MUMPS User Days

4th edition

MUMPS group

CERFACS, CNRS, ENS-Lyon, INRIA, INPT, University of Bordeaux

MUMPS User Days — Montbonnot, June 1-2, 2017

MUMPS overview and recent features

MUMPS group:
CERFACS, CNRS, ENS-Lyon, INRIA, INPT, Univ. Bordeaux

Contents of the presentation:

- MUMPS overview, history, statistics on usage, research links
- Main recent features, software releases
- MUMPS User days 2017

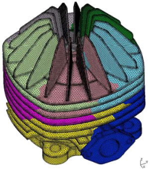

Discretization of a physical problem
(e.g. Code_Aster, finite elements)

$\Longrightarrow$

Solution of sparse systems
$AX = B$

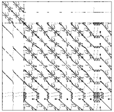

Often the most expensive part in numerical simulation codes

Sparse direct methods:

- Analyse graph of matrix, permutation, memory estimates
- Factor $A = LU$ (LDL^t if A symmetric) using Gaussian elimination
- Triangular solve: $LY = B$, then $UX = Y$

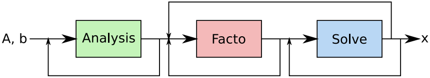

Sometimes preferred to iterative methods for their robustness and ability to solve efficiently multiple/successive right-hand sides.

MUMPS: a MUltifrontal Massively Parallel Solver

Solve **A X = B**,
A is a large sparse matrix, and **B** is dense or sparse
on multiprocessor architectures

MUMPS Background

- Multifrontal methods: Duff, Reid'83
- 1996-1999: MUMPS started in Toulouse from a distributed-memory prototype inspired from a shared memory research code

  Context: European project PARASOL (PARAllel SOLvers, 10 partners, direct and iterative methods developers, industrial end-users, software companies)
- 2000: First "public domain" version of MUMPS
- 2013: Third edition of MUMPS User Group Meeting (EDF-Clamart)

MUMPS: a MUltifrontal Massively Parallel Solver

Solve **A X = B**,

A is a large sparse matrix, and **B** is dense or sparse

on multiprocessor architectures

MUMPS Background

- Multifrontal methods: Duff, Reid'83
- 1996-1999: MUMPS started in Toulouse from a distributed-memory prototype inspired from a shared memory research code

  Context: European project PARASOL (PARAllel SOLvers, 10 partners, direct and iterative methods developers, industrial end-users, software companies)
- 2000: First "public domain" version of MUMPS
- 2013: Third edition of MUMPS User Group Meeting (EDF-Clamart)
- 2014: Consortium of MUMPS users

  Founding members: *CERFACS, INPT, Inria, ENS-Lyon, Bordeaux University*
- 2015: MUMPS 5.0.0, first CeCILL-C version of MUMPS
- 2017: Fourth edition of MUMPS User Group Meeting (Inria Montbonnot)

MUMPS consortium

... to ensure software sustainability and development

Consortium (2014-2022, http://mumps-consortium.org/)

- Manager: Inria; President of Executive Committee: INP Toulouse
- Membership agreement stipulates Member's rights:
  - experiment with versions in advance (latest upgrades, beta releases)
  - exert an influence over future developments and the interface of new features
  - appoint a representative to annual meeting of the Consortium Committee
  - priority access to developers: support, advice, performance analysis which may give rise to a specific study agreement
- Members (10):
  - *EDF, Altair, Michelin, LSTC (USA), SISW-Siemens (Belgium), FFT-MSC Soft. (Belgium), ESI Group, Total, SAFRAN, Lawrence Berkeley Nat. Lab. (USA)*

Membership fees → funding for PhD and engineers

The MUMPS solver – http://mumps-solver.org

- Co-developed in France (Toulouse, Lyon, Bordeaux) by CERFACS, CNRS, ENS Lyon, INPT, Inria, Bordeaux Univ.
  - Address wide classes of problems: various types of matrices/formats, numerical pivoting, many numerical features
  - Asynchronous approach to parallelism

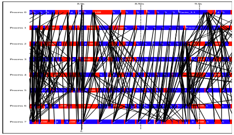

MUMPS

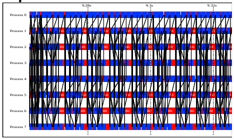

SuperLU_dist

- Software package used worldwide in academic research, R&D departments, and also through
  - commercial software: (*Samcef* from Samtech/Siemens, *Actran* from Free Field Technologies/MSC, *PAM-Crash* from ESI-Group, *Flux* from Altair, *OptiStruct* from Altair, *COMSOL MultiPhysics* from COMSOL, . . .).
  - open-source and research packages: *Code_Aster* (EDF), *IPOPT*, *Petsc*, *Trilinos*, *FreeFEM++*, *OpenSEES* , SOPALE, Kwant, . . .
  - Linux distributions: *Debian*, *CentOS*, . . .

Software download requests: countries around the world

requests 4.2 - 4.6.4 (Dec 2002 - Apr 2007)

requests 4.7 - 4.8 (Apr 2007 - Jul 2009)

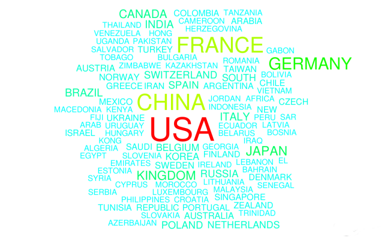

requests 4.9 - 4.10 (Jul 2009 - Feb 2015)

requests 5.0 - 5.1 (Feb 2015 - May 2017)

See requests from the US, Germany, Japan, China

World maps

- 14 232 download requests from Dec 2002 to May 2017 from our website

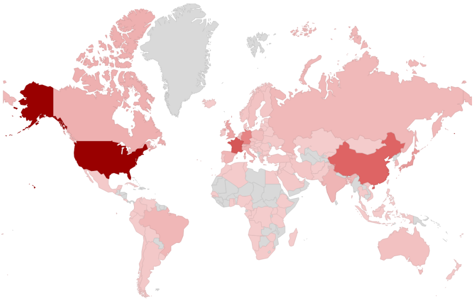

- 177 612 visitors (112 090 unique visits) on our website from Nov 2010 to May 2017

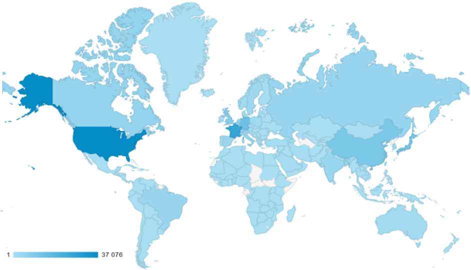

Software download requests: Application Fields

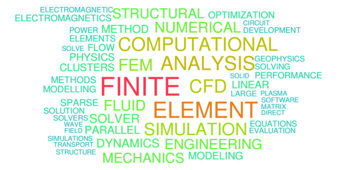

requests 4.2 - 4.6.4 (Dec 2002 - April 2007)

requests 4.7 - 4.8 (April 2007 - July 2009)

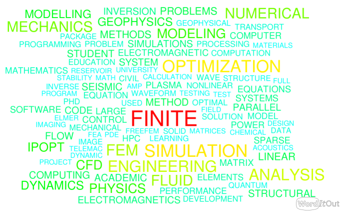

requests 4.9 - 4.10 (July 2009 - Feb 2015)

