On the Complexity of the Block Low-Rank Multifrontal Factorization

P. Amestoy[*,1] A. Buttari[*,2] J.-Y. L'Excellent[†,3] T. Mary[*,4]

[*]Université de Toulouse [†]ENS Lyon
[1]INPT-IRIT [2]CNRS-IRIT [3]INRIA-LIP [4]UPS-IRIT

Sparse Days 2016, Toulouse Jun. 30 - Jul. 1

Introduction

Multifrontal (Duff & Reid '83) with Nested Dissection (George '73)

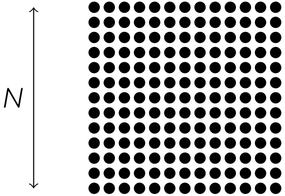

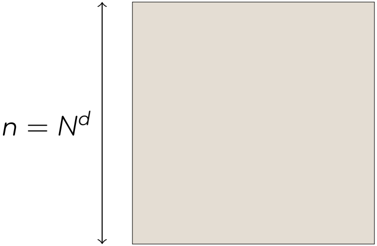

Multifrontal (Duff & Reid '83) with Nested Dissection (George '73)

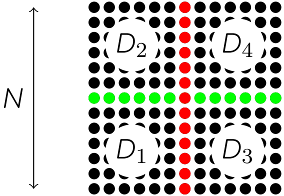

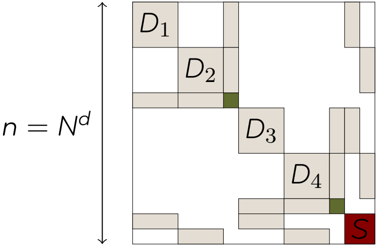

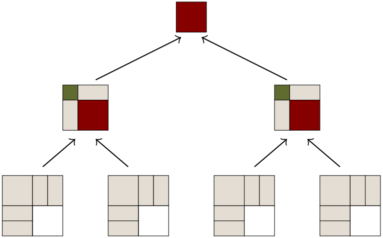

3D problem cost $\propto$

$\rightarrow$ Flops:$O(n^2)$, mem:$O(n^{4/3})$

$\mathcal{H}$ and BLR matrices

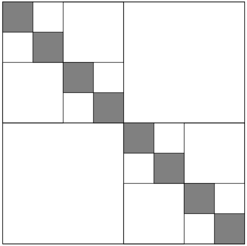

$\mathcal{H}$-matrix

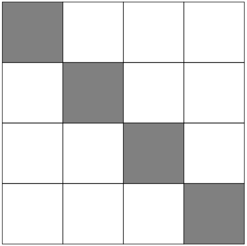

BLR matrix

$\mathcal{H}$ and BLR matrices

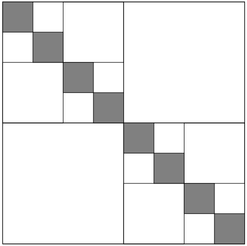

$\mathcal{H}$-matrix

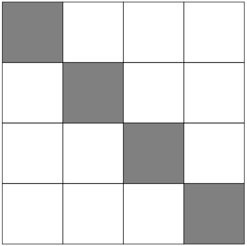

BLR matrix

A block B represents the interaction between two subdomains σ and τ. If they have a small diameter and are far away their interaction is weak $\Rightarrow$ rank is low.

$\mathcal{H}$ and BLR matrices

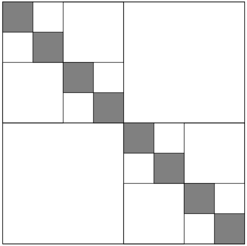

$\mathcal{H}$-matrix

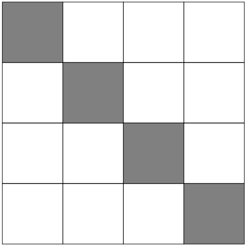

BLR matrix

A block B represents the interaction between two subdomains σ and τ. If they have a small diameter and are far away their interaction is weak $\Rightarrow$ rank is low.

Block-admissibility condition

$$\sigma \times \tau \text{ is admissible } \Leftrightarrow \max(\text{diam}(\sigma), \text{diam}(\tau)) \leq \eta \, \text{dist}(\sigma, \tau)$$

$$\eta = \eta_{max} \Rightarrow \text{ admissibility condition becomes } \text{dist}(\sigma, \tau) > 0$$

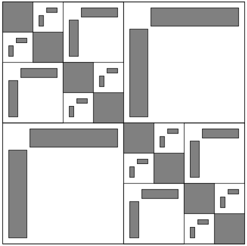

$\mathcal{H}$-matrix

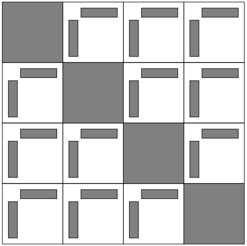

BLR matrix

$$\tilde{B} = XY^T \text{ such that } \mathrm{rank}(\tilde{B}) = k_\varepsilon \text{ and } \|B - \tilde{B}\| \leq \varepsilon$$

If $k_\varepsilon \ll \mathrm{size}(B) \Rightarrow$ memory and flops can be reduced with a controlled loss of accuracy ($\leq \varepsilon$)

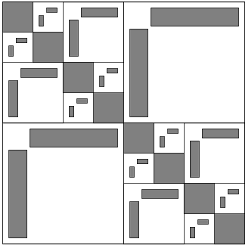

$\mathcal{H}$-matrix

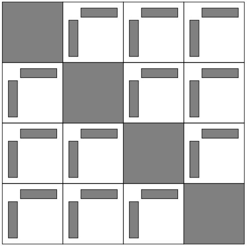

BLR matrix

- Very low theoretical complexity
- Complex, hierarchical structure

- Simple structure
- Theoretical complexity?

