

*Bridging the gap between flat and hierarchical  
low-rank matrix formats: the multilevel BLR  
format*

Amestoy, Patrick and Buttari, Alfredo and L'Excellent,  
Jean-Yves and Mary, Theo

2018

MIMS EPrint: **2018.12**

Manchester Institute for Mathematical Sciences  
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary  
School of Mathematics  
The University of Manchester  
Manchester, M13 9PL, UK

ISSN 1749-9097

# BRIDGING THE GAP BETWEEN FLAT AND HIERARCHICAL LOW-RANK MATRIX FORMATS: THE MULTILEVEL BLR FORMAT

PATRICK R. AMESTOY\*, ALFREDO BUTTARI†, JEAN-YVES L'EXCELLENT‡, AND THEO MARY§

**Abstract.** Matrices possessing a low-rank property arise in numerous scientific applications. This property can be exploited to provide a substantial reduction of the complexity of their  $LU$  or  $LDL^T$  factorization. Among the possible low-rank formats, the flat Block Low-Rank (BLR) format is easy to use but achieves superlinear complexity. Alternatively, the hierarchical formats achieve linear complexity at the price of a much more complex, hierarchical matrix representation. In this paper, we propose a new format based on multilevel BLR approximations: the matrix is recursively defined as a BLR matrix whose full-rank blocks are themselves represented by BLR matrices. We call this format *multilevel BLR* (MBLR). Contrarily to hierarchical matrices, the number of levels in the block hierarchy is fixed to a given constant; while this format can still be represented within the  $\mathcal{H}$  formalism, we show that applying the  $\mathcal{H}$  theory to it leads to very pessimistic complexity bounds. We therefore extend the theory to prove better bounds, and show that the MBLR format provides a simple way to finely control the desired complexity of dense factorizations. By striking a balance between the simplicity of the BLR format and the low complexity of the hierarchical ones, the MBLR format bridges the gap between flat and hierarchical low-rank matrix formats. The MBLR format is of particular relevance in the context of sparse direct solvers, for which it is able to trade off the optimal dense complexity of the hierarchical formats to benefit from the simplicity and flexibility of the BLR format while still achieving  $O(n)$  sparse complexity. We finally compare our MBLR format with the related BLR- $\mathcal{H}$  (or Lattice- $\mathcal{H}$ ) format; our theoretical analysis shows that both formats achieve the same asymptotic complexity for a given top level block size.

**Key words.** low-rank approximations, matrix factorization, sparse linear algebra, hierarchical matrices

**AMS subject classifications.** 15a06, 15a23, 65f05, 65f50, 65n30, 65y20

**1. Introduction.** Efficiently computing the solution of a dense linear system is a fundamental building block of numerous scientific computing applications. Let us refer to such a system as

$$Fu_F = v_F, \tag{1.1}$$

where  $F$  is a dense matrix of order  $m$ ,  $u_F$  is the unknown vector of size  $m$ , and  $v_F$  is the right-hand side vector of size  $m$ .

This paper focuses on solving (1.1) with direct approaches based on Gaussian elimination, which consist in factorizing matrix  $F$  as  $F = LU$  or  $F = LDL^T$ , depending on whether the matrix is unsymmetric or symmetric, respectively.

In many applications (e.g., Schur complements arising from the discretization of elliptic partial differential equations), the matrix  $F$  has been shown to have a low-rank property: many of its off-diagonal blocks can be approximated by low-rank matrices [9].

Several formats have been proposed to exploit this property. The simplest one is the Block Low-Rank (BLR) format [2], which partitions the matrix with a flat, 2D blocking and approximates its off-diagonal blocks by low-rank submatrices, as illustrated in Figure 1.1a. Compared with the cubic  $O(m^3)$  complexity of the dense full-rank  $LU$  or  $LDL^T$  factorizations, the complexity of the dense BLR factorization can be as low as  $O(m^2)$  [4].

More advanced formats are based on a hierarchical partitioning of the matrix: the matrix  $F$  is partitioned with a  $2 \times 2$  blocking and the two diagonal blocks are recursively refined, as illustrated in Figures 1.1c and 1.1d. Different hierarchical formats can be defined depending on whether the off-diagonal blocks are directly approximated (so-called *weakly-admissible* formats) or further refined (so-called *strongly-admissible* formats). The most general of the hierarchical formats is the strongly-admissible  $\mathcal{H}$ -matrix format [9, 19, 10]; the HODLR format [5] is its weakly-admissible counterpart. These hierarchical formats can factorize a dense matrix in near-linear complexity  $O(m \log^q m)$ , where  $q$  is a small integer that depends on which factorization algorithm is used; in the following, we will consider  $q = 2$ . The log factor can be removed by using a so-called *nested-basis*

\*University of Toulouse, INP-IRIT

†University of Toulouse, CNRS-IRIT

‡University of Lyon, CNRS, ENS de Lyon, Inria, UCBL, LIP UMR5668, France

§University of Manchester

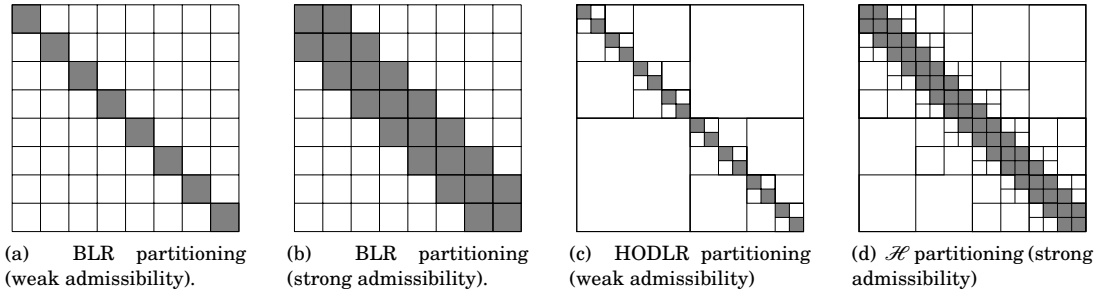


Fig. 1.1: Illustration of different low-rank formats. Gray blocks are stored in full-rank whereas white ones are approximated by low-rank matrices.

47 structure. The strongly-admissible  $\mathcal{H}^2$ -matrix format [10] and the weakly-admissible HSS [29, 11]  
 48 and HBS [17] formats exploit such nested basis structures to achieve linear complexity  $O(m)$ .

49 In this paper, we propose a new format based on multilevel BLR approximations. The matrix  
 50  $F$  is recursively represented as a BLR matrix whose full-rank blocks are themselves BLR matrices.  
 51 We call this format *multilevel BLR* (MBLR). In the hierarchical format, the matrix is refined until  
 52 the diagonal blocks are of constant size; this therefore leads to a number of levels in the block  
 53 hierarchy which is nonconstant (usually  $O(\log_2 m)$ ). With the MBLR format, we propose to make  
 54 this number of levels a tunable parameter to be set to a given value  $\ell$  that does not asymptotically  
 55 depend on the matrix size  $m$ . We prove that this parameter provides a simple way to finely control  
 56 the complexity of the dense MBLR factorization. The complexity varies from  $O(m^2)$  for monolevel  
 57 BLR down to nearly  $O(m)$  for an infinite number of levels. By striking a balance between the  
 58 simplicity of the BLR format and the low complexity of the hierarchical ones, the MBLR format  
 59 bridges the gap between flat and hierarchical low-rank matrix formats.

60 We will show that the MBLR format is of particular relevance in the context of sparse direct  
 61 solvers, which aim to compute the solution of a sparse linear system

$$62 \quad Au_A = v_A, \quad (1.2)$$

63 where  $A$  is a sparse matrix of order  $n$ ,  $u_A$  is the unknown vector of size  $n$ , and  $v_A$  is the right-hand  
 64 side vector of size  $n$ . Two widely studied classes of sparse direct methods are the multifrontal [13,  
 65 22] and supernodal [8, 12] approaches.

66 Sparse direct methods rely on a sequence of partial factorizations of dense matrices  $F$ , referred  
 67 to as supernodes or fronts. Therefore, the complexity of sparse direct methods is directly derived  
 68 from the complexity of the factorization of each dense matrix  $F$ . For example, with an adequate  
 69 reordering, a well-known result is that the dense standard full-rank  $O(m^3)$  factorization leads to a  
 70  $O(n^2)$  sparse complexity for regular 3D problems [14]. The low-rank formats described above can  
 71 be efficiently exploited within sparse solvers to provide a substantial reduction of their complexity.

72 The potential of BLR sparse solvers has been first investigated in [2]; the simplicity and flexi-  
 73 bility of the BLR format makes it easy to use in the context of a general purpose, algebraic solver,  
 74 as presented in [?, 3, 26, 24]. [?] focuses on the multicore performance of BLR multifrontal solvers,  
 75 while [3] and [26] present their use in two real-life industrial applications coming from geosciences.  
 76 [24] present the use of the BLR format in supernodal solvers. Furthermore, it has been proved in [4]  
 77 that the theoretical complexity of the BLR multifrontal factorization may be as low as  $O(n^{4/3})$  (for  
 78 3D problems with constant ranks).

79 Alternatively, most sparse solvers based on the more complex hierarchical formats have been  
 80 shown to possess near-linear complexity. To cite a few, [28, 27, 16, 15] are HSS-based, [17] is  
 81 HBS-based, [6] is HODLR-based, and [25] is  $\mathcal{H}^2$ -based.

82 However, a critical observation is that achieving  $O(n)$  sparse complexity does not actually  
 83 require a linear dense complexity  $O(m)$ . For instance, for 3D problems, all that is required is a  
 84 dense complexity lower than  $O(m^{1.5})$ . Therefore, we will prove that the MBLR format is able to

85 trade off the optimal dense complexity of the hierarchical formats to benefit from the simplicity  
 86 and flexibility of the flat BLR format while still achieving  $O(n)$  sparse complexity.

87 We now describe the organization of the rest of this paper. In Section 2, we provide some  
 88 background on the BLR factorization and its complexity and we motivate the key idea behind  
 89 MBLR approximations. We explain in Section 3 how the MBLR format can be described using the  
 90 cluster tree representation commonly used in the  $\mathcal{H}$  literature; this provides a convenient way to  
 91 explain the key difference between the MBLR and  $\mathcal{H}$  formats. We show that, similarly to the BLR  
 92 case, the  $\mathcal{H}$  theoretical formalism leads to MBLR complexity bounds that are very pessimistic. We  
 93 therefore extend the theory, beginning by the two-level case in Section 4. We prove that two levels  
 94 can already significantly improve the theoretical complexity of the factorization. In Section 5, we  
 95 generalize the previous proof to the MBLR format with an arbitrary number of levels; we prove  
 96 that, for constant ranks, only four levels are already enough to reach  $O(n)$  sparse 3D complexity  
 97 (and three levels already achieve near-linear  $O(n \log n)$  complexity). In Section 6, we validate our  
 98 theoretical results with numerical experiments. We provide our concluding remarks in Section 7.  
 99 In the main body of this article, we consider for the sake of simplicity the weakly-admissible case, in  
 100 which only the diagonal blocks are refined. In the appendix, we provide the extension of the MBLR  
 101 format to the strongly-admissible case, in which off-diagonal full-rank blocks are also recursively  
 102 refined. We prove that our complexity bounds remain valid in this context.