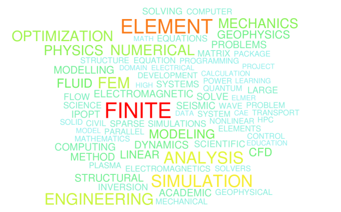

requests 5.0 - 5.1 (Feb 2015 - May 2017)

Users' mailing list (mumps-users@listes.ens-lyon.fr)

- 550 subscribers, $\sim$ 1 message per day on average

Main topics of exchanges between users (2013-2017)

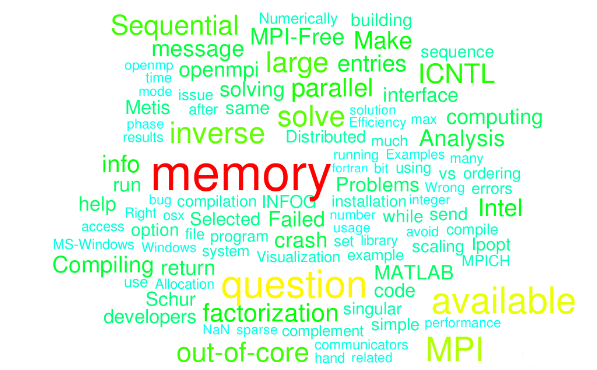

Research links

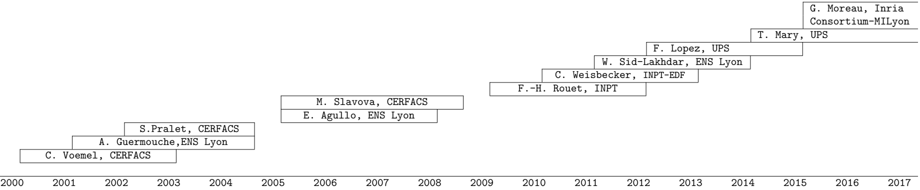

- Robust Memory-Aware Mappings (memory scalability and quality of memory estimates) (PhD Agullo, ENS Lyon, 2005-2008 and Rouet, INPT-IRIT, 2009-2012)
- Shared and distributed memory parallelism on NUMA clusters (initiated with PhD Sid-Lakdhar, ENS Lyon, 2011-2014)
- Synchronisation avoidance and deadlock prevention, in context of dynamic distributed scheduling with asynchronous p2p & collective communications (initiated with PhD's Rouet and Sid-Lakhdar)

Research links

- Block Low-Rank (BLR):
  algebraic solver based on BLR approximation (PhDs Weisbecker, INPT-IRIT, EDF funding, 2010-2013 and Mary, UPS-IRIT, 2014-2017);
  - Collab. O. Boiteau (EDF), C. Ashcraft (LSTC, Livermore, USA)
    $\rightarrow$ See talk by Théo Mary (Toulouse University)
  - Application to geophysics applications (SEISCOPE, EMGS)
    $\rightarrow$ See talk by Daniil Shantsev (EMGS, Norway)

- Performance of solution phase (PhDs Rouet INPT-IRIT, 2009-2012 and Moreau, ENS Lyon, 2015-): entries of A^{-1}, exploit sparsity of right-hand sides/partial solution, performance of (BLR) solve, ...
  $\rightarrow$ See talk by Gilles Moreau this afternoon
- Continuous collaborations and feedback from applications
  ...crucial for MUMPS future research and developments

Research links

Scientific themes for recent papers:

- Agullo, Amestoy, Buttari, Guermouche, L'Excellent, Rouet: Robust memory-aware mappings, SIAM Journal on Scientific Computing (2016).
- Amestoy, Buttari, L'Excellent, Mary: Theoretical complexity of BLR+practical validation, to appear in SIAM Journal on Scientific Computing.
- Amestoy, Buttari, L'Excellent, Mary: Implementation, Performance, scalability of multithreaded BLR, ACM TOMS (submitted)
- EMGS Norway, Amestoy, Buttari, L'Excellent, Mary: Application and performance of BLR to 3D EM modeling, to appear in Geophysical J. Inter.
- Amestoy, Brossier, Buttari, L'Excellent, Mary, Métivier, Miniussi, Operto: Application and performance of BLR to 3D full waveform inversion: Geophysics, 2016

Strong interaction with software work: M. Durand, G. Joslin, C. Puglisi (Inria, supported by MUMPS consortium)

performance tuning, scalability studies, parallel performance with respect to other solvers, multithreading, reduce memory consumption, validation on real-life applications, development/stabilization of new features, follow-up applications feedback/consortium/users' community, user support, etc.

Software history (1996-2013)

- 22 internal PARASOL releases, 1996→ MUMPS 4.0.4 (1999): LU, LDL^T, elemental input, distributed matrix input, Schur complement
- MUMPS 4.1.6 (2000): first freely distributed stabilized release!
- MUMPS 4.2 beta (2002), MUMPS 4.3 (2003): "SDCZ" arithmetics, many orderings (Scotch, Pord, Metis, AMD, QAMD, AMF), candidate processors (PhD Voemel), multiple Right-Hand sides (RHS), inertia
- MUMPS 4.5 (2005): progress on symmetric indefinite matrices (PhD Pralet), 2D block-cyclic Schur complement, first API for sparse RHS, distributed solution
- MUMPS 4.6 (2006), MUMPS 4.7 (2007): hybrid scheduling (PhD's Guermouche+Pralet), reduced/condensed RHS, detection of zero pivots
- MUMPS 4.8 (2008): Parallel scalings (postdoc Uçar), memory reductions, out-of-core (PhD's Agullo+Slavova)
- MUMPS 4.9 (2009), MUMPS 4.10 (2011): Parallel analysis (postdoc Buttari), 64-bit addressing for factors, A^{-1} entries (PhD's Slavova+Rouet), determinant (collaboration A. Salzman)

Software releases since last User Days

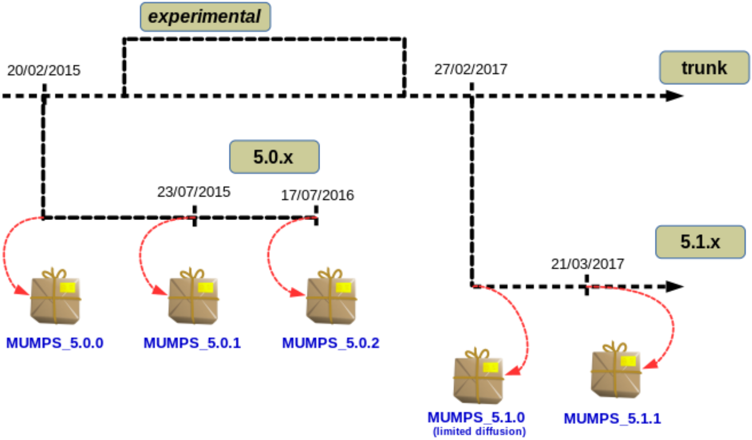

MUMPS 5.0.0 (2015): a major release

- First version under Cecill-C license
- Userguide considerably improved/redesigned
- New features:
  - First version with OpenMP directives (significant performance gains)
  - Forward elimination during factorization, use workspace from user, deterministic parallel analysis, ...
  - Solve phase revisited (memory scalability and performance)
- Evolutions: MUMPS 5.0.1, MUMPS 5.0.2 (stabilization of 5.0.0 and performance improvements for specific matrices)
- Received positive feedback from users, e.g.:
  - *"You know, you've made a huge contribution to the scientific community here. A modern, parallel sparse linear solver that runs on pretty much any platform is enormously useful."*
  - *"The OpenMP is definitely faster overall, especially with larger problems with a few hundred thousand nonzeroes and up"*

MUMPS 5.0.0 vs MUMPS 4.10.0: user feedback

- Computation of $A^{-1}B$ (B sparse) by blocks of 32: 72 MPI processes

	Factorization (seconds)	Solve (seconds)
MUMPS 4.10.0	158.9	13923.3
MUMPS 5.0.0	60.3	9806.0

- Time for factorization (shift and invert method, numerical issues). 360 cores (36 MPI and 10 threads/MPI).

