76</sub>	A <sub>77</sub>

• We can disasemble the permuted matrix into 4 submatrices

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}}^T & \underline{A_{71}}^T \\ \underline{A_{51}} & \underline{A_{55}}^T & \underline{A_{75}}^T \\ A_{71} & \underline{A_{55}}^T & \underline{A_{77}}^T \end{bmatrix}, \begin{bmatrix} \underline{A_{22}} & \underline{A_{52}}^T & \underline{A_{72}}^T \\ \underline{A_{52}} & \underline{A_{52}}^T & \underline{A_{52}}^T \\ A_{72} & \underline{A_{52}}^T & \underline{A_{52}}^T & \underline{A_{75}}^T \end{bmatrix}, \begin{bmatrix} \underline{A_{33}} & \underline{A_{53}}^T & \underline{A_{73}}^T \\ \underline{A_{63}} & \underline{A_{66}}^T & \underline{A_{76}}^T \\ A_{73} & \underline{A_{76}}^T & \underline{A_{77}}^T \end{bmatrix}, \begin{bmatrix} \underline{A_{44}} & \underline{A_{64}}^T & \underline{A_{74}}^T \\ \underline{A_{64}} & \underline{A_{76}}^T & \underline{A_{76}}^T \\ \underline{A_{76}} & \underline{A_{76}}^T & \underline{A_{77}}^T \end{bmatrix}$$

$$\begin{aligned} A_{55} &= A_{55}^1 + A_{55}^2 \ , \ A_{75} &= A_{75}^1 + A_{75}^2 \ | \ A_{66} &= A_{66}^3 + A_{66}^4 \ , \ A_{76} &= A_{76}^3 + A_{76}^4 \\ A_{77} &= A_{77}^1 + A_{77}^2 \ + A_{77}^2 + A_{77}^4 + A_{77}^4 \end{aligned}$$

Ĵ

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

In parallel, we can compute the partial IC factorization of each submatrix

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}}^T & \underline{A_{71}}^T \\ \underline{A_{51}} & \underline{A_{15}}^T & \underline{A_{75}}^T \\ A_{71} & \underline{A_{75}}^T & \underline{A_{77}}^T \end{bmatrix} = \begin{bmatrix} \underline{L_{11}} & 0 & 0 \\ \hline \underline{L_{51}} & \overline{1} & 0 \\ L_{71} & \overline{0} & \overline{1} \end{bmatrix} \begin{bmatrix} \underline{D_{11}} & 0 & 0 \\ \hline 0 & \underline{S_{15}}^T & \underline{S_{75}}^T \\ 0 & \overline{S_{15}}^T & S_{77}^T \end{bmatrix} \begin{bmatrix} \underline{L_{11}}^T & \underline{L_{71}}^T & \underline{L_{71}}^T \\ \hline 0 & \overline{1} & 0 \\ 0 & \overline{0} & \overline{1} \end{bmatrix}$$

A partial IC factorization of matrix A is computed:



and restarting, we obtain a "structural" multilevel factorization.



The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

In parallel, we can compute the partial IC factorization of each submatrix

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}^T} & \underline{A_{71}^T} \\ \underline{A_{51}} & \underline{A_{15}^T} & \underline{A_{75}^T} \\ A_{71} & \underline{A_{75}^T} & \underline{A_{77}^T} \end{bmatrix} = \begin{bmatrix} \underline{L_{11}} & 0 & 0 \\ \underline{L_{51}} & \underline{I} & \underline{0} \\ L_{71} & 0 & 1 \end{bmatrix} \begin{bmatrix} \underline{D_{11}} & 0 & 0 \\ 0 & \underline{S_{15}^T} & \underline{S_{15}^T} \\ 0 & \underline{S_{15}^T} & S_{77}^T \end{bmatrix} \begin{bmatrix} \underline{L_{11}^T} & \underline{L_{51}^T} & \underline{L_{71}^T} \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}.$$

The Schur complements must be properly acumulated to continue

$$\begin{bmatrix} \underline{S}_{55}^{-} & \underline{S}_{75}^{-} \\ \overline{S}_{75}^{-} & \overline{S}_{77}^{-} \end{bmatrix} + \begin{bmatrix} \underline{S}_{55}^{2} & \underline{S}_{75}^{-} \\ \overline{S}_{75}^{2} & \overline{S}_{77}^{-} \end{bmatrix} = \begin{bmatrix} \underline{S}_{55} & \underline{S}_{75}^{-} \\ \overline{S}_{75}^{2} & \overline{S}_{77}^{-} \end{bmatrix} \\ \begin{bmatrix} \underline{S}_{66}^{3} & \underline{S}_{76}^{-} \\ \overline{S}_{76}^{3} & \overline{S}_{77}^{-} \end{bmatrix} + \begin{bmatrix} \underline{S}_{66}^{4} & \underline{S}_{76}^{4} \\ \overline{S}_{76}^{4} & \overline{S}_{77}^{-} \end{bmatrix} = \begin{bmatrix} \underline{S}_{66} & \underline{S}_{76}^{-} \\ \overline{S}_{76}^{4} & \underline{S}_{77}^{4} \end{bmatrix}$$

Assembling two submatrices we obtain

A partial IC factorization of matrix A is computed:



The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

In parallel, we can compute the partial IC factorization of each submatrix

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}^T} & \underline{A_{71}^T} \\ \underline{A_{51}} & \underline{A_{15}^T} & \underline{A_{75}^T} \\ A_{71} & \underline{A_{75}^T} & \underline{A_{77}^T} \end{bmatrix} = \begin{bmatrix} \underline{L_{11}} & 0 & 0 \\ \underline{L_{51}} & \underline{I} & \underline{0} \\ L_{71} & 0 & 1 \end{bmatrix} \begin{bmatrix} \underline{D_{11}} & 0 & 0 \\ 0 & \underline{S_{15}^T} & \underline{S_{15}^T} \\ 0 & \underline{S_{15}^T} & S_{77}^T \end{bmatrix} \begin{bmatrix} \underline{L_{11}^T} & \underline{L_{51}^T} & \underline{L_{71}^T} \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix}.$$

The Schur complements must be properly acumulated to continue

$$\begin{bmatrix} \underline{S}_{55}^{1} & \underline{S}_{75}^{1} & T \\ \overline{S}_{75}^{1} & \overline{S}_{77}^{1} \end{bmatrix} + \begin{bmatrix} \underline{S}_{55}^{2} & \underline{S}_{75}^{2} & T \\ \overline{S}_{75}^{2} & \overline{S}_{77}^{2} \end{bmatrix} = \begin{bmatrix} \underline{S}_{55} & \underline{S}_{75}^{T} \\ \overline{S}_{75} & \overline{S}_{77}^{2} \end{bmatrix} \\ \begin{bmatrix} \underline{S}_{66}^{3} & \underline{S}_{76}^{3} \\ \overline{S}_{76}^{3} & \overline{S}_{77}^{3} \end{bmatrix} + \begin{bmatrix} \underline{S}_{66}^{4} & \underline{S}_{76}^{4} \\ \overline{S}_{76}^{4} & \underline{S}_{77}^{4} \end{bmatrix} = \begin{bmatrix} \underline{S}_{66} & \underline{S}_{76}^{T} \\ \overline{S}_{76} & \overline{S}_{77}^{3} \end{bmatrix}$$

Assembling two submatrices we obtain

$$\begin{bmatrix} S_{55} & 0 & S_{75}^T \\ 0 & S_{66} & S_{76}^T \\ \hline S_{75} & S_{76} & S_{77} \end{bmatrix} , \quad S_{77} = S_{77}^{12} + S_{77}^{34}$$

A partial IC factorization of matrix A is computed:

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

In parallel, we can compute the partial IC factorization of each submatrix

$$\begin{bmatrix} \underline{A_{11}} & \underline{A_{51}}^T & \underline{A_{71}}^T \\ \underline{A_{51}} & \underline{A_{75}}^T & \underline{A_{77}}^T \\ A_{71} & \underline{A_{75}}^T & \underline{A_{77}}^T \end{bmatrix} = \begin{bmatrix} \underline{L_{11}} & 0 & 0 \\ \hline \underline{L_{51}} & \overline{1} & 0 \\ L_{71} & \overline{0} & \overline{1} \end{bmatrix} \begin{bmatrix} \underline{D_{11}} & 0 & 0 \\ \hline 0 & \underline{S_{15}}^T & \underline{S_{75}}^T \\ 0 & \overline{S_{15}}^T & S_{77}^T \end{bmatrix} \begin{bmatrix} \underline{L_{11}}^T & \underline{L_{51}}^T & \underline{L_{71}}^T \\ \hline 0 & \overline{1} & 0 \\ 0 & \overline{0} & \overline{1} \end{bmatrix}$$

• A partial IC factorization of matrix A is computed:

$ \begin{bmatrix} L_{11} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & L_{22} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & L_{33} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & L_{44} & 0 & 0 & 0 \end{bmatrix} $		$ \begin{bmatrix} D_{11} & 0 \\ 0 & D_2 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} $	$\begin{array}{ccc} 0 & 0 \\ 2 & 0 & 0 \\ D_{33} & 0 \\ 0 & D_{44} \end{array}$	0 0 0 0 0 0 0 0	0 0 0 0	$\begin{bmatrix} L_{11}^T & 0 & 0 & 0 \\ 0 & L_{22}^T & 0 & 0 \\ 0 & 0 & L_{33}^T & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$	$L_{51}^{T} 0$ $L_{52}^{T} 0$ $0 L_{63}^{T}$ $0 L^{T}$	$\begin{bmatrix} L_{71}^T \\ L_{72}^T \\ L_{73}^T \end{bmatrix}$
$L_{51} L_{52} 0 0 I 0$	0	0 0	0 0	S <sub>55</sub> 0	$S_{75}^T$	0 0 0 0 244	I 0	0
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	I		0 0	0 S <sub>66</sub> S <sub>75</sub> S <sub>76</sub>	$\frac{S'_{76}}{S_{77}}$	0 0 0 0	0 I	0

and restarting, we obtain a "structural" multilevel factorization.

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



- Each node refers to a block to be factorized, while the edges define the data dependency.
- The nodes without any dependency can be factorized in parallel.
- The non-leaf nodes have to wait until its children have finalized.

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



 Each node refers to a block to be factorized, while the edges define the data dependency.

The nodes without any dependency can be factorized in parallel.

• The non-leaf nodes have to wait until its children have finalized.

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



- Each node refers to a block to be factorized, while the edges define the data dependency.
- The nodes without any dependency can be factorized in parallel.
- The non-leaf nodes have to wait until its children have finalized.

The key to parallelize the preconditioner computation is to decompose the IC factorization into independent tasks:

• The task dependency tree represents the behaviour of the computation



• The nodes without any dependency can be factorized in parallel.

<sup>1</sup>55

• The non-leaf nodes have to wait until its children have finalized.





