Hierarchical Sparse Hybrid Solvers for Manycores Platforms

Emmanuel Agullo Luc GIRAUD Abdou GUERMOUCHE Stojce NAKOV Jean ROMAN

Third Workshop of the CNPq-Inria HOSCAR Project, September 2nd-5th, 2013





S. NAKOV

Inria Bordeaux - Sud-Ouest Research Center, France



Hierarchical Sparse Hybrid Solvers for Manycores Platforms

Outline

1. Introduction

2. Hybrid MPI + threads paralelization

3. Task-based paralelization

4. Conclusion and future work



Usual trades off

Direct

- Robust/accurate for general problems
- ★ BLAS-3 based implementations
- Memory/CPU prohibitive for large 3D problems

- Problem dependent efficiency / accuracy
- ★ Sparse computational kernels
- Less memory requirements and possibly faster



Usual trades off

Direct

- Robust/accurate for general problems
- ★ BLAS-3 based implementations
- Memory/CPU prohibitive for large 3D problems

- Problem dependent efficiency / accuracy
- ★ Sparse computational kernels
- Less memory requirements and possibly faster



Usual trades off

Direct

- ★ Robust/accurate for general problems
- ★ BLAS-3 based implementations
- Memory/CPU prohibitive for large 3D problems

- Problem dependent efficiency / accuracy
- ★ Sparse computational kernels
- Less memory requirements and possibly faster



Usual trades off

Direct

- ★ Robust/accurate for general problems
- ★ BLAS-3 based implementations
- Memory/CPU prohibitive for large 3D problems

- Problem dependent efficiency / accuracy
- ★ Sparse computational kernels
- Less memory requirements and possibly faster



Usual trades off

Direct

- ★ Robust/accurate for general problems
- ★ BLAS-3 based implementations
- Memory/CPU prohibitive for large 3D problems

- Problem dependent efficiency / accuracy
- ★ Sparse computational kernels
- ★ Less memory requirements and possibly faster

Sparse hybrid (direct/iterative) domain decomposition linear solvers

Partitioning the global matrix

* Global hybrid decomposition:

$$\mathcal{A} = egin{pmatrix} \mathcal{A}_{\mathcal{I}\mathcal{I}} & \mathcal{A}_{\mathcal{I}\Gamma} \ \mathcal{A}_{\Gamma\mathcal{I}} & \mathcal{A}_{\Gamma\Gamma} \end{pmatrix}$$

★ Global Schur complement:

$$\mathcal{S} = \mathcal{A}_{\Gamma\Gamma} - \mathcal{A}_{\Gamma\mathcal{I}}\mathcal{A}_{\mathcal{I}\mathcal{I}}^{-1}\mathcal{A}_{\mathcal{I}\Gamma}$$



Sparse hybrid (direct/iterative) domain decomposition linear solvers

Partitioning the global matrix

* Global hybrid decomposition:

$$\mathcal{A} = egin{pmatrix} \mathcal{A}_{\mathcal{I}\mathcal{I}} & \mathcal{A}_{\mathcal{I}\Gamma} \ \mathcal{A}_{\Gamma\mathcal{I}} & \mathcal{A}_{\Gamma\Gamma} \end{pmatrix}$$

* Global Schur complement:

$$\mathcal{S} = \mathcal{A}_{\Gamma\Gamma} - \mathcal{A}_{\Gamma\mathcal{I}}\mathcal{A}_{\mathcal{I}\mathcal{I}}^{-1}\mathcal{A}_{\mathcal{I}\Gamma}$$



Sparse hybrid (direct/iterative) domain decomposition linear solvers

Partitioning the global matrix

* Global hybrid decomposition:

$$\mathcal{A} = egin{pmatrix} \mathcal{A}_{\mathcal{I}\mathcal{I}} & \mathcal{A}_{\mathcal{I}\Gamma} \ \mathcal{A}_{\Gamma\mathcal{I}} & \mathcal{A}_{\Gamma\Gamma} \end{pmatrix}$$

* Global Schur complement:

$$\mathcal{S} = \mathcal{A}_{\Gamma\Gamma} - \mathcal{A}_{\Gamma\mathcal{I}}\mathcal{A}_{\mathcal{I}\mathcal{I}}^{-1}\mathcal{A}_{\mathcal{I}\Gamma}$$



Sparse hybrid (direct/iterative) domain decomposition linear solvers

Global preconditioner based solvers

- ★ HIPS (Inria, HiePACS)
- ★ PDSLin (LBNL, USA)
- ★ ShyLU (Sandia NL)

Local preconditioner based solver

★ MaPHyS (Inria, HiePACS)