	Factorization Phase (LU) (seconds)	Factorization Phase (LDLT) (seconds)
MUMPS 4.10.0	568	652
MUMPS 5.0.0	388	294

MUMPS 5.1.1 (March 2017)

- A release with two new major features:
  - 64-bit integers where needed (O(NZ) data on top of O(|L|) data)
    - Specifications guided by industrial partners, backward-compatible,
    - Metis, Scotch, PORD → 64-bit integers (32-bit also possible)
    - full 64-bit integer version also possible but more resource consuming

    → See presentation by Kostas Sikelis (Altair)
  - First public version with low-rank compression (i.e. BLR) (work initiated in 2010!) → See presentation by Théo Mary on Block-Low-Rank and HSS multifrontal solvers
- Many other issues concerning robustness and performance (e.g., of solve phase with many RHS) → See presentation by Marie Durand

. . . and quite a lot of ongoing work . . .

Timings in seconds on 900 cores (90 MPI x 10 threads)			
EOS computer, CALMIP mesocenter (https://www.calmip.univ-toulouse.fr/)			
Full Rank		Block Low Rank	
5.1.1	FR +	5.1.1	BLR +
3D Full Waveform Inversion (Helmholtz equations) N=17 M			
		BLR precision $\epsilon_{BLR} = 10^{-3}$	
937	**548**	**267**	**206**
3D Electromagnetism (Maxwell equations) N=21 M			
		BLR precision $\epsilon_{BLR} = 10^{-7}$	
2 587	**1 255**	**486**	**319**
3D Structural Mechanics N=8 M			
		BLR precision $\epsilon_{BLR} = 10^{-9}$	
722	**266**	**199**	**117**

→ See closing presentation today *"MUMPS perspectives and discussions"*

The 2017 MUMPS User Days

Objectives:

- Bring together some MUMPS users from both academia and industry and MUMPS developers (61 participants from 10 countries)
- Have time for discussions and informal exchanges
- Present some aspects of MUMPS activities by MUMPS group members / share users' feedback and experience with MUMPS → *cf. next 2 slides*
- And also:
  - benefit from experts knowledge on impact of computer evolutions:
    →presentation by François Courteille (NVIDIA, France) this afternoon
    →presentation by Patrick Demichel (HPE, France) tomorrow
  - share experience with other developers of sparse solvers:
    - sparse direct solver "MF2" → presentation by Bob Lucas tomorrow
    - domain decomposition methods (often using direct methods):
      → presentation by Pierre Jolivet (CNRS, France) tomorrow
      → presentation by Augustin Parret-Fréaud (SAFRAN, France) tomorrow

MUMPS group members present today & presentations

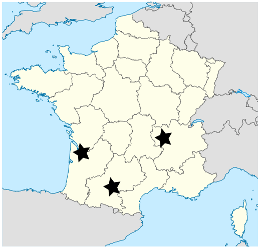

General presentations/discussions:

- Overview and recent features (now)
- MUMPS perspectives (last talk today)
- Closing discussion (tomorrow afternoon)

Members on permanent academic positions:

- Patrick Amestoy (INPT-IRIT, Toulouse)
- Jean-Yves L'Excellent (Inria-LIP, Lyon)
- Abdou Guermouche (LaBRI, Bordeaux)
- Alfredo Buttari (CNRS-IRIT, Toulouse) → qr_mumps: a runtime-based Sequential Task Flow parallel solver (this afternoon)

Engineers:

- Guillaume Joslin (MUMPS Consortium, Inria, Lyon)
- Chiara Puglisi (MUMPS Consortium, Inria, Toulouse)
- Marie Durand (MUMPS Consortium, Inria, Lyon) → Discussion of MUMPS parallel performance in multithreaded environments (this morning)

PhD Students:

- Théo Mary (UPS, Toulouse) → On the comparison of sparse multifrontal hierarchical and Block Low-Rank solvers (this morning)
- Gilles Moreau (MUMPS Consortium, MILyon) → Recent advances on the solution phase of direct solvers with multiple sparse right-hand sides (this afternoon)

MUMPS Users' presentations: summary

- Olivier Boiteau (EDF, France): Use of MUMPS in EDF codes (thermomechanics, material structure, electromagnetics, hydrodynamics)
- Kostas Sikelis (Altair, Greece): Comparison of 32bit vs 64bit integer MUMPS in Optistruct (linear and nonlinear structural and thermal analysis, . . .)
- Daniil Shantsev (EMGS, Norway): Large-scale 3D Controlled source EM modeling with a Block Low-Rank MUMPS solver
- Eveline Rosseel (FFT-MSC Software Belgium): Improving (aero/vibro-)acoustic simulations using MUMPS - Evaluation of Block Low-Rank factorizations
- Rémy Perrin-Bit (Altair, France): Brief history of time in FLUX (Electromagnetic and thermal simulations)
- Luis E. García Castillo (University Carlos III of Madrid, Spain): Higher-Order Finite Element Code for Electromagnetic Simulation
- Yuri Feldman (Ben-Gurion University, Israel): Two phase flow simulations based on Immersed boundary method, by utilizing MUMPS solver

Enjoy those two days!