- Task-level parallelism revealed by nested dissection on adjacency graph  $G_A \rightarrow$  fast and efficient **multilevel** (MLND) variants in METIS and SCOTCH
- Parallelization driven by the dependencies "captured" in the task tree
  - $\rightarrow$  subdomains first eliminated independently, then separators (2, 1) and (2, 2) in parallel, finally the root separator (3, 1)

パ

The parallelization of ILUPACK must consider two multilevel procedures

- ILUPACK computes Multilevel IC factorization
  - $\rightarrow$  Numerical levels related to inverse-based IC factorizations
- Parallelization of IC factorization
  - $\rightarrow$  Structural levels computed by nested disection algorithm

The solution is to combine both multilevel procedures

- Structural levels are initially used to reveal the concurrence
- Numerical levels appears during the MIC factorization of each submatrix

The Multilevel IC factorization must be adjusted when it is applied on a submatrix

- Where to move the rejected pivots
- When to stop the factorization of a submatrix

The parallelization of ILUPACK must consider two multilevel procedures

- ILUPACK computes Multilevel IC factorization
  - $\rightarrow$  Numerical levels related to inverse-based IC factorizations
- Parallelization of IC factorization
  - $\rightarrow$  Structural levels computed by nested disection algorithm

The solution is to combine both multilevel procedures

- Structural levels are initially used to reveal the concurrence
- Numerical levels appears during the MIC factorization of each submatrix

The Multilevel IC factorization must be adjusted when it is applied on a submatrix

- Where to move the rejected pivots
- When to stop the factorization of a submatrix
# Parallel Inverse-based Multilevel IC factorization

Ĵ

The parallelization of ILUPACK must consider two multilevel procedures

- ILUPACK computes Multilevel IC factorization
  - $\rightarrow$  Numerical levels related to inverse-based IC factorizations
- Parallelization of IC factorization
  - $\rightarrow$  Structural levels computed by nested disection algorithm

The solution is to combine both multilevel procedures

- Structural levels are initially used to reveal the concurrence
- Numerical levels appears during the MIC factorization of each submatrix

The Multilevel IC factorization must be adjusted when it is applied on a submatrix

- Where to move the rejected pivots
- When to stop the factorization of a submatrix

# Parallel Inverse-based Multilevel IC factorization

The tasks compute a local MIC restricted to the leading block of the submatrix



- IC decomp. employs restricted Inverse-based pivoting
  - $\rightarrow$  Rejected pivots are only moved to the bottom/right-end of the leading block
- IC decomp. stops if the trailing of the **leading block** only contains rejected pivots → Rejected pivots still "resolved" inside the task by entering next MIC level

# Parallel Inverse-based Multilevel IC factorization

The tasks compute a local MIC restricted to the leading block of the submatrix



● When the set of rejected pivots is "small enough" the local MIC stops → The task "sends" the local contributions to the parent task

#### Parallel inverse-based MIC factorization





- The parent task merges contributions resulting from its children
- A new submatrix is constructed for the local MIC decomp. of the separator task
  - $\rightarrow$  Bad pivots resulting from its children copied on top of the leading block
  - $\rightarrow$  Contributions blocks added



inverse-based Multilevel IC factorization



Parallel inverse-based Multilevel IC factorization

ILUPACK with OMPSS

# Parallel vs Sequential Inverse-based MIC factoriz. 🎽

Factorization of five-point matrix arising from Laplace PDE discretization



ILUPACK with OMPSS

J.Aliaga, R.Badía, M.Barreda, M.Bollhöfer, E.Quintana-Ortí

# Parallel vs Sequential Inverse-based MIC factoriz.

Factorization of five-point matrix arising from Laplace PDE discretization



ILUPACK with OMPSS

J.Aliaga, R.Badía, M.Barreda, M.Bollhöfer, E.Quintana-Ortí



On the way to Exascale, one of the main challenges is the parallel programming revolution:

- In the classical programming models, the code defines
  - Where to place data
  - What to run and where
  - How to communicate
- In the new programming models, the questions have changed,
  - What do I need to compute?
  - What data do I need to use?
  - Hints (not necessarily very precise) on potential concurrency, locality, ...
- For the programmers, the main feature of the new programming models is to write a single source useful for any target (multicores, multicomputers, heterogeneous systems with GPU?s, ...)



On the way to Exascale, one of the main challenges is the parallel programming revolution:

- In the classical programming models, the code defines
  - Where to place data
  - What to run and where
  - How to communicate
- In the new programming models, the questions have changed,
  - What do I need to compute?
  - What data do I need to use?
  - Hints (not necessarily very precise) on potential concurrency, locality, ...
- For the programmers, the main feature of the new programming models is to write a single source useful for any target (multicores, multicomputers, heterogeneous systems with GPU?s, ...)



On the way to Exascale, one of the main challenges is the parallel programming revolution:

- In the classical programming models, the code defines
  - Where to place data
  - What to run and where
  - How to communicate
- In the new programming models, the questions have changed,
  - What do I need to compute?
  - What data do I need to use?
  - Hints (not necessarily very precise) on potential concurrency, locality, ...
- For the programmers, the main feature of the new programming models is to write a single source useful for any target (multicores, multicomputers, heterogeneous systems with GPU?s, ...)



OmpSs is a new programming model, which extends OpenMP:

- Define a single address space
- The code can spawn tasks that will be executed in the future
- Tasks annotated with directionality clauses,

```
pragma omp task in(lvalue_expr-list) \
out(lvalue_expr-list) \
inout(lvalue_expr-list)
```

used to build dependences between tasks.

- The tasks wait to data to be produced by other tasks (in clause).
- Basis for memory management functionalities (replication, locality, movement, ... )



OmpSs is a new programming model, which extends OpenMP:

- Define a single address space
- The code can spawn tasks that will be executed in the future
- Tasks annotated with directionality clauses,

used to build dependences between tasks.