Sparse hybrid (direct/iterative) domain decomposition linear solvers

Global preconditioner based solvers

- ★ HIPS (Inria, HiePACS)
- ★ PDSLin (LBNL, USA)
- ★ ShyLU (Sandia NL)

Local preconditioner based solver

* MaPHyS (Inria, HiePACS)

Method used in MAPHYS

Factorization of local interiors

* Global matrix:

$$\mathcal{A} = egin{pmatrix} \mathcal{A}_{\mathcal{I}\mathcal{I}} & \mathcal{A}_{\mathcal{I}\Gamma} \ \mathcal{A}_{\Gamma\mathcal{I}} & \mathcal{A}_{\Gamma\Gamma} \end{pmatrix}$$

* Global Schur Complement:

$$\mathcal{S}^{(i)} = \mathcal{A}_{\Gamma\Gamma}^{(i)} - \mathcal{A}_{\Gamma_i \mathcal{I}_i} \mathcal{A}_{\mathcal{I}_i \mathcal{I}_i}^{-1} \mathcal{A}_{\mathcal{I}_i \Gamma_i}$$



Method used in MAPHYS

Factorization of local interiors

* Local matrix:

$$\mathcal{A}^{(i)} = egin{pmatrix} \mathcal{A}_{\mathcal{I}_i \mathcal{I}_i} & \mathcal{A}_{\mathcal{I}_i \Gamma_i} \ \mathcal{A}_{\Gamma_i} \mathcal{I}_i & \mathcal{A}_{\Gamma\Gamma}^{(i)} \end{pmatrix}$$

★ Local Schur Complement:

$$\mathcal{S}^{(i)} = \mathcal{A}_{\Gamma\Gamma}^{(i)} - \mathcal{A}_{\Gamma_i \mathcal{I}_i} \mathcal{A}_{\mathcal{I}_i \mathcal{I}_i}^{-1} \mathcal{A}_{\mathcal{I}_i \Gamma_i}$$



Method used in MAPHYS

Constructing of the preconditioner

- * Local Schur Complement: $\mathcal{S}^{(i)} = \mathcal{A}_{\Gamma\Gamma}^{(i)} - \mathcal{A}_{\Gamma_{i}\mathcal{I}_{i}}\mathcal{A}_{\mathcal{T}_{i}\mathcal{T}_{i}}^{-1}\mathcal{A}_{\mathcal{I}_{i}\Gamma_{i}}$
- * Algebraic Additive Schwarz Preconditioner $\mathcal{M} = \sum_{i=1}^{N} \mathcal{R}_{\Gamma_i}^T (\bar{\mathcal{S}}^{(i)})^{-1} \mathcal{R}_{\Gamma_i}$ where $\bar{\mathcal{S}}^{(i)}$ is obtained from $\mathcal{S}^{(i)}$ via neighbor to neighbor communications.
- ★ Global Schur complement:
 - $\mathcal{S} = \sum_{i=1}^{N} \mathcal{R}_{\Gamma_i}^T(\mathcal{S}^{(i)}) \mathcal{R}_{\Gamma_i}$





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- ★ Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - MKL library
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MKL library for the reduced system





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - MKL librar
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MkL library for the reduced system





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - MKL library
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MKL library for the reduced system





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - MKL library
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MKL library for the reduced system





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - ► MKL library
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MKL library for the reduced system





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - MKL library
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MKL library for the reduced system





- Partitioning the global matrix in several local matrices
 - SCOTCH [Pellegrini and al.]
 - METIS [G. Karypis and V. Kumar]
- Local factorization
 - MUMPS [P. Amestoy and al.] (with Schur option)
 - PASTIX [P. Ramet and al.] (with Schur option)
- * Constructing of the preconditioner
 - MKL library
- * Solving the reduced system
 - CG/GMRES/FGMRES [V.Fraysse and L.Giraud] using MKL library for the reduced system





Outline

1. Introduction

2. Hybrid MPI + threads paralelization

3. Task-based paralelization

4. Conclusion and future work

Different type of paralelization of the sparse hybrid linear solvers

Global preconditioner based solvers

- ★ HIPS
- ⋆ PDSLIN
- ★ SHYLU

Local preconditioner based solver

★ MAPHYS



Different type of paralelization of the sparse hybrid linear solvers

Global preconditioner based solvers

- ★ HIPS
- ★ PDSLIN
- ★ SHYLU

Local preconditioner based solver * MAPHYS Process MPI Domain Domain Node 1 S. NAKOV Market Solvers for Manycores Platforms