MUMPS User Days
Thursday, June 1st and Friday, June 2nd 2017

Inria centre, Montbonnot Saint-Martin (near Grenoble, France)

Programme

Thursday, June 1st

8.30 - 8.45	Registration and welcome coffee
8.45 - 9.00	Welcome and presentation of the two day meeting
9.00 - 9.30	Patrick Amestoy (INPT(ENSEEIHT)-IRIT), Abdou Guermouche (Univ. de Bordeaux), Jean-Yves L'Excellent (Inria-LIP-ENS Lyon) MUMPS overview and recent features
9.30 - 10.00	Olivier Boiteau (EDF Lab Paris-Saclay, France) Feedback in the use of MUMPS in EDF codes
10.00 - 10.30	Théo Mary (University of Toulouse, France) On the comparison of sparse multifrontal hierarchical and Block Low-Rank solvers
10.30 - 11.00	Coffee Break
11.00 - 11.30	Eveline Rosseel (FFT-Msc_Software, Belgium) Improving (aero/vibro-)acoustic simulations using MUMPS: evaluation of Block Low-Rank factorizations
11.30 - 12.00	Marie Durand (MUMPS Consortium/Inria, France) Discussion of MUMPS parallel performance in multithreaded environments
12.00 - 12.30	Kostas Sikelis (Altair, Greece) Comparison of 32bit vs 64bit integer MUMPS in Optistruct
12.30 - 14.30	Lunch
14.30 - 15.00	Daniil Shantsev (EMGS, Norway) Large-scale 3D Controlled source EM modeling with a Block Low-Rank MUMPS solver
15.00 - 15.30	Gilles Moreau (MUMPS Consortium/LabEx MILYON/Inria, France) Recent advances on solution phase of sparse solvers with multiple RHS
15.30 - 16.00	Break
16.00 - 16.30	Alfredo Buttari (CNRS, France) qr_mumps: a runtime-based Sequential Task Flow parallel solver
16.30 - 17.00	François Courteille (NVIDIA, France) Programming heterogeneous architecture with libraries: a survey of NVIDIA linear algebra libraries
17.00 - 17.30	Patrick Amestoy (INPT(ENSEEIHT)-IRIT), Abdou Guermouche (Univ. de Bordeaux), Jean-Yves L'Excellent (Inria-LIP-ENS Lyon) MUMPS 5.1, perspectives and discussions
19.30 - 22.00	Banquet at "Le Garage" (134 Chemin de l'étoile 383330 Montbonnot)

Friday, June 2nd

9.00 - 9.30	Bob Lucas (LSTC, USA) Block Low-Rank approximations in LS-DYNA
9.30 - 10.00	Augustin Parret-Fréaud (SAFRAN, France) Robust domain decomposition methods for high performance computation of industrial structures
10.00 - 10.30	Pierre Jolivet (CNRS, France) MUMPS on thousands of cores: feedback on the use of direct solvers in domain decomposition methods
10.30 - 11.00	Coffee Break
11.00 - 11.30	Rémy Perrin-Bit (Altair, France) Brief history of time in Flux
11.30 - 12.00	Luis E. García Castillo (University Carlos III of Madrid, Spain) Higher-Order Finite Element Code for Electromagnetic Simulation on HPC Environments
12.00 - 12.30	Patrick Demichel (HPE, France) The Machine and genZ implications for extreme scale solver problems
12.30 - 14.00	Lunch
14.00 - 14.30	Yuri Feldman (Ben-Gurion University, Israel) Two phase flow simulations based on Immersed boundary method, by utilizing MUMPS solver
14.30 - 15.00	Closing session (MUMPS team)

Credits

This event is supported by:

INPT-ENSEEIHT Inria MUMPS-Consortium IRIT

Labex Milyon

Discussion of MUMPS parallel performance in multithreaded (MT) environments

Marie Durand
(Inria-MUMPS Consortium)

Performance study

Presentation of the study

- focus on the factorization phase (few results on solve)
- a lot of matrices tested
- comparison with direct solvers dedicated to MT environments
  - MKL Pardiso
  - HSL ma86, HSL ma87
- MUMPS evaluated over several #MPI × #OpenMP threads
- LIP grunch computer: 2x14 cores @2.30GHz - Intel Xeon E5-2695 v3

MUMPS vs others - scalability study

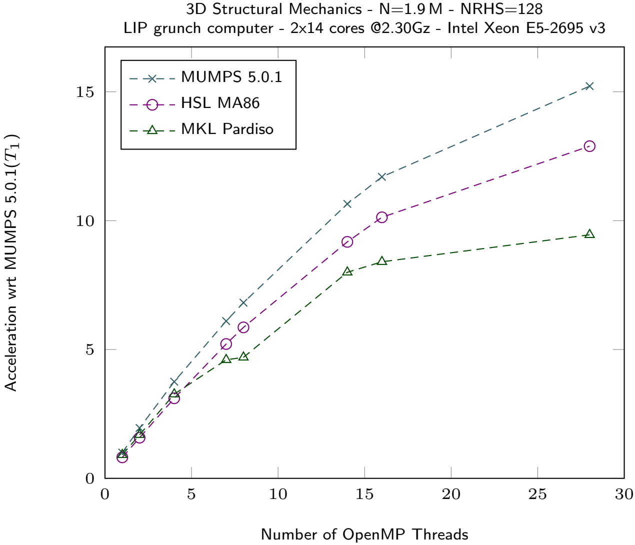

MUMPS vs Others

Each point corresponds to a matrix and to the best time obtained on each configuration of cores.

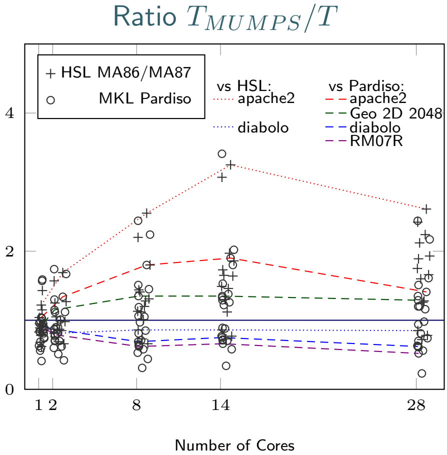

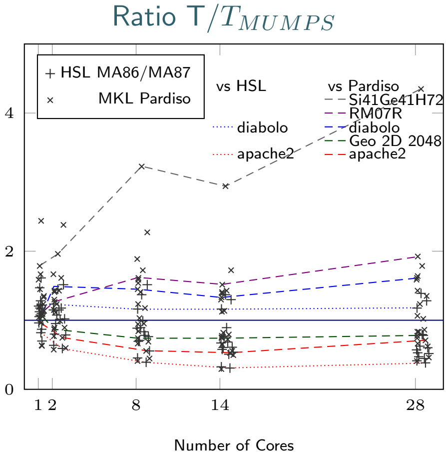

Using MUMPS in a shared memory environment

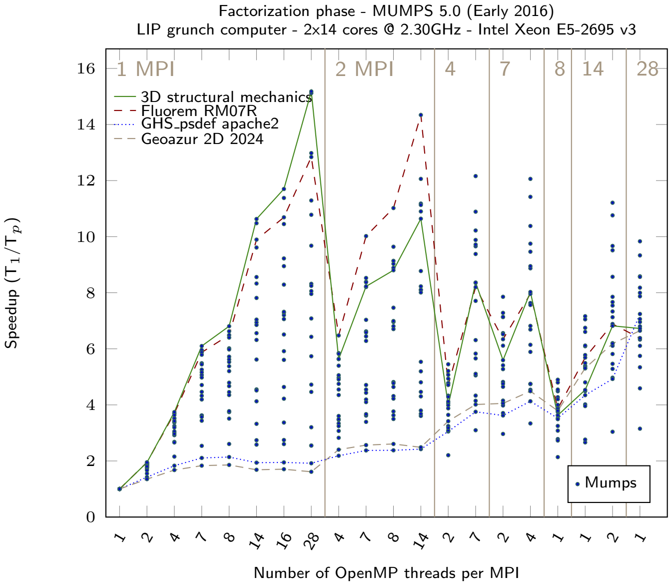

Speed up your factorization with MUMPS 5.1.1

Generic points

- 2D problems require tree-based parallelism: use MPI if possible
- for 3D problems in MT, the number of MPI can be reduced