- The tasks wait to data to be produced by other tasks (in clause).
- Basis for memory management functionalities (replication, locality, movement, ... )



OmpSs is a new programming model, which extends OpenMP:

- Define a single address space
- The code can spawn tasks that will be executed in the future
- Tasks annotated with directionality clauses,

used to build dependences between tasks.

- The tasks wait to data to be produced by other tasks (in clause).
- Basis for memory management functionalities (replication, locality, movement, ... )

# OmpSs memory model



From the point of view of the programmer a single naming space exists

• The standard sequential/shared memory address space

From the point of view of the runtime and target platform, different possible scenarios

- Pure SMP:
  - Single address space
- Distributed/heterogeneous systems (cluster, gpus, ...):
  - Multiple address spaces exists, such that versions of same data may exist in several ones.
  - Data consistency is ensured by the implementation

# OmpSs memory model



From the point of view of the programmer a single naming space exists

• The standard sequential/shared memory address space

From the point of view of the runtime and target platform, different possible scenarios

- Pure SMP:
  - Single address space
- Distributed/heterogeneous systems (cluster, gpus, ...):
  - Multiple address spaces exists, such that versions of same data may exist in several ones.
  - Data consistency is ensured by the implementation

# OmpSs memory model



From the point of view of the programmer a single naming space exists

• The standard sequential/shared memory address space

From the point of view of the runtime and target platform, different possible scenarios

- Pure SMP:
  - Single address space
- Distributed/heterogeneous systems (cluster, gpus, ...):
  - Multiple address spaces exists, such that versions of same data may exist in several ones.
  - Data consistency is ensured by the implementation



Task dependency tree:



We define a structure called dag[][] to define the dependencies in OmpSs

					A <sub>44</sub>	A <sub>55</sub>	A <sub>66</sub>	
			T <sub>1</sub>			T <sub>4</sub>		
Task id. <i>j</i>			1	2		4		6
Left descendant	dag[0][ <b>j</b> ]	-		—	_		2	4
Right descendant						1		5
Ancestor		4	4	5	5	6	6	

The dag is created and initialized before the parallel computation begins

ILUPACK with OMPSS



Task dependency tree:



We define a structure called dag[][] to define the dependencies in OmpSs

Block		A <sub>11</sub>	A <sub>22</sub>	A <sub>33</sub>	A <sub>44</sub>	A <sub>55</sub>	A <sub>66</sub>	A <sub>77</sub>
Task		T <sub>0</sub>	T <sub>1</sub>	T <sub>2</sub>	T <sub>3</sub>	$T_4$	$T_5$	$T_6$
Task id. <i>j</i>		0	1	2	3	4	5	6
Left descendant	dag[0][ <b>j</b> ]	-	-	-	-	0	2	4
Right descendant	dag[1][ <b>j</b> ]	-	-	-	-	1	3	5
Ancestor	dag[2][ <b>j</b> ]	4	4	5	5	6	6	-

The dag is created and initialized before the parallel computation begins



Task dependency tree:



We define a structure called dag[][] to define the dependencies in OmpSs

Block		A <sub>11</sub>	A <sub>22</sub>	A <sub>33</sub>	A <sub>44</sub>	A <sub>55</sub>	A <sub>66</sub>	A <sub>77</sub>
Task		T <sub>0</sub>	T <sub>1</sub>	T <sub>2</sub>	T <sub>3</sub>	T <sub>4</sub>	$T_5$	$T_6$
Task id. <i>j</i>		0	1	2	3	4	5	6
Left descendant	dag[0][ <b>j</b> ]	0	1	2	3	0	2	4
Right descendant	dag[1][ <b>j</b> ]	0	1	2	3	1	3	5
Ancestor	dag[2][ <b>j</b> ]	4	4	5	5	6	6	6

The dag is created and initialized before the parallel computation begins



Task dependency tree:



Given the complexity of the data structure, we use a vector of sentinels to define the data dependency.



Task dependency tree:



Given the complexity of the data structure, we use a vector of sentinels to define the data dependency.



PARAVER trace of the preconditioner computation (32 leaves & 16 threads):



Analyzing the trace, we observe that:

- The leaf tasks (orange) are larger than the non-leaf tasks (green)
- Some small tasks are processed before the larger tasks

Therefore, some priorities could be defined to improve performances:

To process leaf tasks before non-leaf tasks

...



PARAVER trace of the preconditioner computation (32 leaves & 16 threads):



Analyzing the trace, we observe that:

- The leaf tasks (orange) are larger than the non-leaf tasks (green)
- Some small tasks are processed before the larger tasks

Therefore, some priorities could be defined to improve performances:

• To process leaf tasks before non-leaf tasks

• . .



Task dependency tree:



The priority could be used to improve performances



#### PARAVER trace without priority:



#### PARAVER trace with priority:

		-			
THREAD 1.1.1			1		
					-
181240 1.1.2					
THREAD 1.1.3			· · · · · · · · · · · · · · · · · · ·	(	
TUREAD 1 1 A			-	-	
CONTRACTOR AND THE		-			
THREAD 1.1.5					
THREAD 1.1.6		-		<u> </u>	-
		17 C			
THREAD 1.1.7					
THREAD 1.1.8					
TU2C40 1 1 0			-		
THACAD 1.1.9					-
THREAD 1.1.10				(	
THREAD 1.1.11					
Contractor and and					
THREAD 1.1.12					
THREAD 1.1.13		<u></u>			
	· · · · · · · · · · · · · · · · · · ·				
THREAD 1.1.14					
THREAD 1.1.15					
THE AL 1 1 16	· · · · · · · · · · · · · · · · · · ·	-	_		_
Inne/0 1.1.10					