## 103 2. Background and motivation.

104 **2.1. Block Low-Rank approximations.** The BLR format is based on a flat, non-hierarchical  
 105 blocking of the matrix which is defined by conveniently clustering the associated unknowns. A BLR  
 106 representation  $\tilde{F}$  of a dense matrix  $F$  is shown in (2.1), where we assume that  $p \times p$  blocks have been  
 107 defined. Off-diagonal blocks  $F_{ij}$  ( $i \neq j$ ) of size  $m_i \times n_j$  and numerical rank  $k_{ij}^\varepsilon$  are approximated by  
 108 a low-rank matrix  $\tilde{F}_{ij} = X_{ij}Y_{ij}^T$  at accuracy  $\varepsilon$ , where  $X_{ij}$  is a  $m_i \times k_{ij}^\varepsilon$  matrix and  $Y_{ij}$  is a  $n_j \times k_{ij}^\varepsilon$   
 109 matrix. The diagonal blocks  $F_{ii}$  are stored as full-rank matrices ( $\tilde{F}_{ii} = F_{ii}$ ).

$$110 \quad \tilde{F} = \begin{bmatrix} \tilde{F}_{11} & \tilde{F}_{12} & \cdots & \tilde{F}_{1p} \\ \tilde{F}_{21} & \cdots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \vdots \\ \tilde{F}_{p1} & \cdots & \cdots & \tilde{F}_{pp} \end{bmatrix}. \quad (2.1)$$

111 Throughout this article, we will note  $F_{ij}$  the  $(i, j)$ -th block of  $F$  and  $F_{:,k}$  its  $k$ -th block-column.  
 112 We will also assume that all blocks have a size of order  $b$ , i.e.,  $m_i = n_j = O(b)$ .

113 Computing the low-rank approximation  $\tilde{F}_{ij}$  to each block, referred to as the *compression* step,  
 114 can be performed in different ways. We have chosen to use a truncated  $QR$  factorization with  
 115 column pivoting; this corresponds to a  $QR$  factorization with pivoting which is truncated as soon  
 116 as a diagonal coefficient of the  $R$  factor falls below the prescribed threshold  $\varepsilon$ . For a block of size  
 117  $b \times b$  and rank  $r$ , the cost of the compression is  $O(b^2r)$ , whereas computing the exact singular value  
 118 decomposition of the block would require  $O(b^3)$  operations. This choice thus allows for a convenient  
 119 compromise between cost and accuracy of the compression operation.

120 **2.2. Block Low-Rank dense factorization.** We describe in Algorithm 2.1 the CUFS variant  
 121 (the ‘‘CUFS’’ acronym is explained below) of the BLR factorization algorithm for dense matrices,  
 122 introduced in [4]. Algorithm 2.1 is presented in its  $LU$  version, but it can easily be adapted to the  
 123 symmetric case.

124 In order to perform the  $LU$  or  $LDL^T$  factorization of a dense BLR matrix, the standard block  
 125  $LU$  or  $LDL^T$  factorization has to be modified so that the low-rank blocks can be exploited to per-  
 126 form fast operations. Many such algorithms can be defined depending on where the compression  
 127 step is performed. As described in [4], the CUFS variant achieves the lowest complexity of all BLR  
 128 variants by performing the compression as early as possible.

129 This algorithm is referred to as CUFS (standing for Compress, Update, Factor, and Solve), to  
 130 indicate the order in which the steps are performed. All low-rank updates of a given block  $\tilde{F}_{ik}$  are  
 131 accumulated together before being recompressed, in order to achieve the smallest rank possible for  
 132  $\tilde{F}_{ik}$ .

---

**Algorithm 2.1** Dense BLR  $LU$  factorization: CUFS variant.

---

Input: a  $p \times p$  block matrix  $F$  of order  $m$ .

Output:  $F$  overwritten by its BLR LU factors  $\tilde{F}$ .

```

1: for  $k = 1$  to  $p$  do
2:   for  $i = k + 1$  to  $p$  do
3:     Compress ( $L$ ):  $F_{ik} \leftarrow \tilde{F}_{ik} = X_{ik} Y_{ik}^T$ 
4:     Compress ( $U$ ):  $F_{ki} \leftarrow \tilde{F}_{ki} = Y_{ki} X_{ki}^T$ 
5:   end for
6:   for  $i = k$  to  $p$  do
7:     for  $j = 1$  to  $k - 1$  do
8:       Update ( $L$ ):  $\tilde{F}_{ik} \leftarrow \tilde{F}_{ik} - X_{ij} Y_{ij}^T Y_{jk} X_{jk}^T$ 
9:       Update ( $U$ ):  $\tilde{F}_{ki} \leftarrow \tilde{F}_{ki} - X_{kj} Y_{kj}^T Y_{ji} X_{ji}^T$ 
10:    end for
11:     $\tilde{F}_{ik} \leftarrow \text{Recompress}(\tilde{F}_{ik})$ 
12:     $\tilde{F}_{ki} \leftarrow \text{Recompress}(\tilde{F}_{ki})$ 
13:  end for
14:  Factor:  $F_{kk} = L_{kk} U_{kk}$ 
15:  for  $i = k + 1$  to  $p$  do
16:    Solve ( $L$ ):  $\tilde{F}_{ik} \leftarrow \tilde{F}_{ik} U_{kk}^{-1} = X_{ik} Y_{ik}^T U_{kk}^{-1}$ 
17:    Solve ( $U$ ):  $\tilde{F}_{ki} \leftarrow L_{kk}^{-1} \tilde{F}_{ki} = L_{kk}^{-1} X_{ki} Y_{ki}^T$ 
18:  end for
19: end for

```

---

133 The CUFS BLR variant is referred to as *fully-structured*, which means the off-diagonal low-  
134 rank blocks  $\tilde{F}_{ik}$  are never stored in full-rank again after being initially compressed. Furthermore,  
135 in the rest of this article, we will assume that the matrix is already available under compressed  
136 form, that is, that the cost of the Compress step is negligible with respect to the total complexity of  
137 the factorization. This is for example the case when the blocks of the original matrix  $F$  are sparse,  
138 since their low-rank representation can be computed in only  $O(br)$  flops.

139 One of the main results of [4] is that the storage complexity of the factorization of a dense  
140 matrix of order  $m$  with off-diagonal blocks of rank at most  $r$  is equal to

$$141 \quad \mathcal{S}_{ds}^1(m, r) = O(m^{1.5} \sqrt{r}) \quad (2.2)$$

142 and the flop complexity is

$$143 \quad \mathcal{F}_{ds}^1(m, r) = O(m^2 r). \quad (2.3)$$

144 The proof of this result will be recalled in Section 2.4. The 1 superscript refers to the monolevel  
145 BLR factorization. This notation will be generalized in the next sections.

146 **2.3. BLR sparse factorization.** Because sparse direct factorizations such as multifrontal  
147 or supernodal approaches rely on dense factorizations, block low-rank approximations can easily  
148 be incorporated into the sparse factorization by representing the fronts (or supernodes) with the  
149 chosen low-rank format. For example, in the BLR case, the fronts are represented as defined  
150 by (2.1), and Algorithm 2.1 is adapted to perform their partial factorization. This is described in  
151 detail in [2, 23].

152 As a consequence, the complexity of the sparse factorization can be directly computed from  
153 the complexity of the dense factorization. We consider a matrix of order  $n = N^d$ , where  $d$  denotes  
154 the problem dimension, that is reordered using nested dissection [14]: the domain is recursively  
155 partitioned by so-called *separators*. The sparse complexity is then computed as follows (assuming  
156 cross-shaped separators): at each level  $\ell$  of the separator tree, we need to factorize  $(2^d)^\ell$  fronts of  
157 order  $O((N/2^\ell)^{d-1})$ , for  $\ell$  ranging from 0 to  $L = \log_2(N)$ . Therefore, the flop complexity  $\mathcal{F}_{sp}(N)$  is  
158 equal to

$$159 \quad \mathcal{F}_{sp}(N) = \sum_{\ell=0}^L (2^d)^\ell \mathcal{F}_{ds}^1\left(\left(\frac{N}{2^\ell}\right)^{d-1}\right), \quad (2.4)$$

Table 2.1: Flop and storage complexities of the factorization of a sparse system of  $n = N \times N$  (2D case) or  $n = N \times N \times N$  (3D case) unknowns, assuming a dense factorization complexity  $O(m^\beta)$ .

2D		3D	
$\beta > 2$	$O(n^{\beta/2})$	$\beta > 1.5$	$O(n^{2\beta/3})$
$\beta = 2$	$O(n \log n)$	$\beta = 1.5$	$O(n \log n)$
$\beta < 2$	$O(n)$	$\beta < 1.5$	$O(n)$

Table 2.2: Flop and storage complexities of the FR, BLR, and  $\mathcal{H}$  factorizations of a sparse system of  $n = N \times N$  (2D case) or  $n = N \times N \times N$  (3D case) unknowns, derived from the complexities of the factorization of a dense matrix of order  $m$  and with a rank bound  $r$ . The BLR variant considered is CUFS.

	$\mathcal{F}_{ds}(m,r)$	$\mathcal{F}_{sp}(n,r)$		$\mathcal{S}_{ds}(m,r)$	$\mathcal{S}_{sp}(n,r)$	
		2D	3D		2D	3D
FR	$O(m^3)$	$O(n^{3/2})$	$O(n^2)$	$O(m^2)$	$O(n \log n)$	$O(n^{4/3})$
BLR	$O(m^2 r)$	$O(n \max(r, \log n))$	$O(n^{4/3} r)$	$O(m^{3/2} r^{1/2})$	$O(n)$	$O(n \max(r^{1/2}, \log n))$
$\mathcal{H}$	$O(m r^2 \log^2 m)$	$O(\max(n, n^{1/2} r^2))$	$O(\max(n, n^{2/3} r^2))$	$O(m r \log m)$	$O(n)$	$O(\max(n, n^{2/3} r))$

160 where  $\mathcal{F}_{ds}(m)$  is the dense flop complexity. In BLR, it is given by (2.3). Similarly, the storage  
 161 complexity  $\mathcal{S}_{sp}(N)$  is equal to

$$162 \quad \mathcal{S}_{sp}(N) = \sum_{\ell=0}^L (2^d)^\ell \mathcal{S}_{ds} \left( \left( \frac{N}{2^\ell} \right)^{d-1} \right), \quad (2.5)$$

163 where  $\mathcal{S}_{ds}(m)$  is the dense storage complexity. In BLR, it is given by (2.2).

164 Assuming a dense complexity  $O(m^\beta)$ , it can easily be shown from (2.4) and (2.5) that the sparse  
 165 complexities only depend on the dense complexity exponent  $\beta$  and the dimension  $d$ . This correspon-  
 166 dence is reported in Table 2.1. The key observation is that a *linear*  $O(m)$  dense complexity is *not*  
 167 *required to achieve a linear*  $O(n)$  sparse complexity. In fact, a dense complexity lower than  $O(m^2)$   
 168 and  $O(m^{1.5})$  suffices for 2D and 3D problems, respectively.

169 Then, the sparse complexities are reported in Table 2.2, for the FR, BLR, and  $\mathcal{H}$  factorizations.  
 170 Thanks to the key observation above, the BLR sparse complexities are not that far from the  $\mathcal{H}$   
 171 complexities. For example, the BLR 2D storage complexity is already optimal; furthermore, the 2D  
 172 flop and 3D storage complexities are nearly linear, with only an additional factor depending only  
 173 on  $r$  and  $\log n$  compared with  $\mathcal{H}$ . More importantly, thanks to the same key observation, we only  
 174 need a modest improvement of the dense complexity to reach  $O(n)$  complexity. Specifically:

- 175 • To drop the  $\max(r, \log n)$  factor in the 2D flop complexity, the dense complexity  $\mathcal{F}_{ds}(m)$  only  
 176 needs to be strictly inferior to  $O(m^2)$ ;
- 177 • Similarly, to drop the  $\max(r^{1/2}, \log n)$  factor in the 3D storage complexity, the dense com-  
 178 plexity  $\mathcal{S}_{ds}(m)$  only needs to be strictly inferior to  $O(m^{1.5})$ ;
- 179 • Finally, the superlinear 3D flop complexity can be made linear with a dense complexity  
 180  $\mathcal{F}_{ds}(m)$  strictly inferior to  $O(m^{1.5})$ .

181 The main motivation behind the MBLR format is to find a *simple* modification of the BLR fac-  
 182 torization that preserves its simplicity and flexibility, while improving the complexity *just enough*  
 183 to get the desired exponent. We will prove in Section 5 that this complexity improvement can be  
 184 controlled by the number of levels.