$\mathcal{H}$ and BLR matrices

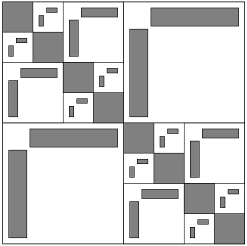

$\mathcal{H}$-matrix

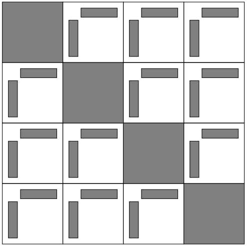

BLR matrix

- Very low theoretical complexity
- Complex, hierarchical structure

- Simple structure
- Theoretical complexity?

$\Rightarrow$ Our hope is to find a good comprise between theoretical complexity and performance/usability

Questions that will be answered in this talk

- What theoretical bound on the ranks of the blocks can we derive? How does it compare to the $\mathcal{H}$ case?
- What is the complexity of the BLR factorization? In particular, is it asymptotically better than the full-rank one? (i.e., in $O(n^\alpha)$, with $\alpha < 2$ and where n is the number of unknowns)
- What are the different variants of the BLR factorization? Do they improve its complexity?
- Can we validate these theoretical results with experimental ones? In particular, does the theory hold in a purely algebraic context?
- How does the low-rank threshold ε influence the complexity? How about the block size b?

Variants of the BLR factorization

Variants of the BLR LU factorization

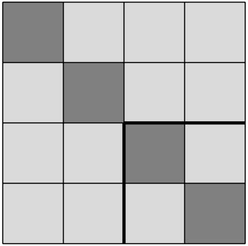

- FSCU

Variants of the BLR LU factorization

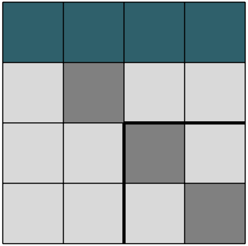

- FSCU (Factor,

Variants of the BLR LU factorization

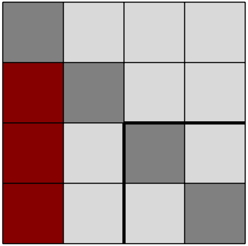

- FSCU (Factor, Solve,

Variants of the BLR LU factorization

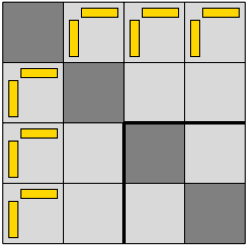

- FSCU (Factor, Solve, Compress,

Variants of the BLR LU factorization

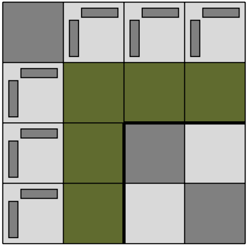

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

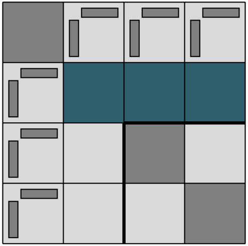

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

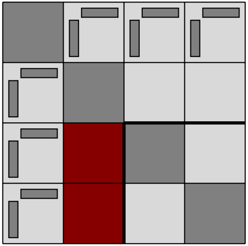

- FSCU (Factor, Solve, Compress, Update)

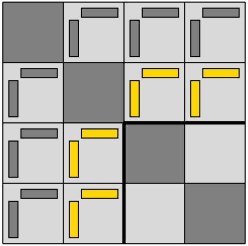

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

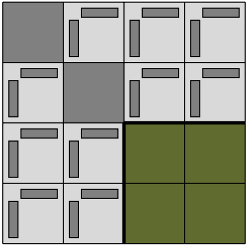

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

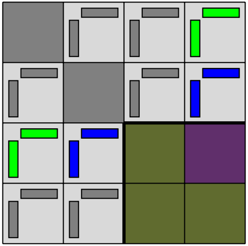

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

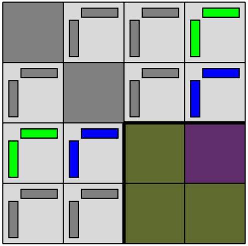

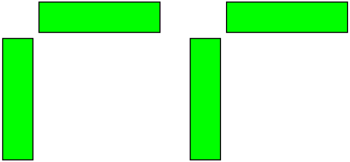

- FSCU (Factor, Solve, Compress, Update)

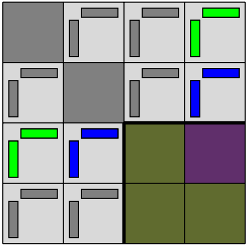

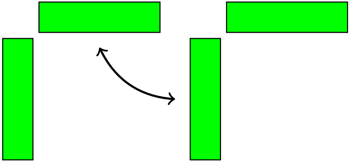

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

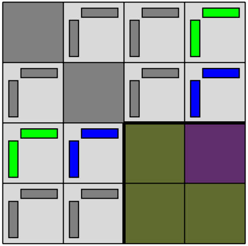

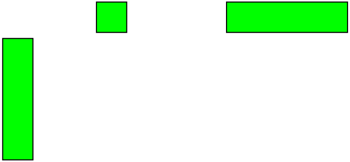

- FSCU (Factor, Solve, Compress, Update)