ILUPACK with OMPSS

J.Aliaga, R.Badía, M.Barreda, M.Bollhöfer, E.Quintana-Ortí

# Outline



- Motivation and Introduction
- ILUPACK inverse-based MIC preconditioners
- 3 Parallel inverse-based MIC factorization preconditioners
- Experimental Results
- 5 Parallel Preconditioned Conjugate Gradient
- 6 Conclusions and Future Work

#### **Experimental Framework**

- Target platforms:
  - 4 AMD Opteron 6276 Processors (16 cores, total of 64):
    - 2.10 GHz
    - 64 Gbytes RAM
  - 2 Intel Xeon E5-2670 (8 cores, total of 16)
    - 2.60 GHz
    - 32 Gbytes RAM
  - All executions use up to 16 cores.
- IEEE double precision
- OmpSs version: Nanox 0.7a, mcxx 1.99

#### **Example Benchmark**

We consider the Laplacian PDE equation

 $-\Delta u = f$ 

in a 3D unit cube  $\Omega = [0, 1]^3$  with Dirichlet boundary conditions u = g on  $\partial \Omega$ 

For the discretization ...

- $\Omega$  replaced by a  $N_x \times N_y \times N_z$  uniform grid
- - $\rightarrow$  seven-point-star difference stencil

In this talk, we only consider the largest benchmark problem that fitted into the main memory of the target machines

$N_x  imes N_y  imes N_z$	п	nnz <sub>A</sub> /n		
$200\times200\times200$	8,000,000	7		



#### OMPSS preconditioner speed-up in AMD



Speed-up of preconditioner computation on AMD Opteron 6276 platform



### OMPSS preconditioner speed-up in INTEL



Speed-up of preconditioner computation on Intel Xeon E5-2670 platform



# Outline



- Motivation and Introduction
- ILUPACK inverse-based MIC preconditioners
- 3 Parallel inverse-based MIC factorization preconditioners
- 4 Experimental Results
- Parallel Preconditioned Conjugate Gradient
- 6 Conclusions and Future Work

#### Preconditioned Conjugate Gradient

Initialize k, x<sub>0</sub>, r<sub>0</sub>, z<sub>0</sub>, d<sub>0</sub>, 
$$\beta_0$$
,  $\tau_0$   
while ((k < k<sub>max</sub>)&&( $\tau_k > \tau_{min}$ ))  
 $w_k := Ad_k$   
 $\gamma_k := d_k^T w_k$   
 $\rho_k := \beta_k / \gamma_k$   
 $x_{k+1} := x_k + \rho_k d_k$   
 $r_{k+1} := r_k - \rho_k w_k$   
 $z_{k+1} := M^{-1}r_{k+1}$   
 $\alpha_k := \beta_k$   
 $\beta_{k+1} := r_{k+1}^T z_{k+1}$   
 $\alpha_k := \beta_{k+1} / \alpha_k$   
 $d_{k+1} := z_{k+1} + \alpha_k d_k$   
 $\tau_{k+1} := ||r_{k+1}||_2$   
 $k := k + 1$   
end while



#### Preconditioned Conjugate Gradient



Initialize k, x<sub>0</sub>, r<sub>0</sub>, z<sub>0</sub>, d<sub>0</sub>, 
$$\beta_0$$
,  $\tau_0$   
while ((k < k<sub>max</sub>)&&( $\tau_k > \tau_{min}$ )  
 $w_k := Ad_k$   
 $\gamma_k := d_k^T w_k$   
 $\rho_k := \beta_k / \gamma_k$   
 $x_{k+1} := x_k + \rho_k d_k$   
 $r_{k+1} := r_k - \rho_k w_k$   
 $z_{k+1} := M^{-1} r_{k+1}$   
 $\alpha_k := \beta_k$   
 $\beta_{k+1} := r_{k+1}^T z_{k+1}$   
 $\alpha_k := \beta_{k+1} / \alpha_k$   
 $d_{k+1} := z_{k+1} + \alpha_k d_k$   
 $\tau_{k+1} := ||r_{k+1}||_2$   
 $k := k + 1$   
end while

DOT AXPY AXPY Preconditioning

DOT

SPMV

AXPY-like 2-norm

#### Preconditioned Conjugate Gradient



Initialize 
$$k$$
,  $x_0$ ,  $r_0$ ,  $z_0$ ,  $d_0$ ,  $\beta_0$ ,  $\tau_0$   
while  $((k < k_{max})\&\&(\tau_k > \tau_{min})$   
 $w_k := Ad_k$   
 $\gamma_k := d_k^T w_k$   
 $\rho_k := \beta_k / \gamma_k$   
 $x_{k+1} := x_k + \rho_k d_k$   
 $r_{k+1} := r_k - \rho_k w_k$   
 $z_{k+1} := M^{-1}r_{k+1}$   
 $\alpha_k := \beta_k$   
 $\beta_{k+1} := r_{k+1}^T z_{k+1}$   
 $\alpha_k := \beta_{k+1} / \alpha_k$   
 $d_{k+1} := z_{k+1} + \alpha_k d_k$   
 $\tau_{k+1} := ||r_{k+1}||_2$   
 $k := k + 1$   
end while

 $\begin{array}{l} \text{SPMV} \rightarrow \text{on leaves} \\ \text{DOT} \rightarrow \text{on leaves} + \text{reduction} \end{array}$ 