185 **2.4. Complexity analysis for BLR.** To motivate the key idea behind MBLR, let us sum-  
 186 marize the dense storage and flop complexity analysis found in [4] that leads to the formulas (2.2)  
 187 and (2.3). We consider the CUFS variant of the BLR factorization. As indicated in the introduction,  
 188 we first consider the weakly-admissible case: we assume that all off-diagonal blocks are low-rank.  
 189 The extension to the strongly-admissible case is provided in the appendix.

Table 2.3: Main operations for the BLR factorization of a dense matrix of order  $m$ , with blocks of size  $b$ , and low-rank blocks of rank at most  $r$ . We note  $p = m/b$ . Type: type of the block(s) on which the operation is performed.  $cost_{Step}^1$ : cost of performing the operation once.  $number_{Step}$ : number of times the operation is performed.  $\mathcal{F}_{Step}^1$ : obtained by multiplying the  $cost_{Step}^1$  and  $number_{Step}$  columns (equation (2.11)). The first expression is given as a function of  $b$ ,  $p$ , and  $r$ , while the second is obtained with the assumption that  $b = O(m^x)$  (and thus  $p = O(m^{1-x})$ ) and  $r = O(m^\alpha)$ , for  $x, \alpha \in [0, 1]$ .

Step	Type	$cost_{Step}^1$	$number_{Step}$	$\mathcal{F}_{Step}^1(b, p, r)$	=	$\mathcal{F}_{Step}^1(m, x, \alpha)$
Factor	FR	$O(b^3)$	$O(p)$	$O(pb^3)$	=	$O(m^{1+2x})$
Solve	FR-LR	$O(b^2r)$	$O(p^2)$	$O(p^2b^2r)$	=	$O(m^{2+\alpha})$
Product	LR-LR	$O(br^2)$	$O(p^3)$	$O(p^3br^2)$	=	$O(m^{3-2x+2\alpha})$
Recompress	LR	$O(bpr^2)$	$O(p^2)$	$O(p^3br^2)$	=	$O(m^{3-2x+2\alpha})$

190 We consider a dense BLR matrix of order  $m$ . We note  $b$  the block size and  $p = m/b$  the number  
 191 of blocks per row and/or column. The amount of storage required to store the factors of such a  
 192 matrix can be computed as the sum of the storage for the full-rank diagonal blocks and that for the  
 193 low-rank off-diagonal blocks:

$$194 \quad \mathcal{S}_{ds}^1(b, p, r) = O(pb^2) + O(p^2br). \quad (2.6)$$

195 Then, we assume that the block size  $b$  is of order  $O(m^x)$ , where  $x$  is a real value in  $[0, 1]$ , and thus  
 196 the number of blocks  $p$  per row and/or column is of order  $O(m^{1-x})$ . We also assume that the rank  
 197 bound is of the form  $r = O(m^\alpha)$ . By replacing  $b$ ,  $p$ , and  $r$  by their expression in (2.6), we obtain an  
 198 expression of  $\mathcal{S}_{ds}^1$  which depends on  $(m, x, \alpha)$  instead of  $(b, p, r)$ :

$$199 \quad \mathcal{S}_{ds}^1(m, x, \alpha) = O(m^{1+x}) + O(m^{2-x+\alpha}). \quad (2.7)$$

200 We then define  $x^*$  as the optimal choice of  $x$  which minimizes the asymptotic complexity of (2.7).  
 201  $x^*$  can be computed as the value which makes each term in (2.7) asymptotically equal. We obtain

$$202 \quad x^* = (1 + \alpha)/2, \quad (2.8)$$

203 which means the optimal choice of block size is

$$204 \quad b^* = O(\sqrt{mr}). \quad (2.9)$$

205 This leads to the final dense complexity (already given in (2.2))

$$206 \quad \mathcal{S}_{ds}^1(m, r) = O(m^{1.5}\sqrt{r}). \quad (2.10)$$

207 Next, to compute the flop complexity, we compute the cost of the Factor, Solve, Product, and  
 208 Recompress steps and report them in Table 2.3 (third column). This cost depends on the type (full-  
 209 rank or low-rank) of the block(s) on which the operation is performed (second column). Note that  
 210 the Product operation can only take the form of a product of two low-rank blocks (LR-LR), because  
 211 it involves only off-diagonal blocks, which are low-rank in the weakly-admissible case. We also  
 212 remind that we have assumed that the matrix is already available under compressed form and  
 213 thus we do not report the Compress step in Table 2.3.

214 In the weakly-admissible case, there is only one full-rank block on each block-row (the diagonal  
 215 one); therefore, we can easily count the number of blocks on which each step is performed; we report  
 216 it on the fourth column of Table 2.3. The BLR factorization cost of each step is then equal to

$$217 \quad \mathcal{F}_{Step}^1 = cost_{Step}^1 \times number_{Step} \quad (2.11)$$

218 and is reported in the fifth and sixth columns of Table 2.3. In the fifth column, its expression  
 219 depends on  $b$ ,  $p$ , and  $r$ , while in the sixth column it is given as a function of  $m$ ,  $x$ , and  $\alpha$  by  
 220 substituting  $b$ ,  $p$ , and  $r$  by  $O(m^x)$ ,  $O(m^{1-x})$ , and  $O(m^\alpha)$ , respectively.




Fig. 2.1: Comparison of BLR, MBLR, and hierarchical formats in the weakly-admissible case.

221 The total flop complexity of the dense BLR factorization is equal to the sum of the cost of all  
 222 steps

$$223 \mathcal{F}_{ds}^1(m, x, \alpha) = O(m^{1+2x} + m^{2+\alpha} + m^{3-2x+2\alpha}). \quad (2.12)$$

224 We compute  $x^*$ , the optimal choice of  $x$  which minimizes the complexity, and find again  $x^* = (1 +$   
 225  $\alpha)/2$ , which means that the same  $x^*$  value minimizes both the storage and flop complexities, a  
 226 valuable property. We finally obtain

$$227 \mathcal{F}_{ds}^1(m, r) = O(m^2 r). \quad (2.13)$$

228 **2.5. Key idea of the MBLR format.** Let us now consider each step of Table 2.3 with the  
 229 objective of reducing the total cost of the factorization. The Product and Recompress steps involve  
 230 exclusively low-rank blocks and their cost is already optimal as it is linear with respect to the block  
 231 size  $b$ . Therefore, we focus on the Factor and Solve steps. These steps have superlinear cost with  
 232 respect to  $b$  because they involve the diagonal full-rank blocks.

233 Thus, the key idea of the MBLR format is to further refine these full-rank diagonal blocks by  
 234 replacing them by BLR matrices. This is illustrated in Figure 2.1b. Compared with the BLR format  
 235 (Figure 2.1a), we will show in Section 4 that this more compressed representation decreases the  
 236 cost of performing the Factor and Solve steps, which allows to increase the block size to achieve a  
 237 lower complexity. However, it also remains very different from a hierarchical format (Figure 2.1c).  
 238 First, the diagonal blocks are represented by the simple BLR format, rather than a more complex  
 239 hierarchical format. Second, while larger than in the BLR case, the off-diagonal blocks are in  
 240 general still much smaller than in the  $\mathcal{H}$  case. This has several advantages:

- 241 • No relative order is needed between blocks; this allows the clustering to easily be computed  
 242 and delivers a great flexibility to distribute the data in a parallel environment.
- 243 • The size of the blocks can be small enough to fit on a single shared-memory node; there-  
 244 fore, in a parallel environment, each processor can efficiently and concurrently work on  
 245 different blocks.
- 246 • To perform numerical pivoting, the quality of the pivot candidates in the off-diagonal  
 247 blocks  $F_{ik}$  can be estimated with the entries of its low-rank basis  $Y_{ik}$ , provided that  $X_{ik}$   
 248 is an orthonormal matrix. However, as explained in [23], the quality of this estimation  
 249 depends on the size of the block; smaller blocks lead to a tighter estimate. This makes the  
 250 BLR format particularly suitable to handle numerical pivoting, a critical feature lacking  
 251 in most hierarchical solvers presented in the literature.

252 While many advantages of the BLR and MBLR formats lie in their efficiency and flexibility in  
 253 the context of a parallel execution, the parallel implementation of the MBLR format is out of the  
 254 scope of this paper. Similarly, we omit a detailed description of the algorithms designed to handle  
 255 numerical pivoting; see [23] for a thorough discussion. In this paper, we focus on the theoretical  
 256 complexity analysis of the MBLR factorization.

257 **3. Difference between MBLR and  $\mathcal{H}$  matrices.** In this section, we first explain how  
 258 MBLR matrices can be represented using the cluster tree modelization typically used in the  $\mathcal{H}$






Fig. 3.1: An example of cluster tree and its associated HODLR block-clustering.

259 literature; this provides a convenient tool to formalize the key difference between the MBLR and  
 260 hierarchical formats that was informally presented in the previous section: the number of levels  
 261 in the cluster tree is a constant in the MBLR format, while it is logarithmically dependent on the  
 262 problem size in the  $\mathcal{H}$  format.

263 Using this formalism,  $\mathcal{H}$  theory is thus applicable to the MBLR format; however, we show in  
 264 Section 3.3 that it leads to very pessimistic complexity bounds. We must therefore develop a new  
 265 theory to compute satisfying bounds, which is the object of Sections 4 and 5. Since the proofs and  
 266 computations in these sections are *not* based on the  $\mathcal{H}$  theory, this section may be skipped by the  
 267 reader, at least on a first read.

268 **3.1. The hierarchical case.** Here, we briefly remind the definition of cluster trees, and how  
 269 they are used to represent hierarchical partitionings. We refer to [9, 20] for a formal and detailed  
 270 presentation.

271 Let us note  $\mathcal{I}$  the set of unknowns. We assume that the sets of row and column indices of the  
 272 matrix are the same, for the sake of simplicity, and because we do not need to distinguish them for  
 273 the purpose of this section.

274 Computing a recursive partition  $\mathfrak{S}(\mathcal{I})$  of  $\mathcal{I}$  can be modeled with a so-called *cluster tree*.

275 **DEFINITION 3.1 (Cluster tree).** Let  $\mathcal{I}$  be a set of unknowns and  $T_{\mathcal{I}}$  a tree whose nodes  $v$  are  
 276 associated with subsets  $\sigma_v$  of  $\mathcal{I}$ .  $T_{\mathcal{I}}$  is said to be a cluster tree iff:

- 277 • The root of  $T_{\mathcal{I}}$ , noted  $r$ , is associated with  $\sigma_r = \mathcal{I}$ ;
- 278 • For each non-leaf node  $v \in T_{\mathcal{I}}$ , with children noted  $C_{T_{\mathcal{I}}}(v)$ , the subsets associated with the  
 279 children form a partition of  $\sigma_v$ , that is, the  $\sigma_c$  subsets are disjoint and satisfy

$$280 \quad \bigcup_{c \in C_{T_{\mathcal{I}}}(v)} \sigma_c = \sigma_v.$$

281 An example of cluster tree is provided in Figure 3.1a. As illustrated, cluster trees establish a  
 282 *hierarchy* between clusters. A given cluster tree uniquely defines a weakly-admissible, hierarchical  
 283 block-clustering (so-called HODLR matrix), as illustrated in Figure 3.1b.

284 Note that the cluster tree is not enough to define general, strongly-admissible hierarchical  
 285 block-clusterings such as an  $\mathcal{H}$  block-clustering. In that case, a new tree structure, so-called  
 286 block-cluster tree, must be introduced. For the sake of simplicity, we do not discuss them here,  
 287 since cluster trees (and weakly-admissible hierarchical block-clusterings) are enough to explain  
 288 the difference between MBLR and hierarchical matrices.