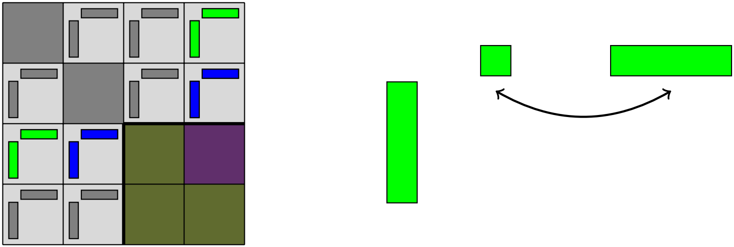

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

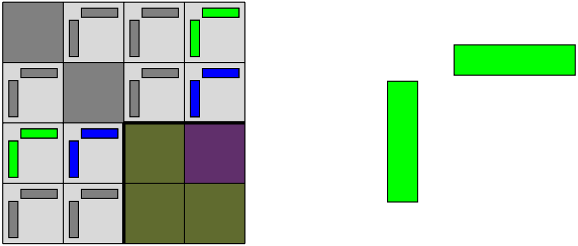

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

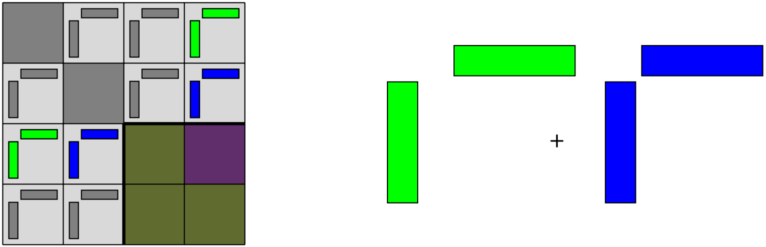

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

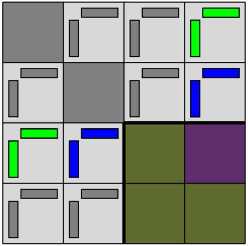

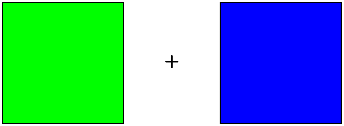

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

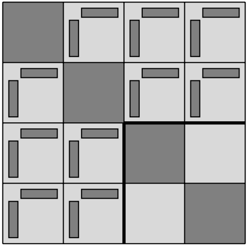

- FSCU (Factor, Solve, Compress, Update)

Variants of the BLR LU factorization

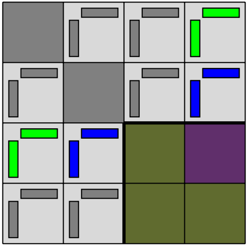

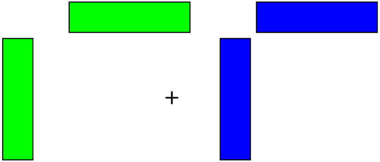

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR

Variants of the BLR LU factorization

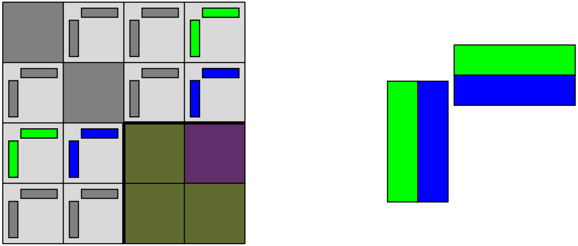

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR

Variants of the BLR LU factorization

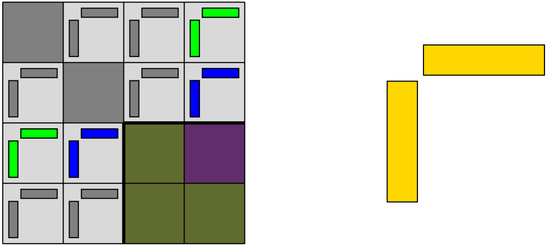

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression ⇒ better complexity?

Variants of the BLR LU factorization

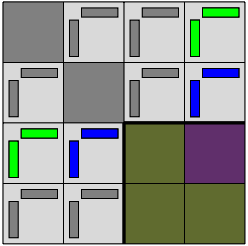

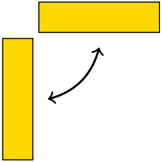

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression ⇒ better complexity?

Variants of the BLR LU factorization

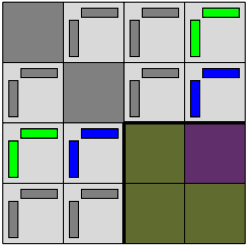

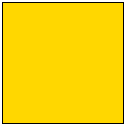

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression ⇒ better complexity?

Variants of the BLR LU factorization

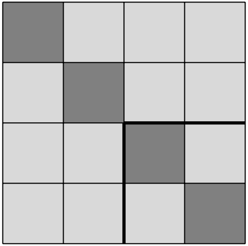

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression ⇒ better complexity?
- FCSU(+LUAR)

Variants of the BLR LU factorization

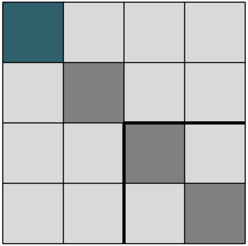

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression $\Rightarrow$ better complexity?
- FCSU(+LUAR)
  - Restricted pivoting, e.g. to diagonal blocks

Variants of the BLR LU factorization

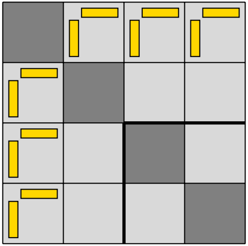

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression $\Rightarrow$ better complexity?
- FCSU(+LUAR)
  - Restricted pivoting, e.g. to diagonal blocks

Variants of the BLR LU factorization

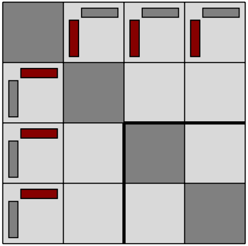

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression ⇒ better complexity?
- FCSU(+LUAR)
  - Restricted pivoting, e.g. to diagonal blocks
  - Low-rank Solve ⇒ better complexity?