Getting more parallelism: the tree-based multithreading

MUMPS 5.1.1: Node parallelism (MPI, OpenMP), tree parallelism (MPI only)

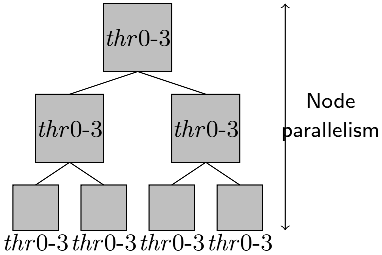

Getting more parallelism: the tree-based multithreading

in future MUMPS (FR++): tree parallelism (MPI, OpenMP)

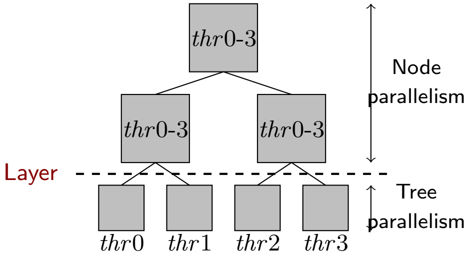

FR ++: Tree-Based Multithreading

- work based on W. M. Sid-Lakhdar's PhD thesis (defended in 2014)
- under the Layer, sub-trees are distributed to threads
- above the Layer, threads work together on each node

Importance of the tree-based multithreading

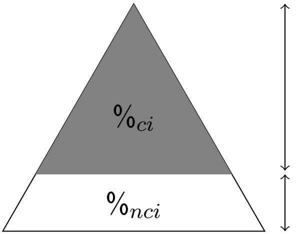

	1 thread			
	time	$\%_{nci}$		
FR	62660s (1)	1%		

3D Poisson; $n = 256^3$ (16M); $\varepsilon = 10^{-6}$; *PhD W. Sid Lakhdar (2014)

Importance of the tree-based multithreading

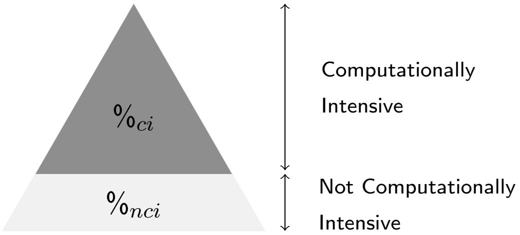

	1 thread			
	time	$\%_{nci}$		
FR	62660s (1)	1%		
BLR	7823s (8)	11%		

3D Poisson; $n = 256^3$ (16M); $\varepsilon = 10^{-6}$; *PhD W. Sid Lakhdar (2014)

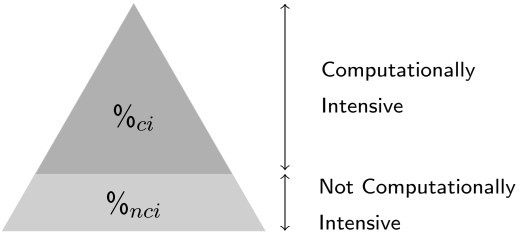

	1 thread			
	time	$\%_{nci}$		
FR	62660s (1)	1%		
BLR	7823s (8)	11%		
BLR++	2464s (25)	38%		

3D Poisson; $n = 256^3$ (16M); $\varepsilon = 10^{-6}$; *PhD W. Sid Lakhdar (2014)

Importance of the tree-based multithreading

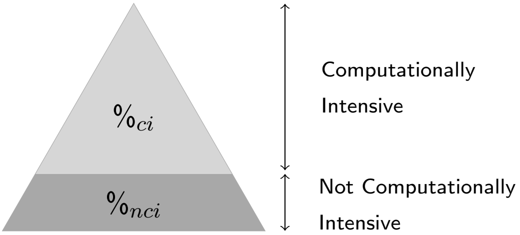

	1 thread		28 threads	
	time	$\%_{nci}$	time	$\%_{nci}$
FR	62660s (1)	1%		
BLR	7823s (8)	11%		
BLR++	2464s (25)	38%	557s (7)	68%

3D Poisson; $n = 256^3$ (16M); $\varepsilon = 10^{-6}$; *PhD W. Sid Lakhdar (2014)

Importance of the tree-based multithreading

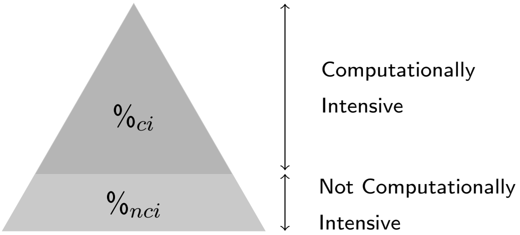

	1 thread		28 threads		28 threads + tree-based MT*	
	time	$\%_{nci}$	time	$\%_{nci}$	time	$\%_{nci}$
FR	62660s (1)	1%				
BLR	7823s (8)	11%				
BLR++	2464s (25)	38%	557s (7)	68%	310s (11)	42%

3D Poisson; $n = 256^3$ (16M); $\varepsilon = 10^{-6}$; *PhD W. Sid Lakhdar (2014)

Importance of the tree-based multithreading

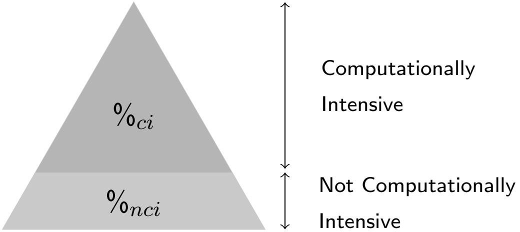

	1 thread		28 threads		28 threads + tree-based MT*	
	time	$\%_{nci}$	time	$\%_{nci}$	time	$\%_{nci}$
FR	62660s (1)	1%	3805s (1)	9%	3430s (1)	0%
BLR	7823s (8)	11%	1356s (3)	26%	1160s (3)	14%
BLR++	2464s (25)	38%	557s (7)	68%	310s (11)	42%

3D Poisson; $n = 256^3$ (16M); $\varepsilon = 10^{-6}$; *PhD W. Sid Lakhdar (2014)