289 **3.2. The BLR and MBLR formats, explained with cluster trees.** It is interesting to first  
 290 mention how the BLR format can be viewed as a very particular kind of  $\mathcal{H}$ -matrix. In [4], we  
 291 made a first attempt to model BLR partitionings with cluster trees, where the BLR partitioning  
 292 is defined using only the leaves of the cluster tree. However, this model is inadequate because  
 293 the other nodes of the tree are unnecessary (and thus there are several possible cluster trees for a  
 294 unique BLR partitioning).

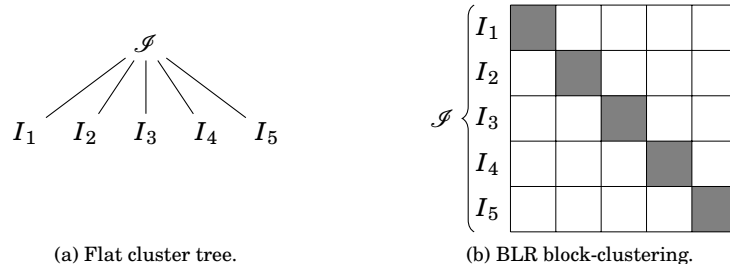


Fig. 3.2: An example of flat cluster tree and its associated BLR block-clustering.




Fig. 3.3: An example of cluster tree of depth 2 and its associated two-level BLR block-clustering.

295 Therefore, we newly propose to define a given BLR partitioning with a unique, flat cluster tree  
 296 of depth 1, as illustrated in Figure 3.2.

297 With this definition, it is straightforward to extend the model to the MBLR format. An  $\ell$ -  
 298 level BLR block-clustering can be represented by a cluster tree of depth  $\ell$ . This is illustrated in  
 299 Figure 3.3 in the two-level case.

300 It is thus clear that, just like the BLR format, the MBLR one can also be viewed as a very  
 301 particular kind of  $\mathcal{H}$ -matrix, since any MBLR matrix can be represented with a cluster tree that  
 302 satisfies Definition 3.1. However, while  $\mathcal{H}$ -matrices are indeed also represented by cluster trees,  
 303 in practice, they are virtually always built with an implicit assumption: the number of children  
 304 of any node in the cluster tree is constant with respect to the size of the matrix. Most commonly,  
 305 cluster trees are binary trees, and thus this number is 2. Since the diagonal blocks of an  $\mathcal{H}$ -matrix  
 306 are of constant size, this leads to  $O(\log m)$  levels in the cluster tree.

307 The MBLR format is based on the converse approach: the number of levels is fixed to some  
 308 constant  $\ell = O(1)$ , which in turn leads to a number of children per node that is asymptotically  
 309 dependent on  $m$ . For example, assuming a constant rank  $r = O(1)$ , we will prove in Section 4 that a  
 310 two-level BLR matrix is represented by a cluster tree with  $O(m^{1/3})$  nodes on the first level (children  
 311 of the root node), and each of these nodes has  $O((m^{2/3})^{1/2}) = O(m^{1/3})$  children on the second level.

312 While this is technically not prohibited by the general  $\mathcal{H}$ -format definition, it has, to the best  
 313 of our knowledge, never been considered in the  $\mathcal{H}$  literature. As a matter of fact, in his recent  
 314 book, Hackbusch writes ([20], p.83): “The partition should contain as few blocks as possible since  
 315 the storage cost increases with the number of blocks”. While that is true, we believe that the  
 316 smaller size and greater number of blocks can provide a gain in flexibility that can be useful in a  
 317 parallel solver.

318 One may in fact find it surprising that this simple idea of having a constant number of levels  
 319 has never been proposed before. We believe that this is mainly explained by the fact that the  
 320  $\mathcal{H}$  theoretical formalism does not provide a satisfying result if applied to the MBLR format, as  
 321 explained in the following.

322 **3.3.  $\mathcal{H}$  theory leads to pessimistic MBLR complexity bounds.** In [4, Section 3.2], we  
 323 explained that applying the  $\mathcal{H}$  theoretical complexity formulas to the BLR format does not provide  
 324 a satisfying complexity bound. Here, we show this remains the case for the MBLR format.

325 The storage and flop complexities of the factorization of a dense  $\mathcal{H}$ -matrix of order  $m$  have  
 326 been shown [18] to be

$$327 \quad \mathcal{S}_{ds}^{\mathcal{H}}(m) = O(mLc_{sp} \max(r_{max}, b_{diag})), \quad (3.1)$$

$$328 \quad \mathcal{F}_{ds}^{\mathcal{H}}(m) = O(mL^2c_{sp}^3 \max(r_{max}, b_{diag})^2). \quad (3.2)$$

330  $L$  is the depth of the cluster tree.  $c_{sp}$  is the so-called sparsity constant, defined as the maximum  
 331 number of blocks of a given level in the cluster tree that are in the same row or column of the  
 332 matrix; it is a measure of how much the original matrix has been refined.  $r_{max}$  is the maximal  
 333 rank of all the blocks of the matrix, while  $b_{diag}$  is the size of the diagonal blocks.

334 Under some assumptions on how the partition  $\mathfrak{S}(\mathcal{S})$  is built [18, Lemma 4.5], the sparsity  
 335 constant can be bounded by  $O(1)$  in the  $\mathcal{H}$  case. By recursively refining the diagonal blocks until  
 336 they become of constant size, we have  $b_{diag} = O(1)$  and  $L = O(\log m)$ . Therefore, (3.1) and (3.2) lead  
 337 to  $\mathcal{S}_{ds}^{\mathcal{H}}(m) = O(rm \log m)$  and  $\mathcal{F}_{ds}^{\mathcal{H}}(m) = O(r^2 m \log^2 m)$ .

338 In the BLR case, we have  $L = 1$  and  $b_{diag} = b$ ; the sparsity constant is equal to the number of  
 339 blocks per row  $m/b$ , a high value that translates the fact that BLR matrices are much more refined  
 340 than hierarchical ones. This leads to  $\mathcal{S}_{ds}^{\mathcal{H},1}(m) = O(m^2)$  and  $\mathcal{F}_{ds}^{\mathcal{H},1}(m) = O(m^4/b) \geq O(m^3)$ , where  
 341 the notation  $\mathcal{S}_{ds}^{\mathcal{H},1}$  and  $\mathcal{F}_{ds}^{\mathcal{H},1}$  signifies “ $\mathcal{H}$  complexity formula applied to the monolevel BLR case”.  
 342 We thus obtain a very pessimistic complexity bound, which is due to the fact that the diagonal  
 343 blocks are of the same size as all the other blocks, i.e.,  $b_{diag} = b$ .

344 Applying the formulas to the MBLR case leads to the same problem. Assuming that  $r = O(1)$   
 345 for the sake of simplicity, let us consider the two-level case and denote by  $b_1$  and  $b_2$  the outer and  
 346 inner blocks sizes, respectively. We have  $L = 2$ ,  $b_{diag} = b_2$ , and  $c_{sp} = \max(m/b_1, b_1/b_2)$ . The storage  
 347 complexity

$$348 \quad \mathcal{S}_{ds}^{\mathcal{H},2}(m) = O(\max(m^2 b_2/b_1, m b_1))$$

349 is minimized for  $b_1 = O(\sqrt{m})$  and  $b_2 = O(1)$ , and equal to  $O(m^{3/2})$ , which is not a satisfying result  
 350 since we have proven in Section 2.4 that this is the BLR complexity. A similarly unsatisfying result  
 351 is obtained for the flop complexity, again minimized for  $b_1 = O(\sqrt{m})$  and  $b_2 = O(1)$  and equal to

$$352 \quad \mathcal{F}_{ds}^{\mathcal{H},2}(m) = O(\max(m^4 b_2^2/b_1^3, m b_1^3/b_2)) = O(m^{5/2}).$$

353 Does the problem persist with a higher number of levels  $\ell$ ? The answer is unfortunately positive.  
 354 Indeed, we have

$$355 \quad \mathcal{S}_{ds}^{\mathcal{H},\ell}(m) = O\left(\max_{i=0,\ell-1} \left(\frac{m b_i b_\ell}{b_{i+1}}\right)\right),$$

356 with the notation  $b_0 = m$ . To minimize the expression, we equilibrate each term in the maximum,  
 357 which leads to the recursive relation  $b_i^2 = b_{i-1} b_{i+1}$ . Noting  $b_1 = b$  the outer block size, it is clear  
 358 that the closed-form expression of  $b_i$  is  $b_i = b^i / m^{i-1}$ . Since the smallest block size  $b_\ell$  must be at  
 359 least  $O(1)$ , this leads to

$$360 \quad \frac{b^\ell}{m^{\ell-1}} \geq O(1)$$

361 and thus  $b \geq O(m^{(\ell-1)/\ell})$ , which in turn leads to the complexity bound

$$362 \quad \mathcal{S}_{ds}^{\mathcal{H},\ell}(m) \geq O\left(\frac{m^2 b_\ell}{b_1}\right) \geq O(m^{(\ell+1)/\ell}).$$

363 In Section 5, we will prove that for  $r = O(1)$  we have  $\mathcal{S}_{ds}^\ell(m) = O(m^{(\ell+2)/(\ell+1)})$ , which is a better  
 364 bound, especially when considering a small number of levels  $\ell$ . The analysis for the flop complexity  
 365 leads to the same expression of  $b_i$  and to the complexity bound  $\mathcal{F}_{ds}^{\mathcal{H},\ell}(m) \geq O(m^{(\ell+3)/\ell})$ , which is  
 366 again overly pessimistic compared with the bound  $\mathcal{F}_{ds}^\ell(m) = O(m^{(\ell+3)/(\ell+1)})$  that we will prove in  
 367 Section 5.

Table 3.1: Applying the  $\mathcal{H}$  theoretical complexity formulas to the BLR or MBLR cases does not provide a satisfying result. We have assumed that  $r = O(1)$  for the sake of clarity, but the bounds would remain pessimistic for general ranks.

	$\mathcal{H}$	BLR	MBLR
$L$	$O(\log m)$	1	$\ell$
$c_{sp}$	$O(1)$	$m/b$	$\max_{i=1,\ell} \frac{b_{i-1}}{b_i}^*$
$b^{diag}$	$O(1)$	$b$	$b_\ell$
$\mathcal{F}_{ds}^{\mathcal{H}}$	$O(m \log m)$	$O(m^2)$	$O(m^{(\ell+1)/\ell})$
$\mathcal{F}_{ds}^{\mathcal{H}}$	$O(m \log^2 m)$	$O(m^3)$	$O(m^{(\ell+3)/\ell})$

\*with the notation  $b_0 = m$

368 We summarize the result of applying the  $\mathcal{H}$  complexity formulas in Table 3.1.

369 It is therefore clear that we must develop a new theory extending the  $\mathcal{H}$  formalism to be able  
 370 to prove better MBLR complexity bounds, just as we did in [4] for the BLR case. We begin by the  
 371 two-level case in the next section.

372 **4. Two-level BLR matrices.** In this section, we compute the theoretical complexity of the  
 373 two-level BLR factorization. The proofs and computations on this particular two-level case are  
 374 meant to be illustrative of those of the general multilevel case with an arbitrary number of levels,  
 375 which is discussed in Section 5.

376 We remind that in this section, we are considering the weakly-admissible case only.

377 **4.1. Two-level kernels description.** In order to adapt Algorithm 2.1 to two-level BLR ma-  
 378 trices, two modifications must be performed.

379 First, the Factor kernel is not a full-rank factorization of the diagonal blocks  $F_{kk}$  anymore, but  
 380 a BLR factorization. Thus, line 14 must be replaced by

$$381 \quad \tilde{F}_{kk} = \tilde{L}_{kk} \tilde{U}_{kk} = \text{BLR-Factor}(F_{kk}), \quad (4.1)$$

382 where BLR-Factor refers to the BLR factorization described in Algorithm 2.1.