 $\begin{array}{l} \text{AXPY} \ \rightarrow \text{on leaves} \\ \text{AXPY} \ \rightarrow \text{on leaves} \\ \text{Preconditioning} \ \rightarrow \text{cross. up} + \text{down} \end{array}$ 

 $\text{DOT } \rightarrow \text{on leaves} + \text{reduction}$ 

## Preconditioned Conjugate Gradient on OmpSs

パ

Task dependency tree:



We define a structure called dag[][] to define the dependencies in OmpSs

Block		A <sub>11</sub>	A <sub>22</sub>	A <sub>33</sub>	A <sub>44</sub>	A <sub>55</sub>	A <sub>66</sub>	A <sub>77</sub>
Task		T <sub>0</sub>	T <sub>1</sub>	T <sub>2</sub>	T <sub>3</sub>	T <sub>4</sub>	$T_5$	$T_6$
Task id. <i>j</i>		0	1	2	3	4	5	6
Left descendant	dag[0][ <b>j</b> ]	0	1	2	3	0	2	4
Right descendant	dag[1][ <b>j</b> ]	0	1	2	3	1	3	5
Ancestor	dag[2][ <b>j</b> ]	4	4	5	5	6	6	6

The dag is created and initialized before the parallel computation begins

ILUPACK with OMPSS

## Preconditioned Conjugate Gradient on OmpSs

パ

Task dependency tree:



In the forward resolution, the tree is crossed from bottom to top:
# Preconditioned Conjugate Gradient on OmpSs

パ

Task dependency tree:



In the backward resolution, the tree is crossed from top to bottom:

# Preconditioned Conjugate Gradient on OmpSs



THERE 1.1.1			100	MAR DA	11		i di		11	1		_	1				-	- 1	
THREAD 1.1.2					1				-	÷.	_	ii i		1	Í			11	-
THEEAD 1.1.3			100		1							-			in the second se		(		1
THREND 1.1.4		-		1111	í i 👘	- I				í ini			1				1		Υ.
THEEAD 1.1.5					1												111	í.	ω.
THREAD 1.1.6	( i i i i i i i i i i i i i i i i i i i											Ĭ.							
THREAD 1.1.T												_							
THREAD 1.1.0														,					
THEEAD 1.1.9										-									
THRAND 1.1.18								1.1											
DECKO 1 1 12															6	7 1 7		1.1	
THERE 1. 1. 1. 13				107 107			1.1.1						•••	i i		1.14	1.11		
THREAD 1.1.14				100 100		1		_	111	1				<u>,</u>	1	1.10	1000	111	11
THEEAD 1.1.13			1	1					1	1					i.			111	i i
THREAD 1.1.16				11 H		í.		1						i -			1		Ĩ.

# Preconditioned Conjugate Gradient on OmpSs



THEFEO 1 1.1	1		17 1.00	10.00				1		11	-		_	-						
	-		A 1947		_				_			-								
THREAD 1.1.2																				
THREAD 1.1.3																				
THREND 1.1.4			ALC: NO	100	10.1				_	_	11.1	-								11
												_								10.0
1991A0 1.1.5								_				_								
THREAD 1.1.6	1							_												
THREAD 1.1.T				1	- 1			-		11	1.1		_	1		i and	1		H 14	10.0
		 	Without a						- ee-	-										1.00
1993.40 1.1.9													-							
THEEAD 1.1.9																			100	
THREAD 1.1.10	1			100	10.07					1	-				1.0				1100	100.00
				- 22	1.00				1.1											
1991A0 1.1.11									-											
THEEAD 1.1.12															100 C			<b>1</b>		
THEEAD 1.1.13				100	100	-		1						11	1			11	10.00	1000
	10			1.11	200				_	10.0					100					110.1
1993.40 1.1.14																1000				
THEEAD 1.1.15	1							_		1								ПГ (		
THREAD 1.1.16	1		-	111	100		1	11	1						i i i			11		100

#### Merging small tasks:



バ

#### OMPSS solver speed-up in AMD





Speed-up of PCG solve on AMD Opteron 6276 platform

### OMPSS solver speed-up in INTEL





Number of leaves

# Parallel vs Sequential Solver



The conjugate gradient method minimizes at the jth step the energy norm of the error on the given j-dimensional Krylov subspace (Hestenes, 1952).

• We can estimate its value to analyze the behavior of each preconditioner.

#### Parallel vs Sequential Solver





## Conclusions and Future work

The ILUPACK parallelization with OmpSs ....

- Requires minor changes in the original code of ILUPACK
- Preserves the semantics of the method in ILUPACK
- Reduces the computation time by using the dynamic scheduling
- Improves the performances when # leaves are the twice # cores



# Conclusions and Future work

The ILUPACK parallelization with OmpSs ...

- Requires minor changes in the original code of ILUPACK
- Preserves the semantics of the method in ILUPACK
- Reduces the computation time by using the dynamic scheduling
- Improves the performances when # leaves are the twice # cores

Future work ...

- Investigate the effect of different options of the OmpSs runtime
  - $\rightarrow$  Schedules that may favor data locality
  - $\rightarrow$  Study how to improve the affinity in NUMA architectures
- Analyze the effect of power saving techniques on the performance
- Integrate energy saving mechanism into the OmpsS runtime scheduler
- Verify the distributed-memory and heterogeneous runtimes of OmpSs



#### Thanks for your attention !

Questions ?