Different type of paralelization of the sparse hybrid linear solvers

Global preconditioner based solvers

- ★ HIPS
- ⋆ PDSLIN
- ★ SHYLU



Different type of paralelization of the sparse hybrid linear solvers

Global preconditioner based solvers

- ★ HIPS
- ★ PDSLIN
- ★ SHYLU

Different type of paralelization of the sparse hybrid linear solvers

Global preconditioner based solvers

- ★ HIPS
- * PDSLIN
- ★ Shyll

Local preconditioner based solver

★ MAPHYS (before)



Different type of paralelization of the sparse hybrid linear solvers

Global preconditioner based solvers

- ★ HIPS
- * PDSLIN
- ★ SHYLU

Local preconditioner based solver

★ MAPHYS (this study)



Software used in MAPHYS (before)

Partitioners

- ⋆ SCOTCH
- ★ Metis

Dense direct solver

★ MKL library

Sparse direct solvers

- ★ Mumps
- ★ PASTIX

Iterative Solvers

★ CG/GMRES/FGMRES using MKL library

Software used in MAPHYS (this study)

Dense direct solver

★ Multi-threaded MKL library

Sparse direct solvers

- ★ MUMPS
- ★ Multi-threaded PASTIX

Iterative Solvers

* CG/GMRES/FGMRES using multi-threaded MKL library

Software used in MAPHYS (this study)

Dense direct solver

★ Multi-threaded MKL library

Sparse direct solvers

- ★ Mumps
- ★ Multi-threaded PASTIX

Iterative Solvers

★ CG/GMRES/FGMRES using multi-threaded MKL library

Challenge

★ Composability, Performance

Experimental set up

Hardware (on each node)

- * Two Quad-core Nehalem Intel® Xeon® X5550
- ★ Memory: 24 GB GDDR3
- ★ Double precision

Matrices								
	Haltere	Audikw_1	tdr455k					
N	1,288,825	943,695	2,738,556					
nnz	10,476,775	39,297,771	112,756,352					
Precision	Complex	Real	Complex					
Description	Electromagnetic problem	3D unstructured	Electromagnetic problem					
		structural mechanics						









Scalability of MAPHYS on multicore nodes

Achieved performance with the Audi matrix on four nodes

All computational steps



Factorization step



Hierarchical Sparse Hybrid Solvers for Manycores Platforms

Flexibility to exploit entire multicore nodes





Flexibility to exploit entire multicore nodes





Node 4

Flexibility to exploit entire multicore nodes



Flexibility to exploit entire multicore nodes



Flexibility to exploit entire multicore nodes

Achieved performance with the Haltere matrix on two nodes



Max

Avg

Flexibility to exploit entire multicore nodes

Achieved performance with the tdr455k matrix on sixteen nodes



Memory used per node



Outline

1. Introduction

2. Hybrid MPI + threads paralelization

3. Task-based paralelization

4. Conclusion and future work

Task-based paralelization

Task-based paralelization

Factorization step

Constructing of the preconditioner



Solve

Task-based paralelization

Task-based paralelization



Approach: Three-layer paradigm

High-level (task-based) algorithm

Runtime (task-based) System STARPU, PARSEC, QUARK, SuperMatrix ...

Device kernels

Approach: Three-layer paradigm

High-level (task-based) algorithm

Runtime (task-based) System STARPU, PARSEC, QUARK, SuperMatrix ...

Device kernels

Approach: Three-layer paradigm

High-level (task-based) algorithm

Runtime (task-based) System STARPU, PARSEC, QUARK, SuperMatrix ...

Device kernels

Runtime systems in linear algebra

Dense linear algebra libraries

- ★ PLASMA (QUARK)
- ★ DPLASMA (PARSEC)
- ★ MAGMA-MORSE (STARPU)
- ★ FLAME (SUPERMATRIX)

Sparse linear algebra libraries

- ★ PASTIX (STARPU and PARSEC)
- ★ qr_mumps (STARPU and PARSEC)

Runtime systems in linear algebra

Dense linear algebra libraries

- ★ PLASMA (QUARK)
- ★ DPLASMA (PARSEC)
- * MAGMA-MORSE (STARPU)
- ★ FLAME (SUPERMATRIX)

Sparse linear algebra libraries

- * PASTIX (STARPU and PARSEC)
- ★ qr_mumps (STARPU and PARSEC)