Variants of the BLR LU factorization

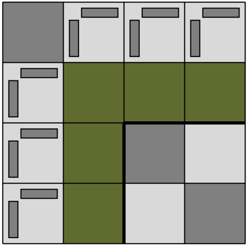

- FSCU (Factor, Solve, Compress, Update)
- FSCU+LUAR
  - Potential for recompression ⇒ better complexity?
- FCSU(+LUAR)
  - Restricted pivoting, e.g. to diagonal blocks
  - Low-rank Solve ⇒ better complexity?

Theoretical complexity of the BLR factorization

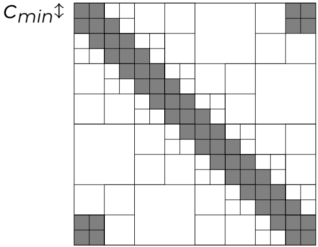

$\mathcal{H}$-admissibility condition

A partition $P \in \mathcal{P}(\mathcal{I} \times \mathcal{I})$ is admissible iff

$$\forall \sigma \times \tau \in P, \ \ \sigma \times \tau \text{ is admissible } \ \text{or} \ \ \min(\#\sigma, \#\tau) \le c_{min} \ \ (Adm_{\mathcal{H}})$$

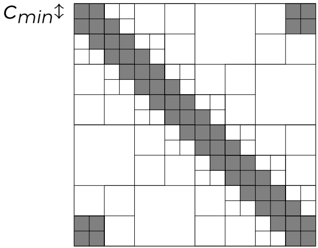

c_{sp} is the max number of blocks of the same size on the same row/column (here, $c_{sp} = 6$)

$\mathcal{H}$-admissibility condition

A partition $P \in \mathcal{P}(\mathcal{I} \times \mathcal{I})$ is admissible iff

$$\forall \sigma \times \tau \in P, \ \ \sigma \times \tau \text{ is admissible} \ \ \text{or} \ \ \min(\#\sigma, \#\tau) \leq c_{min} \ \ (Adm_{\mathcal{H}})$$

The so-called sparsity constant c_{sp} is defined by:

$$c_{sp} = \max(\max_{\sigma \subset \mathcal{I}} \#\{\tau; \sigma \times \tau \in P\}, \max_{\tau \subset \mathcal{I}} \#\{\sigma; \sigma \times \tau \in P\})$$

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(m c_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}		
$r_{\mathcal{H}}$		
$\mathcal{C}_{facto}$		

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	
$r_{\mathcal{H}}$		
$\mathcal{C}_{facto}$		

**Grasedyck & Hackbusch, 2003*

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	
$r_{\mathcal{H}}$	small**	
$\mathcal{C}_{facto}$		

**Grasedyck & Hackbusch, 2003*
***Bebendorf & Hackbusch, 2003*

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	
$r_{\mathcal{H}}$	small**	
$\mathcal{C}_{facto}$	$O(r_{\mathcal{H}}^2 m)$	

**Grasedyck & Hackbusch, 2003*

***Bebendorf & Hackbusch, 2003*

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	m/b
$r_{\mathcal{H}}$	small**	
$\mathcal{C}_{facto}$	$O(r_{\mathcal{H}}^2 m)$	

*Grasedyck & Hackbusch, 2003

**Bebendorf & Hackbusch, 2003

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	m/b
$r_{\mathcal{H}}$	small**	b
$\mathcal{C}_{facto}$	$O(r_{\mathcal{H}}^2 m)$	

**Grasedyck & Hackbusch, 2003*

***Bebendorf & Hackbusch, 2003*

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	m/b
$r_{\mathcal{H}}$	small**	b
$\mathcal{C}_{facto}$	$O(r_{\mathcal{H}}^2 m)$	$O(m^3)$

Grasedyck & Hackbusch, 2003

***Bebendorf & Hackbusch, 2003*

$\mathcal{H}$ vs. BLR complexity

Dense factorization complexity

Complexity: $\mathcal{C}_{facto} = O(mc_{sp}^2 r_{\mathcal{H}}^2)$ (best case)

m	matrix size
c_{sp}	sparsity constant
$r_{\mathcal{H}}$	bound on the maxrank of all blocks

	$\mathcal{H}$	BLR
c_{sp}	$O(1)^*$	m/b
$r_{\mathcal{H}}$	small**	b
$\mathcal{C}_{facto}$	$O(r_{\mathcal{H}}^2 m)$	$O(m^3)$

**Grasedyck & Hackbusch, 2003*
***Bebendorf & Hackbusch, 2003*

BLR: a particular case of $\mathcal{H}$?