Tree-based multithreading: application

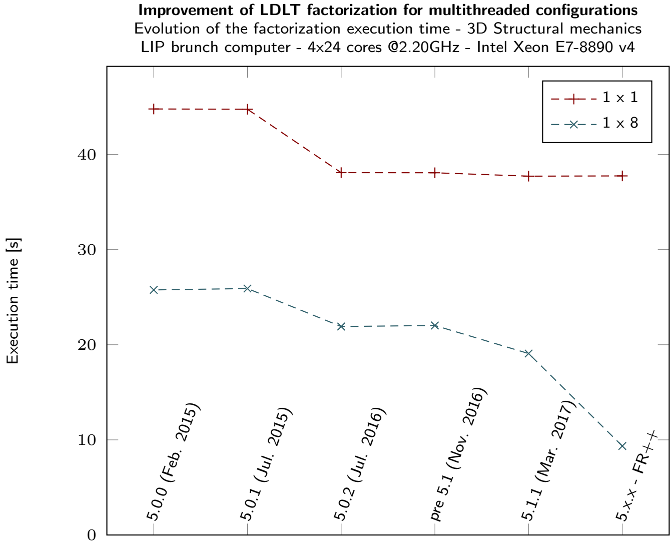

About the solve - 3D

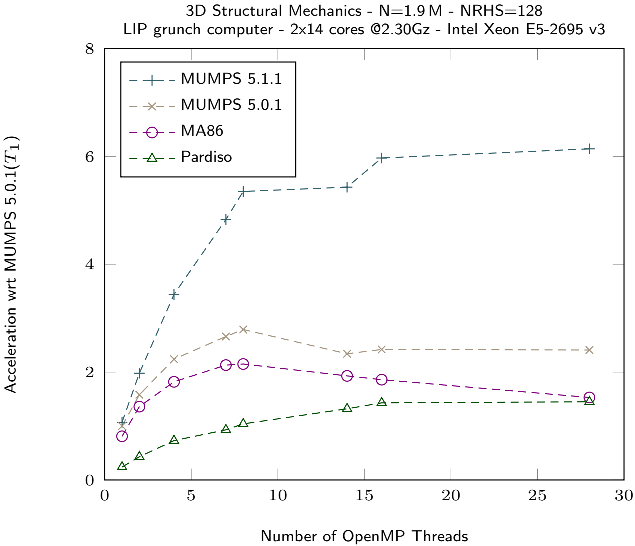

About the solve - 2D

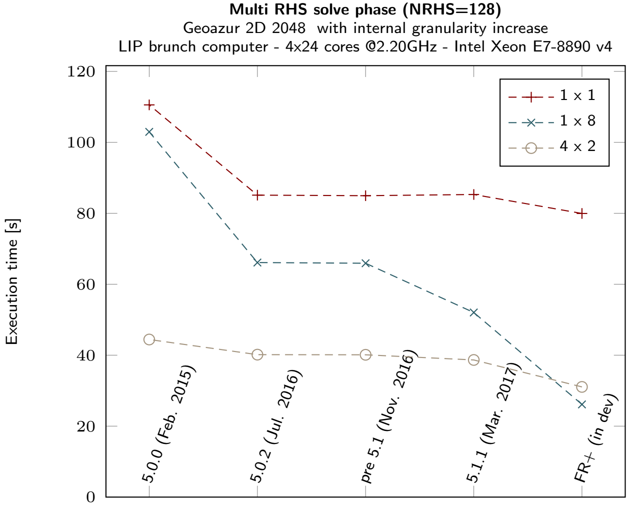

Speed up solve with MUMPS

2D problems

- switch to MUMPS 5.1.1
- improvement increasing internal granularity
- if multiple RHS, increasing the blocking factor may help

Conclusion

We just spoke about

- MUMPS with respect to other MT solvers $\Rightarrow$ not so bad
- what could speed up a lot the factorization and the solve part
  - $\Rightarrow$ tree-based multithreading

We haven't spoken about

- the analyzis part!
- other non computationally intensive parts
  - matrix distribution and scaling (up to 30% of the factor time on some classes of matrices)
  - memory management
  - frontal matrix assembling

MUMPS perspectives and discussions MUMPS group

CERFACS, CNRS, ENS-Lyon, INRIA, INPT, University of Bordeaux

Present and discuss ongoing work that might influence future versions and give new possibilities/perspective to users

Outline

- Preamble: recent work since MUMPS 5.1.1
- Perspectives on Block Low-Rank (BLR)
- BLR memory issues and BLR solve
- Ongoing work

Preamble : recent work since MUMPS 5.1.1

Nodes of parallel computers often have multi/many cores. Good usage of such computers may mix MPI based parallelism with shared memory programming paradigms

- To enhance performance we are working on:
  - Strategies to map MPI tasks on processors
  - Dynamic scheduling, multithreading
  - Multi-level blocking for performance and communication
  - Processing the elimination tree for performance
  - Increase of BLAS3 usage in case pivoting is not requested
- Aggressive optimization setting has been designed
  $\rightarrow$ referred to as **MUMPS FR +**

Timings in seconds		
2x14 cores @ 2.30GHz - Intel Xeon E5-2695 v3		
	Full Rank single	
Electromagnetism, M1ms1 N=0.5 M		
MPI×threads		FR +
1 × 1	**1 080**	1 080
1 × 28	**109**	108
28 × 1	**125**	**88**
RES_∞	1×10^{-5}	

Advanced FR + → gains when using MPI

What about multithreading and BLR?

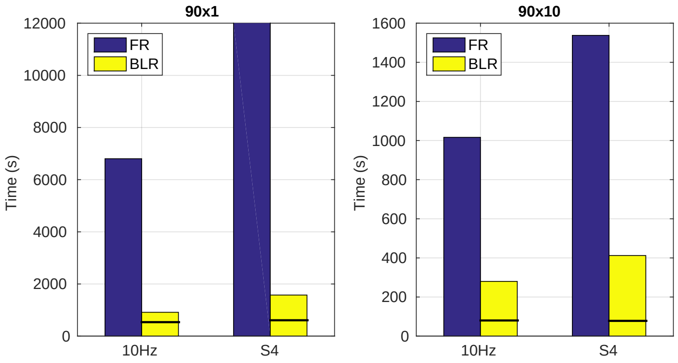

- gain in flops (black line) does not fully translate into gain in time
- average multithreaded efficiency lower in LR than in FR

$\Rightarrow$ *improve efficiency of operations and multithreading with variants*

More Performance?

FR ++: Tree-Based Multithreading (see M. Durand's talk)

- work based on W. M. Sid-Lakhdar's PhD thesis (defended in 2014)

MUMPS is compared here to MKL Pardiso referred to as **Pard**

Timings in seconds - unsymmetric single precision complex Using 1 or 28 OpenMP Threads – No BLR									
2x14 cores @ 2.30GHz - Intel Xeon E5-2695 v3									
	Pard	MUMPS FR +	MUMPS FR ++	Pard	MUMPS FR +	MUMPS FR ++	Pard	MUMPS FR +	MUMPS FR ++
	Matrix M1ms1			Matrix M1ms2			Matrix M3ms1		
1	369	284		1 590	1 080		398	358	
28	21	67	30	110	108	88	29	47	30

- on this class of matrices: performance limited by sequential parts (scaling, matrix distribution)

Multithreading non computationally intensive parts

	FR+ 1 core	FR+ 28 cores	FR++ 28 cores
Factorization	274	**57**	**20**
Scaling+preparation	10	10	10
Total JOB=2	284	67	30

This will become even critical in the future

Even more critical in BLR

S3 (3D EM modelling), times in seconds, 28 cores

FR	BLR	BLR (with FR ++)
585	324	239

More generally, multithread all non computationally intensive parts will become even more critical in the future

Perspectives on Block Low-Rank feature

Expose part of recent work presented in T. Mary's talk

- Amestoy, Buttari, L'Excellent, and Mary. *On the Complexity of the Block Low-Rank Multifrontal Factorization*, SIAM J. Sci. Comput., 2017.
- Amestoy, Buttari, L'Excellent, and Mary. *Block Low-Rank Multifrontal Factorization on Multicore Architectures*, Submitted to ACM Transactions on Mathematical Software.