383 Second, the FR-LR Solve kernel at lines 16 and 17 must be changed to a BLR-LR Solve:

$$384 \quad \tilde{F}_{ik} \leftarrow \text{BLR-LR-Solve}(\tilde{U}_{kk}, \tilde{F}_{ik}) \quad (4.2)$$

$$385 \quad \tilde{F}_{ki} \leftarrow \text{BLR-LR-Solve}(\tilde{L}_{kk}, \tilde{F}_{ki}), \quad (4.3)$$

387 where  $\tilde{F}_{ik}$  and  $\tilde{F}_{ki}$  are LR matrices and  $\tilde{L}_{kk}$  and  $\tilde{U}_{kk}$  are lower and upper triangular BLR matrices.

388 We describe in Algorithm 4.1 the BLR-LR-Solve kernel, in the upper triangular case, omitting  
 389 the lower triangular case which is very similar. The kernel consists in applying a triangular solve  
 390 with an upper triangular BLR matrix  $\tilde{U}$  to a low-rank matrix  $\tilde{B} = \Phi\Psi^T$ . We remind that  $\tilde{U}_{ij}$  denotes  
 391 the  $(i, j)$ -th low-rank sub-block of  $\tilde{U}$  (with the notation  $\tilde{U}_{jj} = U_{jj}$ ), and that  $\Psi_{j,:}$  denotes the  $j$ -th  
 392 block-row of  $\Psi$ .

---

**Algorithm 4.1** BLR-LR-Solve kernel (upper triangular case)

---

Input: a  $p \times p$  block upper triangular BLR matrix  $\tilde{U}$ ;  $\tilde{U} = [\tilde{U}_{ij}]_{i=1:j,j=1:p}$  such that  $\tilde{U}_{ij} = Y_{ij}X_{ij}^T$  for  
 $i \neq j$  and  $\tilde{U}_{jj} = U_{jj}$  is a full-rank matrix; a LR matrix  $\tilde{B} = \Phi\Psi^T$ .

Output: overwritten  $\Psi$  (modified in-place) corresponding to the operation  $\tilde{B} \leftarrow \tilde{B}\tilde{U}^{-1}$ .

```

1: for  $j = 1$  to  $p$  do
2:    $\Psi_{j,:}^T \leftarrow \Psi_{j,:}^T - \sum_{i=1}^{j-1} \Psi_{i,:}^T Y_{ij} X_{ij}^T$ 
3:    $\Psi_{j,:}^T \leftarrow \Psi_{j,:}^T U_{jj}^{-1}$ 
4: end for

```

---

393 Two main operations must be performed: a triangular solve using the full-rank diagonal blocks  
 394  $U_{jj}$  of  $\tilde{U}$ , and an update using the low-rank off-diagonal blocks  $\tilde{U}_{ij} = Y_{ij}X_{ij}^T$  of  $\tilde{U}$ . Both are applied




Fig. 4.1: The BLR-LR-Solve kernel

395 on the low-rank block-columns  $\tilde{B}_{:,j} = \Phi \Psi_{j,:}^T$  of  $\tilde{B}$ . These two operations take place at lines 3 and 2  
 396 of Algorithm 4.1, respectively.

397 The FR-LR triangular solve can be written as

$$398 \quad \tilde{B}_{:,j} \leftarrow \tilde{B}_{:,j} U_{jj}^{-1} = \Phi \left( \Psi_{j,:}^T U_{jj}^{-1} \right) \quad (4.4)$$

399 and thus only  $\Psi_{j,:}^T$  needs to be updated, as shown at line 3.

400 The LR-LR update takes the following form:

$$401 \quad \begin{aligned} \tilde{B}_{:,j} \leftarrow \tilde{B}_{:,j} - \sum_{i=1}^{j-1} \tilde{B}_{:,i} \tilde{U}_{ij} &= \Phi \Psi_{j,:}^T - \sum_{i=1}^{j-1} \Phi \left( \Psi_{i,:}^T Y_{ij} X_{ij}^T \right) \\ &= \Phi \left( \Psi_{j,:}^T - \sum_{i=1}^{j-1} \Psi_{i,:}^T Y_{ij} X_{ij}^T \right) \end{aligned} \quad (4.5)$$

402 and thus, again, only  $\Psi_{j,:}^T$  needs to be updated, as shown at line 2.

403 The BLR-LR-Solve kernel is illustrated in Figure 4.1.

404 It can easily be computed that the cost of applying the BLR-LR-Solve kernel once is equal to  
 405 the storage complexity times the rank bound  $r$ :

$$406 \quad \text{cost}_{\text{Solve}}^2 = O(r) \times \mathcal{S}_{ds}^1(b, r). \quad (4.6)$$

407 Injecting (2.10) into (4.6) leads to

$$408 \quad \text{cost}_{\text{Solve}}^2 = O(b^{3/2} r^{3/2}). \quad (4.7)$$

409 **4.2. Two-level complexity analysis.** We now compute the complexity of the two-level BLR  
 410 factorization of a dense matrix  $F$  of order  $m$ . We reuse the monolevel notations for the top level.  
 411 We denote by  $b$  the size of the first-level blocks and  $p = m/b$  the number of block-rows and block-  
 412 columns. We assume that  $b$  is of the form  $b = m^x$ , for  $x \in [0, 1]$ . We also assume that the ranks of  
 413 the off-diagonal blocks are bounded by  $r = O(m^\alpha)$ .

414 The size required to store the factors is again the sum of the storage for the diagonal and off-  
 415 diagonal blocks, as in (2.6). The difference is that the diagonal blocks are not full-rank but BLR  
 416 matrices. They are further refined into smaller blocks whose size should be chosen of order  $O(\sqrt{br})$ ,  
 417 as determined by (2.9) in the BLR complexity analysis. Therefore, in the two-level case, (2.6)  
 418 becomes

$$419 \quad \mathcal{S}_{ds}^2(b, p, r) = p \times \mathcal{S}_{ds}^1(b, r) + O(p^2 br). \quad (4.8)$$

420 By replacing  $\mathcal{S}_{ds}^1(b, r)$  by its second expression computed in (2.10), we obtain

$$421 \quad \mathcal{S}_{ds}^2(b, p, r) = O(pb^{3/2} r^{1/2}) + O(p^2 br) \quad (4.9)$$

422 Finally, we replace  $b$ ,  $p$ , and  $r$  by their expression  $O(m^x)$ ,  $O(m^{1-x})$ , and  $O(m^\alpha)$ , respectively, to  
 423 obtain

$$424 \quad \mathcal{S}_{ds}^2(m, x, \alpha) = O(m^{1+(x+\alpha)/2} + m^{2-x+\alpha}). \quad (4.10)$$

425 This leads to

$$426 \quad x^* = (2 + \alpha)/3, \quad (4.11)$$

Table 4.1: Two-level equivalent of Table 2.3. The legend of the table is the same. The differences between the two tables are highlighted in gray.

Step	Type	$cost_{Step}^2$	$number_{Step}$	$\mathcal{F}_{Step}^2(b, p, r)$	=	$\mathcal{F}_{Step}^2(m, x, \alpha)$
Factor	BLR	$O(b^2 r)$	$O(p)$	$O(p b^2 r)$	=	$O(m^{1+x+\alpha})$
Solve	BLR-LR	$O(b^{3/2} r^{3/2})$	$O(p^2)$	$O(p^2 b^{3/2} r^{3/2})$	=	$O(m^{2-x/2+3\alpha/2})$
Product	LR-LR	$O(b r^2)$	$O(p^3)$	$O(p^3 b r^2)$	=	$O(m^{3-2x+2\alpha})$
Recompress	LR	$O(b p r^2)$	$O(p^2)$	$O(p^3 b r^2)$	=	$O(m^{3-2x+2\alpha})$

427 where  $x^*$  defines the optimal choice of the first-level block size  $b$ . We thus obtain a final two-level  
428 storage complexity of

$$429 \quad \mathcal{F}_{ds}^2(m, r) = O(m^{4/3} r^{2/3}). \quad (4.12)$$

430 We now compute the flop complexity  $\mathcal{F}_{ds}^2(m, r)$  of the two-level BLR dense factorization. We  
431 compute the cost of each step and report it in Table 4.1, which is the two-level equivalent of Ta-  
432 ble 2.3; the differences between the two tables are highlighted in gray. The Product and Recom-  
433 press steps have not changed and thus have the same cost. The Factor step consists in factorizing  
434 the  $p$  diagonal blocks which are now represented by BLR matrices; thus its cost is directly derived  
435 from the BLR complexity computed in (2.13):

$$436 \quad \mathcal{F}_{Factor}^2(b, p, r) = p \times \mathcal{F}_{ds}^1(b, r) = O(p b^2 r). \quad (4.13)$$

437 It remains to compute the cost of the Solve step, which now takes the form of  $O(p^2)$  calls to the  
438 BLR-LR-Solve kernel whose cost is given by (4.7). Thus, the overall cost of the Solve step is

$$439 \quad \mathcal{F}_{Solve}^2(b, p, r) = O(p^2 b^{3/2} r^{3/2}). \quad (4.14)$$

440 This concludes the computations for the third, fourth, and fifth columns of Table 4.1. Just as  
441 in the monolevel case, the sixth column is obtained by replacing  $b$ ,  $p$ , and  $r$  by their expression  
442  $O(m^x)$ ,  $O(m^{1-x})$ , and  $O(m^\alpha)$ , respectively.

443 We can finally compute the total flop complexity as the sum of the costs of all steps

$$444 \quad \mathcal{F}_{ds}^2(m, x, \alpha) = O(m^{1+x+\alpha} + m^{2-x/2+3\alpha/2} + m^{3-2x+2\alpha}) \quad (4.15)$$

445 We then compute  $x^*$  which is again equal to the same value as the one that minimizes the storage  
446 complexity,  $x^* = (2 + \alpha)/3$ . Therefore, the final two-level dense flop complexity is

$$447 \quad \mathcal{F}_{ds}^2(m, r) = O(m^{5/3} r^{4/3}). \quad (4.16)$$

448 The two-level BLR format therefore significantly improves the asymptotic storage and flop  
449 complexity compared with the monolevel format. Our analysis shows that the top level block  
450 size should be set to  $b^* = O(m^{x^*}) = O(m^{2/3} r^{1/3})$ . This is an asymptotically larger value than the  
451 monolevel optimal block size computed in (2.9), which translates the fact that by refining the diag-  
452 onal blocks, we can afford to take a larger block size to improve the overall asymptotic complexity.  
453 However, contrarily to hierarchical matrices,  $b^*$  remains asymptotically much lower than  $O(m)$ ;  
454 this makes the format much more flexible for the reasons described in Section 2.5. In short, *we*  
455 *have traded off some of the flexibility of the monolevel format to improve the asymptotic complexity.*

456 This improvement of the dense flop and storage complexities is translated into an improvement  
457 of the sparse complexities. Assuming a rank bound in  $O(1)$ , we quantify this improvement in  
458 Table 4.2. Compared with the monolevel BLR format, the two-level BLR format drops the  $O(\log n)$   
459 factor in the 2D flop and 3D storage complexities, which become linear and thus optimal. The two-  
460 level BLR format can thus achieve, in these two cases, the same  $O(n)$  complexity as the hierarchical  
461 formats while being almost as simple and flexible as flat formats.

462 Finally, the 3D flop complexity remains superlinear but is significantly reduced, from  $O(n^{4/3})$   
463 to  $O(n^{10/9})$ . As the problem size gets larger and larger, even small asymptotic improvements can  
464 make a big difference. Therefore, we now generalize the two-level analysis to the multilevel case  
465 with an arbitrary number of levels.

Table 4.2: Flop and storage complexities of the monolevel and two-level BLR factorizations of a sparse system of  $n = N \times N$  (2D case) or  $n = N \times N \times N$  (3D case) unknowns, derived from the complexities of the factorization of a dense matrix of order  $m$ . We consider a constant rank bound  $r = O(1)$ . The BLR variant considered is CUFS.