Problem: in $\mathcal{H}$ formalism, the maxrank of the blocks of a BLR matrix is $r_{\mathcal{H}} = b$ (due to the non-admissible blocks)
Solution: bound the rank of the admissible blocks only, and make sure the non-admissible blocks are in small number

Complexity of dense BLR factorization

BLR-admissibility condition of a partition $\mathcal{P}$

$$\mathcal{P} \text{ is admissible } \Leftrightarrow N_{na} = \#\{\sigma \times \tau \in \mathcal{P}, \ \ \sigma \times \tau \text{ is not admissible}\} \leq q$$

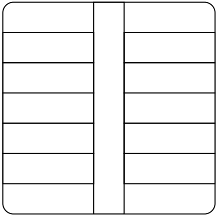

Non-Admissible

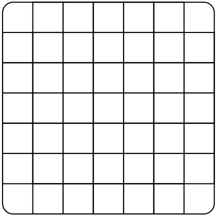

Admissible

Complexity of dense BLR factorization

BLR-admissibility condition of a partition $\mathcal{P}$

$\mathcal{P}$ is admissible $\Leftrightarrow N_{na} = \#\{\sigma \times \tau \in \mathcal{P}, \;\; \sigma \times \tau \text{ is not admissible}\} \leq q$

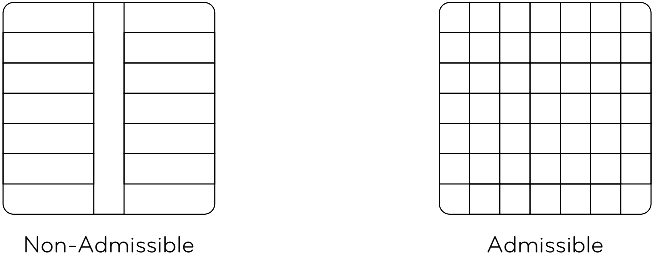

Main result

There exists an admissible $\mathcal{P}$ for $q = O(1)$, s.t. the maxrank of the admissible blocks of A is $r = O(r_{\mathcal{H}})$ *(Amestoy, Buttari, L'Excellent & Mary, 2016)*.

The best case dense factorization complexity thus becomes

$\mathcal{C}_{facto} = O(r^2m^3/b^2 + mb^2q^2) = O(r^2m^3/b^2 + mb^2) = O(rm^2)$ (for $b = O(\sqrt{rm})$)

Element of proof 1: boundedness of N_{na}

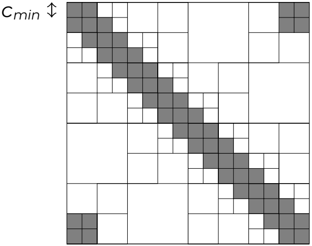

$c_{sp} = 6 = O(1)$
$N_{na} = 4 \leq c_{sp}$

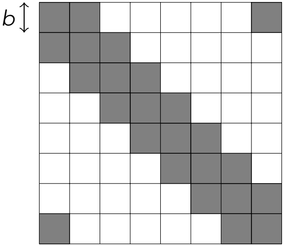

$c_{sp} = m/b \neq O(1)$
$N_{na} = 3 = O(1)$

Secondary result *(Amestoy et al., 2016)*

(a) $N_{na}^{(BLR)} \leq N_{na}^{(\mathcal{H})}$

(b) $N_{na} \leq c_{sp}$

(c) $c_{sp}^{(\mathcal{H})} = O(1)$ *(Grasedyck & Hackbusch, 2003)*

Corollary: $N_{na}^{(BLR)} = O(1)$

Element of proof 2: dense BLR factorization complexity

The computations can be divided in two parts:

- FR part: Factor, Solve (if FSCU), and Update for non-admissible blocks
- LR part: Compress, Solve (if FCSU), and Update for admissible blocks

The relative weight of these two parts changes with the variant ⇒ choose for each variant the optimal block size b^* that minimizes the total

variant	FR part	LR part	b^*	$\mathcal{C}_{facto}$
FSCU	$O(m^2 b)$	$O(rm^3/b)$	$\sqrt{rm}$	$O(\sqrt{r}m^{2.5})$
FSCU+LUAR	$O(m^2 b)$	$O(r^2 m^3/b^2)$	$\sqrt[3]{r^2 m}$	$O(\sqrt[3]{r^2}m^{2.33})$
FCSU+LUAR	$O(mb^2)$	$O(r^2 m^3/b^2)$	$\sqrt{rm}$	$O(rm^2)$

Complexity of multifrontal BLR factorization

Under a nested dissection assumption, the sparse (multifrontal) complexity is directly obtained from the dense complexity

	operations (OPC)		factor size (NNZ)	
	$r = O(1)$	$r = O(N)$	$r = O(1)$	$r = O(N)$
FR	$O(n^2)$	$O(n^2)$	$O(n^{\frac{4}{3}})$	$O(n^{\frac{4}{3}})$
BLR FSCU	$O(n^{\frac{5}{3}})$	$O(n^{\frac{11}{6}})$	$O(n \log n)$	$O(n^{\frac{7}{6}} \log n)$
BLR FSCU+LUAR	$O(n^{\frac{14}{9}})$	$O(n^{\frac{16}{9}})$	$O(n \log n)$	$O(n^{\frac{7}{6}} \log n)$
BLR FCSU+LUAR	$O(n^{\frac{4}{3}})$	$O(n^{\frac{5}{3}})$	$O(n \log n)$	$O(n^{\frac{7}{6}} \log n)$
$\mathcal{H}$	$O(n^{\frac{4}{3}})$	$O(n^{\frac{5}{3}})$	$O(n)$	$O(n^{\frac{7}{6}})$
$\mathcal{H}$ (fully struct.)	$O(n)$	$O(n^{\frac{4}{3}})$	$O(n)$	$O(n^{\frac{7}{6}})$

in the 3D case (similar analysis possible for 2D)