# OMPSS preconditioner computation speed-up in AMD

Speed-up of preconditioner computation on AMD Opteron 6276 platform (16 cores)



# OMPSS preconditioner computation speed-up in INTEL

Speed-up of preconditioner computation on Intel Xeon E5-2670 platform (16 cores)



### OMPSS solver speed-up in AMD





Speed-up of PCG solve on AMD Opteron 6276 platform (16 cores)

### OMPSS solver speed-up in INTEL





Speed-up of PCG solve on Intel Xeon E5-2670 platform (16 cores)

ILUPACK with OMPSS

#### • The data is distributed on the leaves of the elimination tree:

• Consistently: the information is replicated:

$$X_{\mathcal{C}} = \begin{pmatrix} X_1 \\ X_2 \\ \hline X_3 \end{pmatrix} \rightarrow \begin{pmatrix} X_1 \\ \hline X_3 \end{pmatrix}, \begin{pmatrix} X_2 \\ \hline X_3 \end{pmatrix}.$$

$$b_u = \begin{pmatrix} b_1 \\ b_2 \\ \hline b_3 \end{pmatrix} \rightarrow \begin{pmatrix} b_1 \\ \hline b_3^1 \end{pmatrix}, \begin{pmatrix} b_2 \\ \hline b_3^2 \end{pmatrix}, b_3^1 + b_3^2 = b_3.$$

- The data is distributed on the leaves of the elimination tree:
  - Consistently: the information is replicated:

$$x_c = \begin{pmatrix} x_1 \\ x_2 \\ \hline x_3 \end{pmatrix} \rightarrow \begin{pmatrix} x_1 \\ \hline x_3 \end{pmatrix}, \begin{pmatrix} x_2 \\ \hline x_3 \end{pmatrix}.$$

$$b_u = \begin{pmatrix} b_1 \\ b_2 \\ \hline b_3 \end{pmatrix} \rightarrow \begin{pmatrix} b_1 \\ \hline b_3^1 \end{pmatrix}, \begin{pmatrix} b_2 \\ \hline b_3^2 \end{pmatrix}, b_3^1 + b_3^2 = b_3.$$

- The data is distributed on the leaves of the elimination tree:
  - Consistently: the information is replicated:

$$x_c = \begin{pmatrix} x_1 \\ x_2 \\ \hline x_3 \end{pmatrix} \rightarrow \begin{pmatrix} x_1 \\ \hline x_3 \end{pmatrix}, \begin{pmatrix} x_2 \\ \hline x_3 \end{pmatrix}.$$

$$b_u = \begin{pmatrix} b_1 \\ b_2 \\ \hline b_3 \end{pmatrix} \rightarrow \begin{pmatrix} b_1 \\ \hline b_3^1 \end{pmatrix}, \begin{pmatrix} b_2 \\ \hline b_3^2 \end{pmatrix}, b_3^1 + b_3^2 = b_3.$$

- The data is distributed on the leaves of the elimination tree:
  - Consistently: the information is replicated:

$$x_c = \begin{pmatrix} x_1 \\ x_2 \\ \hline x_3 \end{pmatrix} \rightarrow \begin{pmatrix} x_1 \\ \hline x_3 \end{pmatrix}, \begin{pmatrix} x_2 \\ \hline x_3 \end{pmatrix}.$$

$$\begin{split} b_{u} &= \begin{pmatrix} b_{1} \\ \underline{b_{2}} \\ \overline{b_{3}} \end{pmatrix} \rightarrow \begin{pmatrix} \underline{b_{1}} \\ \overline{b_{3}^{1}} \end{pmatrix}, \begin{pmatrix} \underline{b_{2}} \\ \overline{b_{3}^{2}} \end{pmatrix}, b_{3}^{1} + b_{3}^{2} = b_{3}. \\ A_{u} &= \begin{pmatrix} A_{11} & 0 & A_{31}^{T} \\ \underline{0} & A_{22} & A_{32}^{T} \\ \overline{A_{31}} & A_{32}^{T} & A_{33}^{T} \end{pmatrix} \rightarrow \begin{pmatrix} A_{11} & A_{31}^{T} \\ \overline{A_{31}} & A_{33}^{T} \end{pmatrix}, \begin{pmatrix} A_{22} & A_{32}^{T} \\ \overline{A_{32}} & A_{33}^{T} \end{pmatrix}, \\ A_{13}^{T} + A_{33}^{2} = A_{33}. \end{split}$$





#### • The distribution defines how the computations are made:

- The AXPY-like operations are made on similar operators.
- The DOT-like operations are made on different operators:
- The SPMV operations works as DOT-like operations:



- The distribution defines how the computations are made:
  - The AXPY-like operations are made on similar operators.
  - The DOT-like operations are made on different operators:
  - The SPMV operations works as DOT-like operations:

- The distribution defines how the computations are made:
  - The AXPY-like operations are made on similar operators.
  - The DOT-like operations are made on different operators:

$$b_u^T * x_c = \begin{pmatrix} b_1 \\ b_2 \\ \hline b_3 \end{pmatrix}^T * \begin{pmatrix} x_1 \\ x_2 \\ \hline x_3 \end{pmatrix}$$
$$b_u^T * x_c = \begin{pmatrix} b_1 \\ \hline b_3^1 \end{pmatrix}^T * \begin{pmatrix} x_1 \\ \hline x_3 \end{pmatrix} + \begin{pmatrix} b_2 \\ \hline b_3^2 \end{pmatrix}^T * \begin{pmatrix} x_2 \\ \hline x_3 \end{pmatrix}$$
$$b_u^T * x_c = val_1 + val_2$$

- The distribution defines how the computations are made:
  - The AXPY-like operations are made on similar operators.
  - The DOT-like operations are made on different operators:

$$b_u^T * x_c = \left(\frac{b_1}{b_2}\right)^T * \left(\frac{x_1}{x_2}\right) = val_1 + val_2 = val_u$$