- Left-looking factorization:
  lower volume of memory transfers in BLR
- Low-rank Updates with Accumulation (so called LUA) :
  increases the GFlops/s rate of the low-rank based update operations.
- Compression is performed before the solve steps:
  additional reduction in the number of operations

→ referred to as BLR + (it includes FR ++)

Results on complete set of problems on 24 threads

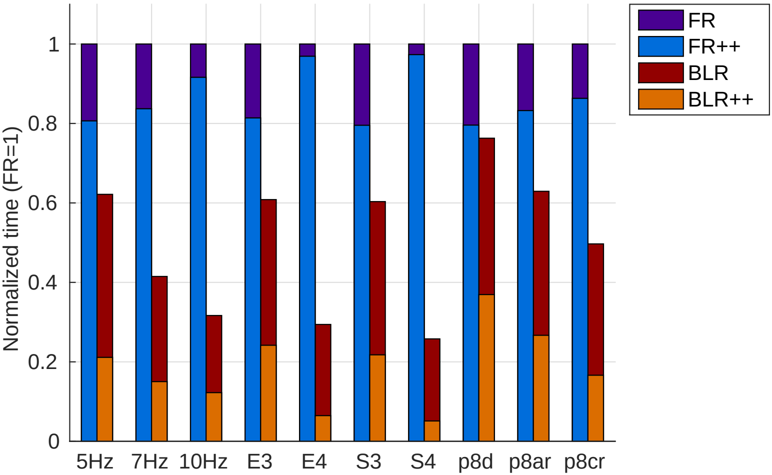

BLR memory issues and BLR solve

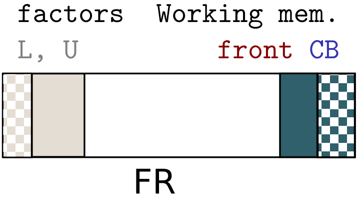

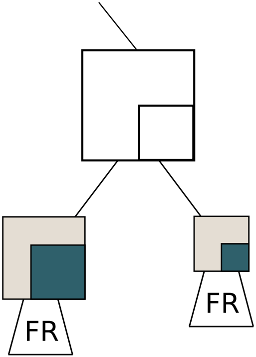

State: FR children have produced a CB (FR subtrees finished)

BLR memory issues and BLR solve

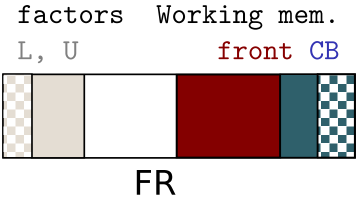

State: Memory for parent reserved

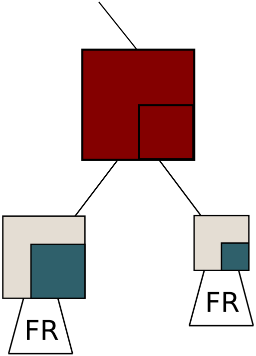

BLR memory issues and BLR solve

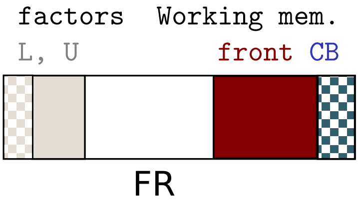

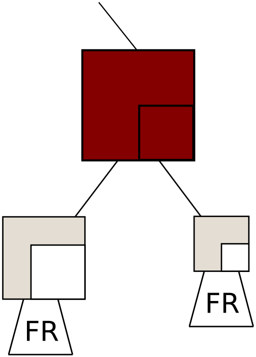

State: Parent assembled (children CB consumed)

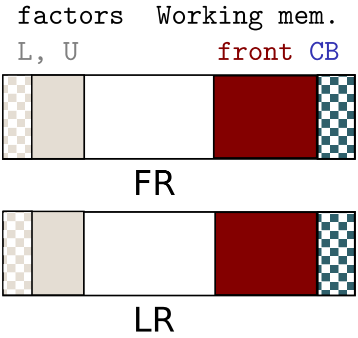

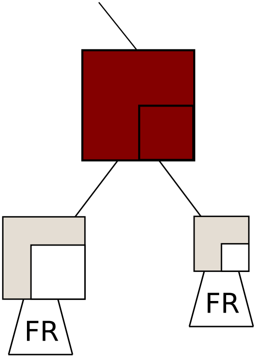

State: Parent assembled (children CB consumed)

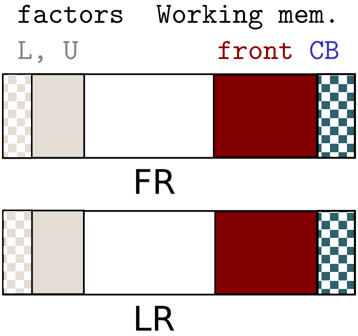

State: FR factorization of parent

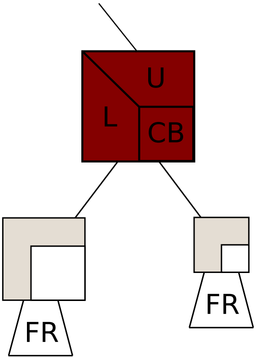

BLR memory issues and BLR solve

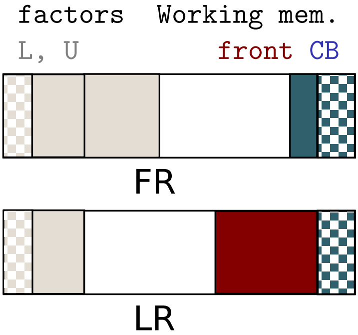

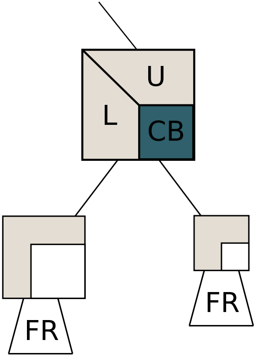

State: Stacked FR factors and FR CB

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Factor size/proc	
	FR	BLR
1 proc	711 GB	
90 procs	8 GB	

Memory efficiency: perfect for factors

BLR memory issues and BLR solve

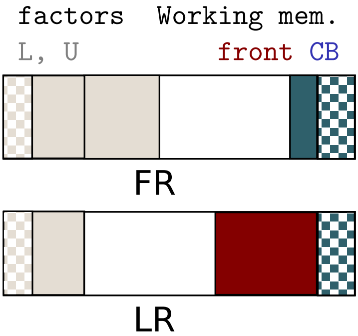

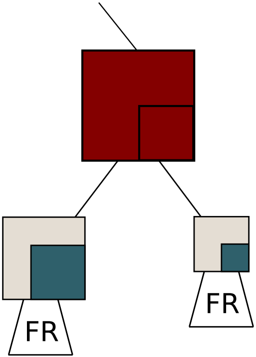

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Factor size/proc	
	FR	BLR
1 proc	711 GB	
90 procs	8 GB	