Important properties: with both $r = O(1)$ or $r = O(N)$

- The complexity of the standard BLR variant (FSCU) has a lower exponent than the full-rank one
- Each variant further improves the complexity, with the best one (FCSU+LUAR) being not so far from the $\mathcal{H}$-case

Experimental complexity of the BLR factorization

Experimental Setting: Matrices

1. **`Poisson`**: N^3 grid with a 7-point stencil with $u = 1$ on the boundary $\partial\Omega$

$$\Delta u = f$$

2. **`Helmholtz`**: N^3 grid with a 27-point stencil, ω is the angular frequency, $v(x)$ is the seismic velocity field, and $u(x, \omega)$ is the time-harmonic wavefield solution to the forcing term $s(x, \omega)$.

$$\left(-\Delta - \frac{\omega^2}{v(x)^2}\right) u(x, \omega) = s(x, \omega)$$

ω is fixed and equal to 4Hz.

Nested Dissection ordering (geometric)

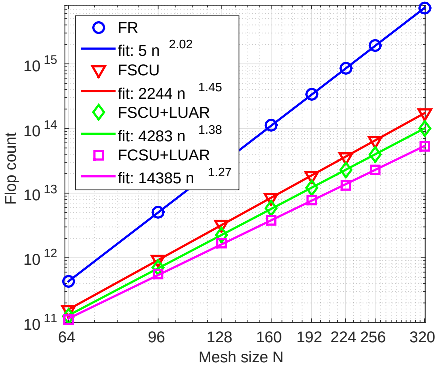

- good agreement with theoretical complexity ($O(n^2)$, $O(n^{1.67})$, $O(n^{1.55})$, and $O(n^{1.33})$)

Experimental MF flop complexity: Poisson ($\varepsilon = 10^{-10}$)

Nested Dissection ordering (geometric)

METIS ordering (purely algebraic)

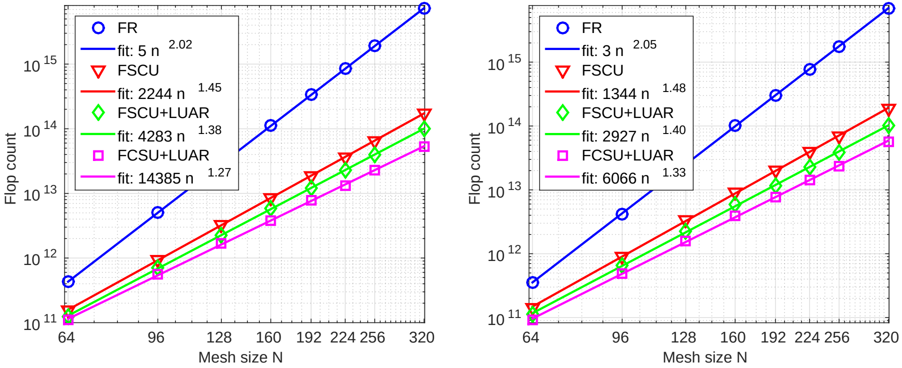

- good agreement with theoretical complexity ($O(n^2)$, $O(n^{1.67})$, $O(n^{1.55})$, and $O(n^{1.33})$)
- remains close to ND complexity with METIS ordering

Experimental MF flop complexity: Helmholtz ($\varepsilon = 10^{-4}$)

Nested Dissection ordering (geometric)

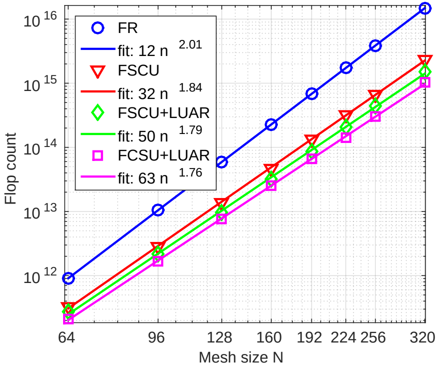

METIS ordering (purely algebraic)

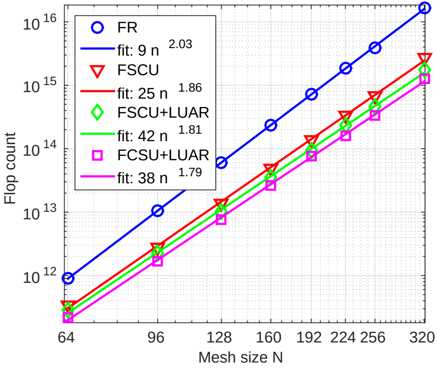

- good agreement with theoretical complexity ($O(n^2)$, $O(n^{1.83})$, $O(n^{1.78})$, and $O(n^{1.67})$)
- remains close to ND complexity with METIS ordering

Experimental MF complexity: factor size

NNZ (Poisson)

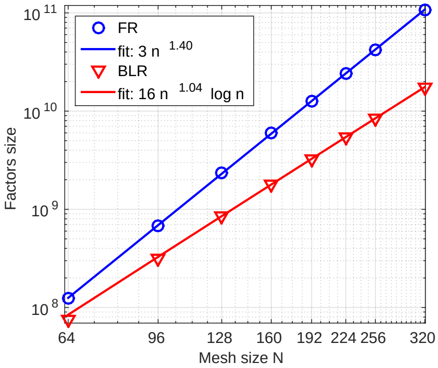

NNZ (Helmholtz)