	$\mathcal{F}_{ds}(m)$	$\mathcal{F}_{sp}(n)$		$\mathcal{S}_{ds}(m)$	$\mathcal{S}_{sp}(n)$	
		2D	3D		2D	3D
Monolevel BLR	$O(m^2)$	$O(n \log n)$	$O(n^{4/3})$	$O(m^{1.5})$	$O(n)$	$O(n \log n)$
Two-level BLR	$O(m^{5/3})$	$O(n)$	$O(n^{10/9})$	$O(m^{4/3})$	$O(n)$	$O(n)$

466 **5. Generalization to  $\ell$ -level BLR matrices.** In this section, we generalize two-level proof  
467 and computations of the previous section to an arbitrary number of levels  $\ell$  by computing recursive  
468 complexity formulas. We remind that in this section, we are still considering the weakly-admissible  
469 case only.

470 **5.1. Recursive complexity analysis.** Just as for the two-level case, one can compute the  
471 three-level asymptotic complexities, and so on until the general formula becomes clear. We state  
472 the result for an arbitrary number of levels  $\ell$  in the following theorem. Note that it is important to  
473 assume that the number of levels  $\ell$  is constant ( $\ell = O(1)$ ), since the constants hidden in the big  $O$   
474 depend on  $\ell$ .

475 **THEOREM 5.1** (Storage and flop complexity of the  $\ell$ -level BLR factorization). *Let us consider*  
476 *a dense  $\ell$ -level BLR matrix of order  $m$ . We note  $b$  the size of the top level blocks, and  $p = m/b$ . Let*  
477  *$r = O(m^\alpha)$  be the bound on the maximal rank of any block on any level. Then, the optimal choice of*  
478 *the top level block size is  $b = O(m^{x^*})$ , with  $x^* = (\ell + \alpha)/(\ell + 1)$ , which leads to the following storage*  
479 *and flop complexities:*

$$480 \quad \mathcal{S}_{ds}^\ell(m, r) = O(m^{(\ell+2)/(\ell+1)} r^{\ell/(\ell+1)}); \quad (5.1)$$

$$481 \quad \mathcal{F}_{ds}^\ell(m, r) = O(m^{(\ell+3)/(\ell+1)} r^{2\ell/(\ell+1)}). \quad (5.2)$$

483

484 *Proof.* The proof is inductive. We carefully track the constants hidden in the big  $O$  so as to  
485 check their asymptotic dependence on  $\ell$ . We therefore seek to prove the recursive bounds

$$486 \quad \mathcal{S}_{ds}^\ell(m, r) \leq C_\ell m^{(\ell+2)/(\ell+1)} r^{\ell/(\ell+1)},$$

$$487 \quad \mathcal{F}_{ds}^\ell(m, r) \leq C'_\ell m^{(\ell+3)/(\ell+1)} r^{2\ell/(\ell+1)},$$

489 where  $C_\ell$  and  $C'_\ell$  are constants independent of  $m$ . The formulas hold for  $\ell = 1$  with  $C_1 = 1$  and  
490  $C'_1 = 4$ . Let us assume that they are true for the  $(\ell - 1)$ -level BLR factorization and prove that they  
491 still hold for the  $\ell$ -level one.

$\mathcal{S}_{ds}^\ell(m, r)$  can be computed as the storage cost for the off-diagonal low-rank blocks (whose number is less than  $p^2$ ) plus that of the  $p$  diagonal blocks, which are represented as  $(\ell - 1)$ -level BLR matrices. Therefore, by induction,

$$\mathcal{S}_{ds}^\ell(b, p, r) \leq p \times \mathcal{S}_{ds}^{\ell-1}(b, r) + p^2 b r \leq p C_{\ell-1} b^{(\ell+1)/\ell} r^{(\ell-1)/\ell} + p^2 b r.$$

We replace  $b$ ,  $p$ , and  $r$  by their expression  $m^x$ ,  $m^{1-x}$ , and  $m^\alpha$ , respectively, to obtain

$$\mathcal{S}_{ds}^\ell(m, x, \alpha) \leq C_{\ell-1} m^{1-x+(\ell+1)x/\ell+\alpha(\ell-1)/\ell} + m^{2-x+\alpha}.$$

For  $x^* = (\ell + \alpha)/(\ell + 1)$ , we obtain

$$\mathcal{S}_{ds}^\ell(m, r) \leq C_\ell m^{(\ell+2)/(\ell+1)} r^{\ell/(\ell+1)},$$

492 with  $C_\ell = C_{\ell-1} + 1$  (and thus  $C_\ell = \ell$ ).

Table 5.1:  $\ell$ -level equivalent of Tables 2.3 and 4.1. The legend of the table is the same. The differences between this table and the previous two are highlighted in gray.

Step	Type	$cost_{Step}^\ell$	$number_{Step}$	$\mathcal{F}_{Step}^\ell(b, p, r)$	$= \mathcal{F}_{Step}^\ell(m, x, \alpha)$
Factor	BLR $_{(\ell-1)}$	$O(b^{(\ell+2)/\ell} r^{2(\ell-1)/\ell})$	$O(p)$	$O(pb^{(\ell+2)/\ell} r^{2(\ell-1)/\ell})$	$= O(m^{1+2x/\ell+2\alpha(\ell-1)/\ell})$
Solve	BLR $_{(\ell-1)}$ -LR	$O(b^{(\ell+1)/\ell} r^{(2\ell-1)/\ell})$	$O(p^2)$	$O(p^2 b^{(\ell+1)/\ell} r^{(2\ell-1)/\ell})$	$= O(m^{2+(1-\ell)x/\ell+\alpha(2\ell-1)/\ell})$
Product	LR-LR	$O(br^2)$	$O(p^3)$	$O(p^3 br^2)$	$= O(m^{3-2x+2\alpha})$
Recompress	LR	$O(bpr^2)$	$O(p^2)$	$O(p^3 br^2)$	$= O(m^{3-2x+2\alpha})$

We now consider the flop complexity.  $\mathcal{F}_{ds}^\ell(m, r)$  can be computed as the sum of the costs of the Factor, Solve, Product and Recompress steps. We provide the  $\ell$ -level equivalent of Tables 2.3 and 4.1 in Table 5.1. We now consider the flop complexity. The Product and Recompress steps do not depend on  $\ell$  and their cost is less than  $p^3 br^2 \leq m^{3-2x+2\alpha}$ . The Solve step consists in applying less than  $p^2$  times the BLR $_{(\ell-1)}$ -LR-Solve kernel, described in Algorithm 5.1, where BLR $_{(\ell-1)}$  denotes a  $(\ell-1)$ -level BLR matrix (with the convention that BLR $_0$  denotes a FR matrix). It can easily be proven by induction that the cost of applying the BLR $_{(\ell-1)}$ -LR-Solve kernel is less than  $r \times \mathcal{S}_{ds}^{\ell-1}(b, r)$ . Therefore, it holds

$$\mathcal{F}_{Solve}^\ell(p, b, r) \leq p^2 r \times \mathcal{S}_{ds}^{\ell-1}(b, r) \leq C_{\ell-1} p^2 b^{(\ell+1)/\ell} r^{(2\ell-1)/\ell},$$

and thus

$$\mathcal{F}_{Solve}^\ell(m, \alpha, x) \leq C_{\ell-1} m^{2+(1-\ell)x/\ell+\alpha(2\ell-1)/\ell}.$$

The Factor step consists in factorizing  $p$  diagonal  $(\ell-1)$ -level BLR matrices and therefore, by induction,

$$\mathcal{F}_{Factor}^\ell(p, b, r) \leq p C'_{\ell-1} b^{(\ell+2)/\ell} r^{2(\ell-1)/\ell} \leq C'_\ell m^{1+2x/\ell+2\alpha(\ell-1)/\ell}.$$

Finally, summing the cost of all steps and taking  $x^* = (\ell + \alpha)/(\ell + 1)$ , we obtain

$$\mathcal{F}_{ds}^\ell(m, r) \leq C'_\ell m^{(\ell+3)/(\ell+1)} r^{2\ell/(\ell+1)},$$

493 with  $C'_\ell = C'_{\ell-1} + C_{\ell-1} + 2$  (and thus  $C'_\ell = \ell^2 + 3$ ). Since the number of levels  $\ell$  is assumed to be  
 494 constant, we have  $C_\ell = O(C'_\ell) = O(1)$  which concludes the proof.  $\square$

---

**Algorithm 5.1** BLR $_{\ell}$ -LR-Solve kernel,  $\ell > 1$  (upper triangular case)

---

Input: a  $p \times p$  block upper triangular BLR $_{\ell}$  matrix  $\tilde{U}$ ;  $\tilde{U} = [\tilde{U}_{ij}]_{i=1:j,j=1:p}$  such that  $\tilde{U}_{ij} = Y_{ij} X_{ij}^T$  for  $i \neq j$  and  $\tilde{U}_{jj}$  is a BLR $_{(\ell-1)}$  matrix; a LR matrix  $\tilde{B} = \Phi \Psi^T$ .

Output: overwritten  $\Psi$  (modified in-place) corresponding to the operation  $\tilde{B} \leftarrow \tilde{B} \tilde{U}^{-1}$ .

- 1: **for**  $j = 1$  **to**  $p$  **do**
  - 2:  $\Psi_{j,:}^T \leftarrow \Psi_{j,:}^T - \sum_{i=1}^{j-1} \Psi_{i,:}^T Y_{ij} X_{ij}^T$
  - 3: BLR $_{(\ell-1)}$ -LR-Solve  $(\tilde{U}_{jj}, \Phi \Psi_{j,:}^T)$
  - 4: **end for**
- 

495 **5.2. Influence of the number of levels  $\ell$ .** With the formulas from Theorem 5.1 proven, we  
 496 now analyze their practical implications. It is clear that both the storage and flop asymptotic com-  
 497 plexities decrease monotonically with the number of levels, while the top level block size increases.

498 In Figure 5.1, we plot the value of the exponent of the asymptotic complexities as a function of  
 499 the number of levels, for different rank bounds  $r = O(m^\alpha)$ .

500 Let us first consider the case of  $r = O(1)$  (i.e.,  $\alpha = 0$ ). We have already shown that, in this case,  
 501 the two-level BLR factorization has  $O(m^{5/3})$  flop complexity, which leads to  $O(n^{10/9})$  3D sparse  
 502 complexity. With a third level, the dense complexity decreases to  $O(m^{3/2})$ , which is precisely the  
 503 3D sparse breaking point (see Table 2.1), and thus leads to  $O(n \log n)$  complexity. The log factor can



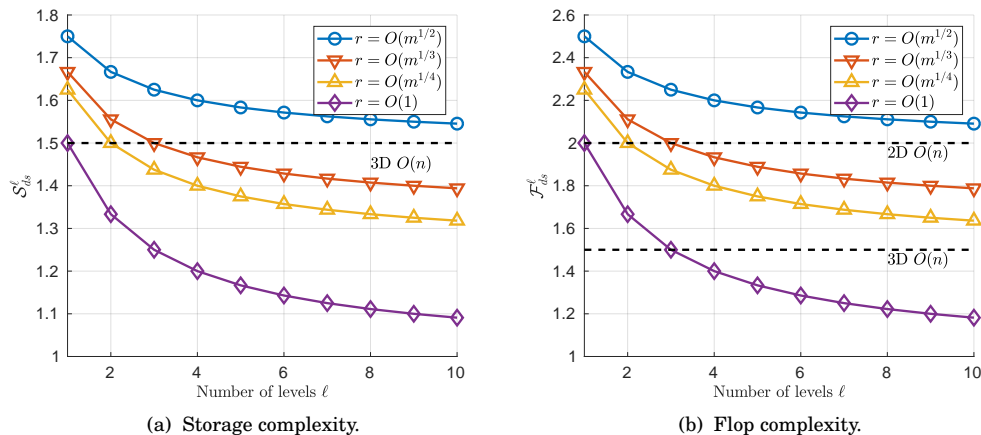


Fig. 5.1: Theoretical asymptotic exponent of the storage and flop complexity of the dense MBLR factorization

504 be dropped by adding a fourth level. The dense complexity tends towards  $O(m)$  as the number of  
 505 levels increases, but the sparse complexity cannot be further improved after the optimal  $O(n)$  has  
 506 been reached. This illustrates that only a small number of levels is necessary to reach low sparse  
 507 complexity. In particular, with four levels, the number of blocks on the top level is  $p = O(m^{1/5})$ ,  
 508 which is still a quite large number which, as previously described, provides more flexibility to  
 509 address issues such as data distribution, parallel implementation, numerical pivoting, etc.