- The distribution defines how the computations are made:
  - The AXPY-like operations are made on similar operators.
  - The DOT-like operations are made on different operators:

$$b_u^T * x_c = \left(\frac{b_1}{b_2}\right)^T * \left(\frac{x_1}{x_2}\right) = val_1 + val_2 = val_u$$

$$\begin{aligned} A_{u} * x_{c} &= \begin{pmatrix} A_{11} & 0 & | A_{31}^{T} \\ 0 & A_{22} & | A_{32}^{T} \\ \hline A_{31} & A_{32} & | A_{33} \end{pmatrix} * \begin{pmatrix} x_{1} \\ x_{2} \\ \hline x_{3} \end{pmatrix} \\ A_{u} * x_{c} &= \begin{pmatrix} A_{11} & | A_{31}^{T} \\ \hline A_{31} & | A_{31}^{1} \end{pmatrix} * \begin{pmatrix} x_{1} \\ \hline x_{3} \end{pmatrix} \oplus \begin{pmatrix} A_{22} & | A_{32}^{T} \\ \hline A_{32} & | A_{32}^{T} \end{pmatrix} * \begin{pmatrix} x_{2} \\ \hline x_{3} \end{pmatrix} \\ A_{u} * x_{c} &= \begin{pmatrix} v_{1} \\ \hline v_{3}^{T} \end{pmatrix} \oplus \begin{pmatrix} v_{2} \\ v_{3}^{T} \end{pmatrix} = v_{u} \end{aligned}$$

- The distribution defines how the computations are made:
  - The AXPY-like operations are made on similar operators.
  - The DOT-like operations are made on different operators:

$$b_u^T * x_c = \left(\frac{b_1}{b_2}\right)^T * \left(\frac{x_1}{x_2}\right) = val_1 + val_2 = val_u$$

$$A_{u} * x_{c} = \begin{pmatrix} A_{11} & 0 & | A_{31}^{T} \\ 0 & A_{22} & | A_{32}^{T} \\ \hline A_{31} & A_{32} & | A_{33} \end{pmatrix} * \begin{pmatrix} x_{1} \\ x_{2} \\ \hline x_{3} \end{pmatrix}$$
$$A_{u} * x_{c} = \begin{pmatrix} v_{1} \\ 0 \\ \hline v_{3}^{1} \end{pmatrix} + \begin{pmatrix} 0 \\ v_{2} \\ \hline v_{3}^{2} \end{pmatrix} = v_{u}$$



Initialize  $k, x_0, r_0, z_0, d_0, \beta_0, \tau_0$ while  $((k < k_{max})\&\&(\tau_k > \tau_{min}))$  $W_k := Ad_k$  $\gamma_k := \boldsymbol{d}_k^T \boldsymbol{w}_k$  $\rho_{\mathbf{k}} := \beta_{\mathbf{k}} / \gamma_{\mathbf{k}}$  $\mathbf{x}_{k+1} := \mathbf{x}_k + \rho_k \mathbf{d}_k$  $\mathbf{r}_{k+1} := \mathbf{r}_k - \rho_k \mathbf{W}_k$  $Z_{k+1} := M^{-1}r_{k+1}$  $\alpha_{\mathbf{k}} := \beta_{\mathbf{k}}$  $\beta_{k+1} := r_{k+1}^T z_{k+1}$  $\alpha_{k} := \beta_{k+1} / \alpha_{k}$  $d_{k+1} := z_{k+1} + \alpha_k d_k$  $\tau_{k+1} := ||r_{k+1}||_2$ k := k + 1end while

At the beginning:  $A_{\mu}$ ,  $x_{c}$ ,  $b_{\mu}$ Loop:  $d_c, w_u$  $\xrightarrow{reduction} \gamma_c$  $\gamma_{u}$  $\rho_c$  $X_{c}$ ru  $\xrightarrow{crossing} Z_c$  $\alpha_{c}$  $\beta_{II} \xrightarrow{\text{reduction}} \beta_{C}$  $\alpha_{c}$  $\xrightarrow{\text{crossing}}$   $\mathbf{r}_{c}$  ,  $\tau_{\mu} \xrightarrow{\text{reduction}} \tau_{c}$ 



Initialize  $k, x, r, z, d, \beta, \tau$ while  $((k < k_{max}) \& \& (\tau > \tau_{min}))$  $w_{\mu} := A_{\mu}d_{c}$  $\gamma_c := \boldsymbol{d}_c^T \boldsymbol{w}_{\mu}$  $\rho_c := \beta_c / \gamma_c$  $\mathbf{x}_c := \mathbf{x}_c + \rho_c \mathbf{d}_c$  $\mathbf{r}_{\mu} := \mathbf{r}_{\mu} - \rho_{c} \mathbf{W}_{\mu}$  $\alpha_c := \beta_c$  $Z_c := M^{-1} r_{\mu} \cdot r_c \leftarrow r_{\mu}$  $\tau_c := sart(r_u^T r_c)$  $\beta_c := r_u^T z_c$  $\alpha_c := \beta_c / \alpha_c$  $d_c := z_c + \alpha_c d_c$ k := k + 1end while

At the beginning:  $A_{\mu}$ ,  $x_{c}$ ,  $b_{\mu}$ Loop:  $\gamma_{\mu} \xrightarrow{\text{reduction}} \gamma_{c}$  $\xrightarrow{\text{crossing}} Z_c, \xrightarrow{\text{crossing}} I_c$  $\tau_{\mu} \xrightarrow{\text{reduction}} \tau_{c}$  $\beta_{\mu} \xrightarrow{\text{reduction}} \beta_{\sigma}$