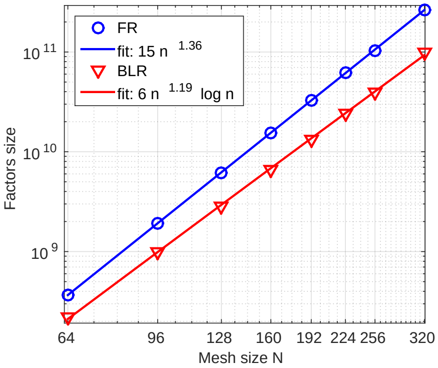

- good agreement with theoretical complexity (FR: $O(n^{1.33})$; BLR: $O(n \log n)$ and $O(n^{1.17} \log n)$)

Experimental MF complexity: low-rank threshold ε

OPC (Poisson)

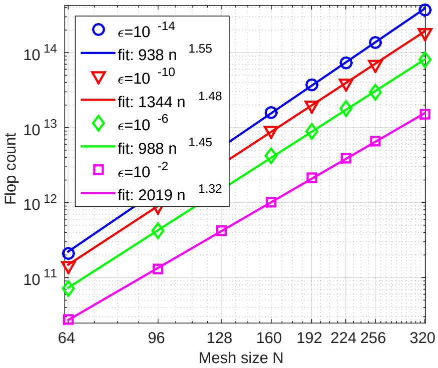

OPC (Helmholtz)

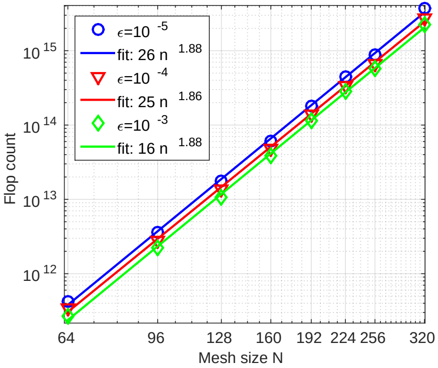

- theory states ε should only play a role in the constant factor
- true for Helmholtz, but not Poisson $\Rightarrow$ why?

Influence of zero-rank blocks on the complexity

		N 64	128	192	256	320
$\varepsilon = 10^{-14}$	N_{FR}	40.8	31.3	26.4	23.6	13.4
	N_{LR}	59.2	68.6	73.6	76.4	86.6
	N_{ZR}	0.0	0.1	0.0	0.0	0.0
$\varepsilon = 10^{-10}$	N_{FR}	21.3	16.6	14.6	12.8	5.8
	N_{LR}	78.6	83.4	85.4	87.1	94.2
	N_{ZR}	0.0	0.1	0.0	0.0	0.0
$\varepsilon = 10^{-6}$	N_{FR}	2.9	3.0	2.5	2.1	0.6
	N_{LR}	97.0	96.7	96.4	95.3	93.3
	N_{ZR}	0.1	0.3	1.0	2.5	6.1
$\varepsilon = 10^{-2}$	N_{FR}	0.0	0.0	0.0	0.0	0.0
	N_{LR}	26.2	12.2	7.6	5.5	3.0
	N_{ZR}	73.8	87.8	92.4	94.5	97.0

Number of full-rank/low-rank/zero-rank blocks
in percentage of the total number of blocks (Poisson problem).

Influence of zero-rank blocks on the complexity

		N				
		64	128	192	256	320
$\varepsilon = 10^{-14}$	N_{FR}	40.8	31.3	26.4	23.6	13.4
	N_{LR}	59.2	68.6	73.6	76.4	86.6
	N_{ZR}	0.0	0.1	0.0	0.0	0.0
$\varepsilon = 10^{-10}$	N_{FR}	21.3	16.6	14.6	12.8	5.8
	N_{LR}	78.6	83.4	85.4	87.1	94.2
	N_{ZR}	0.0	0.1	0.0	0.0	0.0
$\varepsilon = 10^{-6}$	N_{FR}	2.9	3.0	2.5	2.1	0.6
	N_{LR}	97.0	96.7	96.4	95.3	93.3
	N_{ZR}	0.1	0.3	1.0	2.5	6.1
$\varepsilon = 10^{-2}$	N_{FR}	0.0	0.0	0.0	0.0	0.0
	N_{LR}	26.2	12.2	7.6	5.5	3.0
	N_{ZR}	73.8	87.8	92.4	94.5	97.0

Number of full-rank/low-rank/zero-rank blocks
in percentage of the total number of blocks (Poisson problem).

- N_{FR} decreases with N: asymptotically negligible

Influence of zero-rank blocks on the complexity

		N 64	128	192	256	320
$\varepsilon = 10^{-14}$	N_{FR}	40.8	31.3	26.4	23.6	13.4
	N_{LR}	59.2	68.6	73.6	76.4	86.6
	N_{ZR}	0.0	0.1	0.0	0.0	0.0
$\varepsilon = 10^{-10}$	N_{FR}	21.3	16.6	14.6	12.8	5.8
	N_{LR}	78.6	83.4	85.4	87.1	94.2
	N_{ZR}	0.0	0.1	0.0	0.0	0.0
$\varepsilon = 10^{-6}$	N_{FR}	2.9	3.0	2.5	2.1	0.6
	N_{LR}	97.0	96.7	96.4	95.3	93.3
	N_{ZR}	0.1	0.3	1.0	2.5	6.1
$\varepsilon = 10^{-2}$	N_{FR}	0.0	0.0	0.0	0.0	0.0
	N_{LR}	26.2	12.2	7.6	5.5	3.0
	N_{ZR}	73.8	87.8	92.4	94.5	97.0

Number of full-rank/low-rank/zero-rank blocks
in percentage of the total number of blocks (Poisson problem).