510 The picture is different with higher rank bounds. Indeed, the higher the rank, the more diffi-  
 511 cult it is to reach a low complexity and thus more levels are required. For example, with  $r = O(\sqrt{m})$   
 512 ( $\alpha = 1/2$ ), it is actually not possible to reach a  $O(n)$  3D flop complexity since the dense complexity is  
 513 at best  $O(mr^2) = O(m^2)$ . This is the dense complexity achieved by the hierarchical formats, as well  
 514 as the MBLR format with an infinite number of levels, and leads to  $O(n^{4/3})$  sparse complexity. It is  
 515 therefore not possible to achieve this complexity with a constant number of levels. However, a crucial  
 516 observation is that the rate of improvement of the exponent, which follows  $(\ell + 3 + 2\alpha\ell)/(\ell + 1)$ ,  
 517 is rapidly decreasing as  $\ell$  increases. For example, with  $\alpha = 1/2$ , one level decreases the full-rank  
 518  $O(m^3)$  complexity to the BLR  $O(m^{2.5})$  complexity; it would require an infinite number of levels to  
 519 achieve another  $O(m^{0.5})$  factor of gain. Similarly, adding two more levels leads to  $O(m^{2.25})$ , achiev-  
 520 ing a  $O(m^{0.25})$  gain which can only be achieved again with infinitely more levels! This illustrates  
 521 the critical observation that *the first few levels achieve most of the asymptotic gain*. We therefore  
 522 believe that the MBLR factorization with only a small number of levels can be of practical interest,  
 523 even for problems with larger ranks.

524 **5.3. Comparison with the BLR- $\mathcal{H}$  format.** We conclude this section by comparing our  
 525 MBLR format to the related BLR- $\mathcal{H}$  format, sometimes also referred to as “Lattice- $\mathcal{H}$ ”. It consists  
 526 in representing the matrix using the BLR format, and then approximating its diagonal blocks with  
 527  $\mathcal{H}$ -matrices. The BLR- $\mathcal{H}$  format has been considered as a simple way to use hierarchical matrices  
 528 in a distributed-memory setting [21, 1] but has been little studied from a theoretical standpoint.  
 529 The question is whether refining the diagonal blocks with additional levels (as  $\mathcal{H}$  rather than  
 530 MBLR matrices) improves the asymptotic complexity of the format. In the following, we prove that  
 531 this is not the case and therefore recommend the use of the MBLR format over that of the BLR- $\mathcal{H}$   
 532 one.

533 The complexity of the BLR- $\mathcal{H}$  format is entirely determined by the block size  $b$  used for the  
 534 BLR partitioning. Indeed, the storage complexity can be computed as the sum of the storage for  
 535 the off-diagonal low-rank blocks and that of the diagonal  $\mathcal{H}$  blocks:

536 
$$\mathcal{S}_{ds}^{BLR-\mathcal{H}}(p, b, r) = O(p^2 br) + O(pbr \log b) = O(p^2 br),$$

537 where  $p = m/b$ . Thus, the term corresponding to the off-diagonal low-rank blocks is dominant and  
 538 we obtain

$$539 \quad \mathcal{S}_{ds}^{BLR-\mathcal{H}}(m, b, r) = O\left(\frac{m^2 r}{b}\right). \quad (5.3)$$

540 Similarly, the flop complexity of the BLR- $\mathcal{H}$  factorization is dominated by the LR-LR-Product and  
 541 Recompress steps which cost

$$542 \quad \mathcal{F}_{ds}^{BLR-\mathcal{H}}(m, b, r) = O(p^3 b r^2) = O\left(\frac{m^3 r^2}{b^2}\right). \quad (5.4)$$

543 Applying (5.3) and (5.4) with the optimal choice of block size for the  $\ell$ -level BLR format,  $b =$   
 544  $O(m^{\ell/(\ell+1)} r^{1/(\ell+1)})$ , we obtain the same asymptotic complexity as that proved in Section 5 (Theo-  
 545 rem 5.1):

$$546 \quad \mathcal{S}_{ds}^{BLR-\mathcal{H}}(m, r) = \mathcal{S}_{ds}^{\ell}(m, r) = O(m^{(\ell+2)/(\ell+1)} r^{\ell/(\ell+1)});$$

$$547 \quad \mathcal{F}_{ds}^{BLR-\mathcal{H}}(m, r) = \mathcal{F}_{ds}^{\ell}(m, r) = O(m^{(\ell+3)/(\ell+1)} r^{2\ell/(\ell+1)}).$$

This result can be interpreted as follows: for any given block size  $b$ , there exists a constant number of levels  $\ell_b$  such that representing the diagonal blocks of the matrix as  $\ell_b$ -level BLR matrices suffices to achieve the lowest possible complexity. As far as asymptotic complexity is concerned, it is thus not necessary to represent these diagonal blocks with the  $\mathcal{H}$  format. Note that the value of  $\ell_b$  can easily be computed as

$$\ell_b = \min \left\{ \ell : m^{\ell/(\ell+1)} r^{1/(\ell+1)} \geq b \right\} - 1.$$

549 **6. Numerical experiments.** In this section we compare the experimental complexities of the  
 550 full-rank, BLR, and MBLR formats (with different numbers of levels) for the factorization of dense  
 551 matrices arising from Schur complements of sparse problems.

552 We have developed a MATLAB code to perform the BLR and MBLR factorization of a dense  
 553 matrix, which we use to run all experiments. The objective of this experimental section is purely to  
 554 validate the theoretical complexity bounds that we have computed in Theorem 5.1. Assessing the  
 555 practical performance of the MBLR factorization is complex, not our focus in this paper and will be  
 556 the objective of future work; numerical experiments with sparse matrix factorizations are also out  
 557 of the scope of this article.

558 **6.1. Experimental setting.** All the experiments were performed on the brunch computer  
 559 from the LIP laboratory (ENS Lyon), a shared-memory machine equipped with 24 Haswell proces-  
 560 sors and 1.5 TB of memory.

561 To validate our theoretical complexity results, we use a Poisson problem, which generates the  
 562 symmetric positive definite matrix  $A$  from a 7-point finite-difference discretization of equation

$$563 \quad \Delta u = f$$

564 on a 3D domain of size  $n = N \times N \times N$  with Dirichlet boundary conditions. We compute the dense  
 565 MBLR factorization of the matrices  $F$  corresponding to the root separator of the nested dissection  
 566 partitioning, which are of order  $m = N^2$ .

567 We compute the experimental asymptotic complexities by means of the least-squares estima-  
 568 tion of the coefficients  $\{\beta_i\}_i$  of a regression function  $f$  such that  $X_{fit} = f(N, \{\beta_i\}_i)$  fits the observed  
 569 data  $X_{obs}$ . We use the following regression function:

$$570 \quad X_{fit} = e^{\beta_1^*} N^{\beta_2^*} \text{ with } \beta_1^*, \beta_2^* = \underset{\beta_1, \beta_2}{\operatorname{argmin}} \|\log X_{obs} - \beta_1 - \beta_2 \log N\|^2.$$

571 To compute the low-rank form of the blocks, we perform a truncated QR factorization with  
 572 column pivoting (i.e., a truncated version of LAPACK's [7] `_geqp3` routine). We use a a mix of  
 573 fixed-accuracy and fixed-rank truncation: we stop the factorization after either an accuracy of  $\varepsilon$   
 574 has been achieved or at most  $r_{max}$  columns have been computed. In the following experiments, we  
 575 have set  $\varepsilon = 10^{-14}$  and  $r_{max} = 10$ .

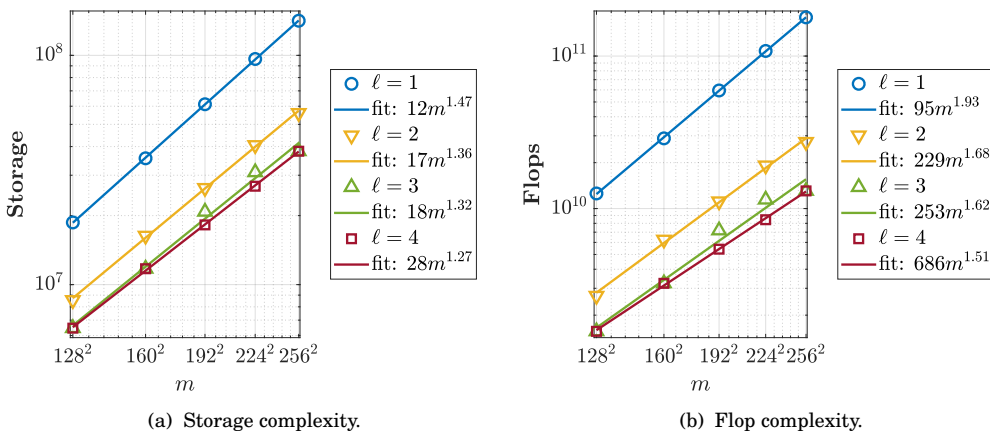


Fig. 6.1: Experimental complexity of the dense MBLR factorization of the root separator of order  $m = N^2$  of a 3D Poisson problem of order  $n = N^3$ .

Table 6.1: Comparison between theoretical and experimental complexities.

	$\ell = 1$	$\ell = 2$	$\ell = 3$	$\ell = 4$
	Storage complexity			
Theoretical	$O(m^{1.50})$	$O(m^{1.33})$	$O(m^{1.25})$	$O(m^{1.20})$
Experimental	$O(m^{1.47})$	$O(m^{1.36})$	$O(m^{1.32})$	$O(m^{1.27})$
	Flop complexity			
Theoretical	$O(m^{2.00})$	$O(m^{1.67})$	$O(m^{1.50})$	$O(m^{1.40})$
Experimental	$O(m^{1.97})$	$O(m^{1.68})$	$O(m^{1.62})$	$O(m^{1.51})$

576 Note that in the weakly admissible context, fixing the rank to 10 yields a limited accuracy  
 577 (measured by the quantity  $\|F - LU\|$ ) of the order of  $10^{-1}$ . While this could be enough to build a  
 578 preconditioner, achieving a higher accuracy would require a strongly admissible implementation.  
 579 The generalization of the MBLR format to the strongly admissible case is therefore of great im-  
 580 portance and, while its practical implementation is outside the scope of this article, we provide its  
 581 algorithmic description in the appendix. We prove that its theoretical complexity is identical to  
 582 that of the weakly admissible case; we also include some very preliminary experiments showing  
 583 that its complexity is in agreement with the theory while leading to a much higher accuracy.

584 **6.2. Experimental complexity results.** We report in Figure 6.1 the experimental storage  
 585 and flop complexity of the dense MBLR factorization, using a number of levels  $\ell$  varying from 1 to  
 586 4. We summarize the asymptotic exponents obtained by fitting the data points in Table 6.1. The  
 587 experimental complexities are in relatively good agreement with the theoretical ones.

588 While there is an almost perfect match for the first two levels, the gap between theory and  
 589 practice increases as more levels are added. This is due to the increasing difficulty of tuning the  
 590 different block sizes at each level. This is a practical limitation that should be further investigated.  
 591 Nevertheless, an asymptotic gain is achieved by each addition of a new level, at least up to  $\ell = 4$ .

592 Overall, these experimental results therefore support the capacity of the MBLR format to  
 593 significantly reduce the asymptotic complexity of the factorization, even when only a small number  
 594 of levels is used.