Memory efficiency: perfect for factors

BLR memory issues and BLR solve

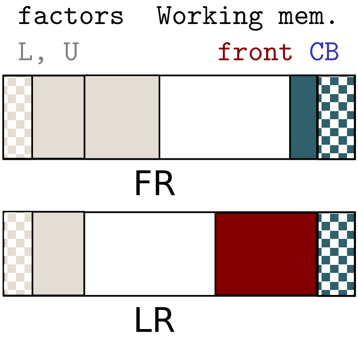

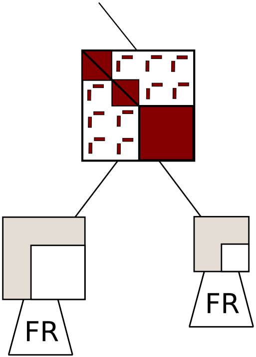

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Factor size/proc	
	FR	BLR
1 proc	711 GB	
90 procs	8 GB	

Memory efficiency: perfect for factors

BLR memory issues and BLR solve

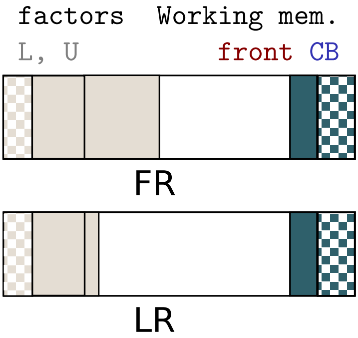

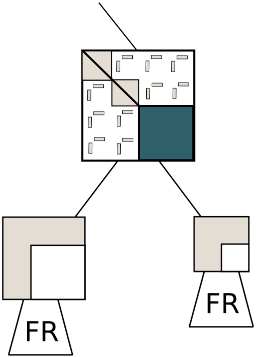

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Factor size/proc	
	FR	BLR
1 proc	711 GB	175 GB
90 procs	8 GB	2 GB

Memory efficiency: perfect for factors

BLR memory issues and BLR solve

Reminding current BLR characteristics (MUMPS 5.1):

- Memory not reduced
- Compression used to accelerate factorization
- Factors in full-rank (approximated) form stored in memory for current FR solve to work

BLR solve

- Reduce memory usage:
  - keep BLR factors during factorization (for use during solve)
  - Memory used: Working memory (tend to OOC size with compression) + FR factors of small fronts + BLR factors
- Reduce flops/memory accesses during solve

BLR-related memory issues (factorization)

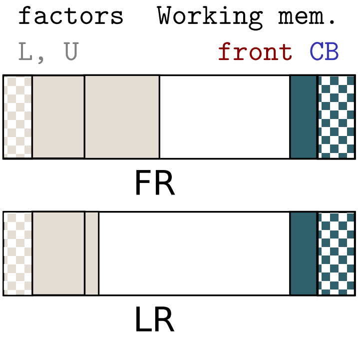

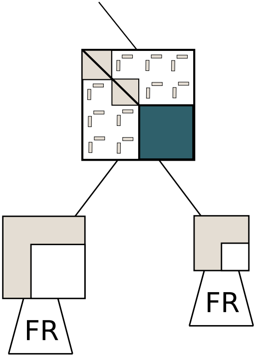

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Factors/proc		Working mem./proc	
	FR	BLR	FR	BLR
1 proc	711 GB	175 GB	87 GB	
90 procs	8 GB	2 GB	7 GB	

Memory efficiency: perfect for factors; bad for working memory (12/90)

BLR-related memory issues (factorization)

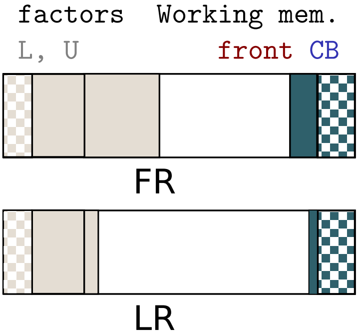

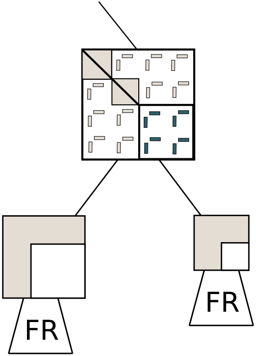

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Factors/proc		Working mem./proc	
	FR	BLR	FR	BLR
1 proc	711 GB	175 GB	87 GB	9 GB
90 procs	8 GB	2 GB	7 GB	1 GB

Memory efficiency: perfect for factors; bad for working memory (12/90)
Good news: working memory has greater potential for compression
BUT at the cost of an increase in compression flops

BLR-related memory issues (factorization)

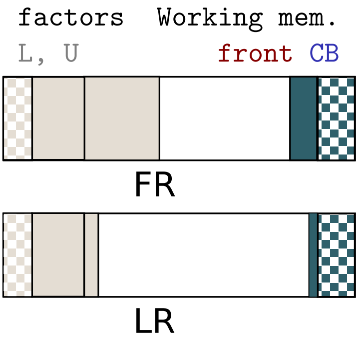

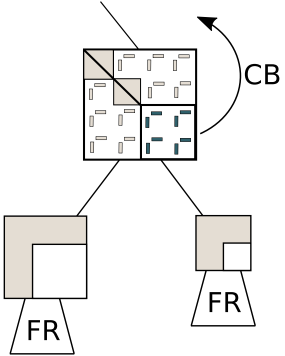

Matrix from SEISCOPE 10Hz (3D seismic imaging, N: 17M, NNZ: 446M)

	Total	Within fronts	Between fronts
FR	5.1 TB	3.0 TB	2.0 TB
BLR (5.1.1)	2.3 TB	0.3 TB	2.0 TB
BLR (if compressed CB)	0.5 TB	0.3 TB	0.2 TB

Volume of communication (90 MPI procs)

BLR-related memory issues (factorization)

- Factors and CB blocks can be compressed BUT compression factor not known in advance !
- Provide memory estimates at analysis?
- Compression of CBs
  - offers a great potential but at the cost of extra compression
  - could also reduce communication volume
  - To be understood: how to compromise flops increase and memory reduction objectives

Much active research and perspectives

- Performance related
  - Performance of the solution phase
  - Improve quality of memory estimates
  - Improve parallel efficiency
- BLR-related
  - New BLR variants, improved compression and further reduce complexity (collaboration with LSTC)
  - Comparisons with HSS (LBNL collab.)
  - MPI/OpenMP performance and scalability
  - Exploit BLR format during solution phase and compress memory during factorization
- Performance of solution phase
  - Sometimes critical (electromagnetism, geophysics, DD, ...)
  - MPI/OpenMP scalability (e.g. mapping of factors for solve?)
  - Sparsity of right-hand sides: flops reduction versus parallelism?

Work often related to the PhD thesis of T. Mary and G. Moreau

Closing Session

Concluding remarks

Statistics about workshop

- 61 participants (Belgium, France, Greece, Israel, Italy, Norway, Saudi Arabia, South Africa, Spain, USA), 45 have participated to the banquet
  - 39 industrials
  - 22 academics
- 18 talks (MUMPS overview and recent features, 2 talks from MUMPS PhD students, 9 talks from industrials, 3 from public researchers using MUMPS, 2 talks from MUMPS team, Perspectives and discussion)

Next MUMPS usersday

- Change format (talks, duration) ?
- When and Where ?

Merci à Inria pour son accueil

INPT-ENSEEIHT IRIT Inria MUMPS-Consortium

Labex Milyon

Merci aux organisateurs:
Marie Durand, Guillaume Joslin, Chiara Puglisi