- N_{FR} decreases with N: asymptotically negligible
- N_{ZR} increases with ε (as one would expect) but also with N: asymptotically dominant

Analysis on the root node (of size $m = N^2$):

FSCU FSCU+LUAR FCSU+LUAR

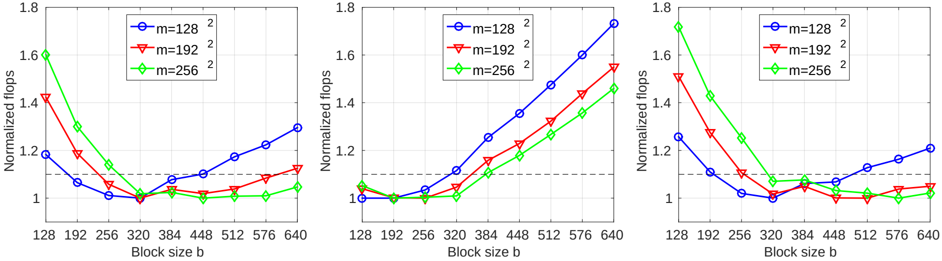

- large range of acceptable block sizes around the optimal b^*
  $\Rightarrow$ flexibility to tune block size for performance
- that range increases with the size of the matrix
  $\Rightarrow$ necessity to have variable block sizes
- necessity to adjust b^* for each new variant

Conclusion and perspectives

Summary

- BLR matrices are a particular kind of $\mathcal{H}$-matrices but $\mathcal{H}$-matrix theory does not provide satisfying results for BLR matrices
- Extended theory to compute complexity bounds of the BLR (multifrontal) factorization
- Theoretical complexity of the BLR (multifrontal) factorization is asymptotically better than FR
- Studied BLR variants to further reduce complexity by achieving higher compression
- Numerical experiments show experimental complexity in agreement with theoretical one
- Identified and analyzed the importance of zero-rank blocks and variable block sizes on the complexity

Perspectives

- Efficient strategies to recompress accumulators
- Pivoting strategies compatible with the BLR variants
- Influence of the BLR variants on the accuracy of the factorization

Acknowledgements

- CALMIP for providing access to the machines
- SEISCOPE for providing the Helmholtz Generator
- LSTC members for scientific discussions

Thanks!

Questions?

Backup Slides

Accumulator recompression

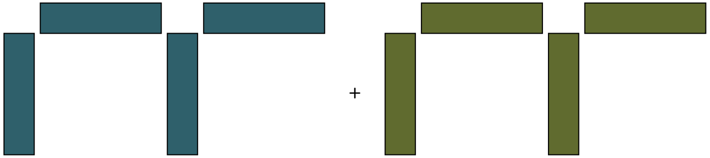

Accumulator recompression

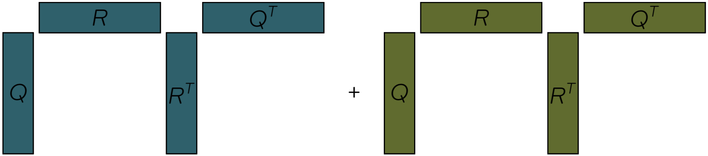

Accumulator recompression

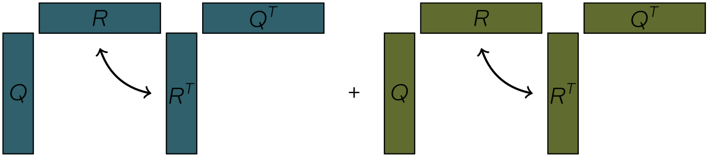

Accumulator recompression

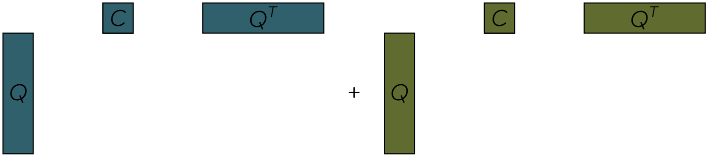

Accumulator recompression

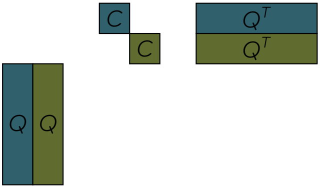

Accumulator recompression

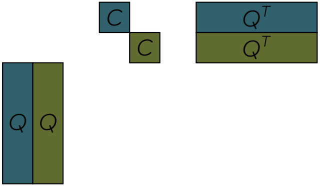

- Weight recompression on $\{C_i\}_i$

  ⇒ With absolute threshold ε, each C_i can be compressed separately

- Redundancy recompression on $\{Q_i\}_i$

  ⇒ Bigger recompression overhead, when is it worth it?

Scope of the talk

Complexity and performance
of the Block Low-Rank multifrontal
factorization and its variants
SIAM PP'16, April 12-15, Paris
(Overview of preliminary work on both complexity and performance aspects)

On the complexity
of the Block Low-Rank
multifrontal factorization
Sparse Days 2016
June 30-July 1, Toulouse
(A detailed complexity study with both theoretical and experimental results)

→ Based on the paper of the same name
(submitted to SIAM SISC)

Performance and scalability
of a Block Low-Rank
multifrontal solver
PMAA'16
July 6-8, Bordeaux
(A detailed performance analysis on real-life applications)