Dense direct solver

★ MAGMA-MORSE library

Sparse direct solvers

* PASTIX

Iterative Solvers

★ CG/GMRES/FGMRES

Dense direct solver

★ MAGMA-MORSE library

with STARPU

Sparse direct solvers

★ PASTIX

with STARPU

Iterative Solvers

★ CG/GMRES/FGMRES

Dense direct solver

★ MAGMA-MORSE library

with STARPU

Sparse direct solvers

★ PASTIX

with STARPU

Iterative Solvers * Cg/GMRES/FGMRES

Tackifying the CG Algorithm

CG Algorithm

 $\begin{array}{l} \begin{array}{l} 1: r \leftarrow b \\ 2: r \leftarrow r - Ax \\ 3: p \leftarrow r \\ 4: \delta_{new} \leftarrow dot(r, r) \\ 5: \delta_{old} \leftarrow \delta_{new} \\ 6: \text{ while } \frac{||b-Ax||}{||b||} \leq eps \text{ do} \\ 7: q \leftarrow Ap \\ 8: \alpha \leftarrow \delta_{new}/dot(p, q) \\ 9: x \leftarrow x + \alpha p \\ 10: r \leftarrow r - \alpha q \\ 11: \delta_{new} \leftarrow dot(r, r) \\ 12: \beta \leftarrow \delta_{new}/\delta_{old} \\ 13: \delta_{old} \leftarrow \delta_{new} \\ 14: p \leftarrow r + \beta p \\ 15: \text{ end while} \end{array}$

Tackifying the CG Algorithm

CG Algorithm

 $\begin{array}{l}
\mathbf{1:} r \leftarrow b \\
\mathbf{2:} r \leftarrow r - Ax \\
\mathbf{3:} p \leftarrow r
\end{array}$ **4:** $\delta_{new} \leftarrow dot(r, r)$ 5: $\delta_{old} \leftarrow \delta_{new}$ 6: while $\frac{\|b-Ax\|}{\|b\|} \leq eps$ do 7: $q \leftarrow Ap$ 8: $\alpha \leftarrow \delta_{new}/dot(p,q)$ 9: $x \leftarrow x + \alpha p$ 10: $r \leftarrow r - \alpha q$ **11:** $\delta_{new} \leftarrow dot(r, r)$ 12: $\beta \leftarrow \delta_{new} / \delta_{old}$ 13: $\delta_{old} \leftarrow \delta_{new}$ 14: $p \leftarrow r + \beta p$ 15: end while

Tackifying the CG Algorithm

Main CG loop1: while $\frac{\|b-Ax\|}{\|b\|} \le eps$ do2: $q \leftarrow Ap$ 3: $\alpha \leftarrow \delta_{new}/dot(p,q)$ 4: $x \leftarrow x + \alpha p$ 5: $r \leftarrow r - \alpha q$ 6: $\delta_{new} \leftarrow dot(r, r)$ 7: $\beta \leftarrow \delta_{new}/\delta_{old}$ 8: $\delta_{old} \leftarrow \delta_{new}$ 9: $p \leftarrow r + \beta p$ 10: end while



Execution of CG ON 3 GPUs





Trace of one iteration of CG with 3 GPUs.

S. NAKOV

Hierarchical Sparse Hybrid Solvers for Manycores Platforms

Execution of CG ON 3 GPUs





1: while $\frac{\|b-Ax\|}{\|b\|} \le eps$ do 2: $q \leftarrow Ap$ 3: $\alpha \leftarrow \delta_{new}/dot(p,q)$ 4: $x \leftarrow x + \alpha p$ 5: $r \leftarrow r - \alpha q$ 6: $\delta_{new} \leftarrow dot(r,r)$ 7: $\beta \leftarrow \delta_{new}/\delta_{old}$ 8: $\delta_{old} \leftarrow \delta_{new}$ 9: $p \leftarrow r + \beta p$ 10: end while

Trace of one iteration of CG with 3 GPUs after optimizing the task-flow.

S. NAKOV

Performance



	Basic CG			Optimized CG		
Matrix	1GPU	3GPU	Speed-up	1GPU	3GPU	Speed-up
	(Gflop/s)	(Gflop/s)		(Gflop/s)	(Gflop/s)	
Audikw_1	9.3	12.97	1.39	9.7	16.1	1.7
11ptd-256-256-256	5.24	7.5	1.43	5.7	16.27	2.85

Data transfers between GPUs



Benefit of using runtime systems







GPU-CPU-GPU communication mechanism

Benefit of using runtime systems







GPU-GPU communication mechanism

Outline

1. Introduction

2. Hybrid MPI + threads paralelization

3. Task-based paralelization

4. Conclusion and future work

Conclusion and future work

Conclusion

- * Two-level parallelism efficiently exploits multicore architectures
- * Synchronizations may be a bottleneck
- * Investigating task-based paralelization to pipeline all the steps

Future work

- ★ Full task-based solver
- * Further experiments in the collaboration with TOTAL