595 **7. Conclusion.** We have proposed a new multilevel BLR (MBLR) format to bridge the gap be-  
 596 tween flat and hierarchical low-rank matrix formats. Contrarily to hierarchical formats for which

597 the number of levels in the block hierarchy is logarithmically dependent on the size of the problem,  
 598 the MBLR format only uses a constant number of levels  $\ell$ .

599 We had previously explained why the  $\mathcal{H}$ -matrix theory, while applicable to the BLR format,  
 600 leads to very pessimistic complexity bounds and is therefore not suitable. Here, we have shown  
 601 that this remains true for the MBLR format and we therefore extended the theory to compute bet-  
 602 ter bounds. We proved that both the storage and flop complexities of the factorization can be finely  
 603 controlled by  $\ell$ . We theoretically showed that the first few levels achieve most of the asymptotic  
 604 gain that can be expected. In particular, for a sparse 3D problem with constant ranks, two lev-  
 605 els suffice to achieve  $O(n)$  storage complexity and three levels achieve  $O(n \log n)$  flop complexity,  
 606 suggesting that a small number of levels may be enough in practice. Our numerical experiments  
 607 confirm this trend.

608 Having a small number of levels leads to a greater freedom to distribute data in parallel;  
 609 in particular blocks are small enough for several of them to fit in shared-memory, allowing an  
 610 efficient parallelization. Finally, a small number of levels greatly simplifies the implementation  
 611 of the format, making it easy to handle important features such as dynamic data structures and  
 612 numerical pivoting. The related BLR- $\mathcal{H}$  (or Lattice- $\mathcal{H}$ ) format targets a similar objective; however,  
 613 our theoretical analysis shows that using more levels to refine the diagonal blocks actually does  
 614 not improve the asymptotic complexity with respect to the MBLR format.

615 In short, the MBLR format aims to strike a balance between asymptotic complexity and actual  
 616 performance on parallel computers by trading off the optimal hierarchical complexity to retrieve  
 617 some of the simplicity and flexibility of flat monolevel formats. We believe that this increased  
 618 simplicity and flexibility will prove to be useful in a parallel, algebraic, fully-featured, general  
 619 purpose sparse direct solver. The implementation of the MBLR format in such a solver will be the  
 620 object of future work.

621 **Acknowledgements.** We thank Cleve Ashcraft for useful discussions.

## 622 REFERENCES

- 623 [1] *Private communication with R. Kriemann.*
- 624 [2] P. R. AMESTOY, C. ASHCRAFT, O. BOITEAU, A. BUTTARI, J.-Y. L'EXCELLENT, AND C. WEISBECKER, *Improving*  
 625 *multifrontal methods by means of block low-rank representations*, SIAM Journal on Scientific Computing, 37  
 626 (2015), pp. A1451–A1474.
- 627 [3] P. R. AMESTOY, R. BROSSIER, A. BUTTARI, J.-Y. L'EXCELLENT, T. MARY, L. MÉTIVIER, A. MINIUSSI, AND S. OP-  
 628 *ERTO, Fast 3D frequency-domain full waveform inversion with a parallel Block Low-Rank multifrontal direct*  
 629 *solver: application to OBC data from the North Sea*, Geophysics, 81 (2016), pp. R363 – R383.
- 630 [4] P. R. AMESTOY, A. BUTTARI, J.-Y. L'EXCELLENT, AND T. MARY, *On the complexity of the Block Low-Rank mul-*  
 631 *tifrontal factorization*, SIAM Journal on Scientific Computing, 39 (2017), pp. A1710–A1740, [https://doi.org/10.](https://doi.org/10.1137/16M1077192)  
 632 [1137/16M1077192](https://doi.org/10.1137/16M1077192).
- 633 [5] A. AMINFAR, S. AMBIKASARAN, AND E. DARVE, *A fast block low-rank dense solver with applications to finite-element*  
 634 *matrices*, Journal of Computational Physics, 304 (2016), pp. 170–188.
- 635 [6] A. AMINFAR AND E. DARVE, *A fast, memory efficient and robust sparse preconditioner based on a multifrontal ap-*  
 636 *proach with applications to finite-element matrices*, International Journal for Numerical Methods in Engineering,  
 637 107 (2016), pp. 520–540.
- 638 [7] E. ANDERSON, Z. BAI, C. BISCHOF, S. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DUCROZ, A. GREENBAUM,  
 639 S. HAMMARLING, A. MCKENNEY, AND D. SORENSEN, *LAPACK Users' Guide*, SIAM Press, Philadelphia, PA,  
 640 third ed., 1995.
- 641 [8] C. ASHCRAFT, R. G. GRIMES, J. G. LEWIS, B. W. PEYTON, AND H. D. SIMON, *Progress in sparse matrix methods for*  
 642 *large linear systems on vector computers*, Int. Journal of Supercomputer Applications, 1(4) (1987), pp. 10–30.
- 643 [9] M. BEBENDORF, *Hierarchical Matrices: A Means to Efficiently Solve Elliptic Boundary Value Problems*, vol. 63 of  
 644 *Lecture Notes in Computational Science and Engineering (LNCSE)*, Springer-Verlag, 2008.
- 645 [10] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Introduction to hierarchical matrices with applications*, Engineering  
 646 analysis with boundary elements, 27 (2003), pp. 405–422, [https://doi.org/10.1016/S0955-7997\(02\)00152-2](https://doi.org/10.1016/S0955-7997(02)00152-2).
- 647 [11] S. CHANDRASEKARAN, M. GU, AND T. PALS, *A fast ULV decomposition solver for hierarchically semiseparable repre-*  
 648 *sentations*, SIAM Journal on Matrix Analysis and Applications, 28 (2006), pp. 603–622.
- 649 [12] J. W. DEMMEL, S. C. EISENSTAT, J. R. GILBERT, X. S. LI, AND J. W. H. LIU, *A supernodal approach to sparse*  
 650 *partial pivoting*, SIAM Journal on Matrix Analysis and Applications, 20 (1999), pp. 720–755.
- 651 [13] I. S. DUFF AND J. K. REID, *The multifrontal solution of indefinite sparse symmetric linear systems*, ACM Transactions  
 652 on Mathematical Software, 9 (1983), pp. 302–325.
- 653 [14] J. A. GEORGE, *Nested dissection of a regular finite-element mesh*, SIAM Journal on Numerical Analysis, 10 (1973),  
 654 pp. 345–363.

- 655 [15] P. GHYSELS, X. S. LI, C. GORMAN, AND F. H. ROUET, *A robust parallel preconditioner for indefinite systems using*  
656 *hierarchical matrices and randomized sampling*, in 2017 IEEE International Parallel and Distributed Processing  
657 Symposium (IPDPS), May 2017, pp. 897–906, <https://doi.org/10.1109/IPDPS.2017.21>.
- 658 [16] P. GHYSELS, X. S. LI, F.-H. ROUET, S. WILLIAMS, AND A. NAPOV, *An efficient multicore implementation of a novel*  
659 *hss-structured multifrontal solver using randomized sampling*, SIAM Journal on Scientific Computing, 38 (2016),  
660 pp. S358–S384, <https://doi.org/10.1137/15M1010117>.
- 661 [17] A. GILLMAN, P. YOUNG, AND P.-G. MARTINSSON, *A direct solver with  $\mathcal{O}(N)$  complexity for integral equations on*  
662 *one-dimensional domains*, Frontiers of Mathematics in China, 7 (2012), pp. 217–247.
- 663 [18] L. GRASEDYCK AND W. HACKBUSCH, *Construction and Arithmetics of  $\mathcal{H}$ -Matrices*, Computing, 70 (2003), pp. 295–  
664 334, <https://doi.org/10.1007/s00607-003-0019-1>.
- 665 [19] W. HACKBUSCH, *A sparse matrix arithmetic based on  $\mathcal{H}$ -matrices. Part I: introduction to  $\mathcal{H}$ -matrices*, Computing,  
666 62 (1999), pp. 89–108, <https://doi.org/http://dx.doi.org/10.1007/s006070050015>.
- 667 [20] W. HACKBUSCH, *Hierarchical matrices : algorithms and analysis*, vol. 49 of Springer series in computational math-  
668 ematics, Springer, Berlin, 2015, <https://doi.org/10.1007/978-3-662-47324-5>.
- 669 [21] A. IDA, *Efficient Low-rank Solver for Integral Equations on Distributed Memory Systems*, in SIAM Conference on  
670 Parallel Processing (SIAM PP18), Tokyo, Japan, March 2018.
- 671 [22] J. W. H. LIU, *The multifrontal method for sparse matrix solution: Theory and Practice*, SIAM Review, 34 (1992),  
672 pp. 82–109.
- 673 [23] T. MARY, *Block Low-Rank multifrontal solvers: complexity, performance, and scalability*, PhD thesis, Université de  
674 Toulouse, November 2017.
- 675 [24] G. PICHON, E. DARVE, M. FAVERGE, P. RAMET, AND J. ROMAN, *Sparse supernodal solver using block low-rank*  
676 *compression*, in 2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW),  
677 May 2017, pp. 1138–1147, <https://doi.org/10.1109/IPDPSW.2017.86>.
- 678 [25] H. POURANSARI, P. COULIER, AND E. DARVE, *Fast hierarchical solvers for sparse matrices using low-rank approxi-*  
679 *mation*, arXiv preprint arXiv:1510.07363, (2015).
- 680 [26] D. SHANTSEV, P. JAYSAVAL, S. DE LA KETHULLE DE RYHOVE, P. R. AMESTOY, A. BUTTARI, J.-Y. L'EXCELLENT,  
681 AND T. MARY, *Large-scale 3D EM modeling with a Block Low-Rank multifrontal direct solver*, Geophysical  
682 Journal International, 209 (2017), pp. 1558–1571.
- 683 [27] J. XIA, *Efficient structured multifrontal factorization for general large sparse matrices*, SIAM Journal on Scientific  
684 Computing, 35 (2013), pp. A832–A860, <https://doi.org/10.1137/120867032>.
- 685 [28] J. XIA, S. CHANDRASEKARAN, M. GU, AND X. S. LI, *Superfast multifrontal method for large structured linear*  
686 *systems of equations*, SIAM Journal on Matrix Analysis and Applications, 31 (2009), pp. 1382–1411, <https://doi.org/10.1137/09074543X>.
- 687 [29] J. XIA, S. CHANDRASEKARAN, M. GU, AND X. S. LI, *Fast algorithms for hierarchically semiseparable matrices*,  
688 Numerical Linear Algebra with Applications, 17 (2010), pp. 953–976.

690 **Appendix A. Appendix: extension to the strongly-admissible case.** In Sections 4 and 5,  
691 we were considering the weakly-admissible case only. In this section, we extend the proofs and com-  
692 putations to the strongly-admissible case. This requires the introduction of new kernels involving  
693 computations on off-diagonal non-admissible blocks that are represented as MBLR matrices. We  
694 prove that the theoretical complexity results established in the weakly-admissible case still hold  
695 in the strongly-admissible one. The proof involves two main ingredients. First, the fact that the  
696 number of non-admissible blocks per row or column on any level of the block hierarchy can be  
697 bounded by a constant. Second, the computation of the cost of the new kernels, which we show do  
698 not asymptotically increase the overall cost of the factorization.

699 **A.1. New kernels description.** In the strongly admissible case, some off-diagonal blocks  
700 may be non-admissible and therefore represented as MBLR matrices. This introduces five new  
701 kernels: the  $\text{BLR}_\ell$ -LR and  $\text{BLR}_\ell$ - $\text{BLR}_\ell$  additions, the  $\text{BLR}_\ell$ -LR and  $\text{BLR}_\ell$ - $\text{BLR}_\ell$  products, and the  
702  $\text{BLR}_\ell$ - $\text{BLR}_\ell$  triangular solve. These are described in Algorithms A.1 through A.5.

703 Note that the  $\text{BLR}_\ell$ -LR (addition or product) algorithms can easily be adapted to define similar  
704 LR- $\text{BLR}_\ell$  algorithms. The  $\text{BLR}_\ell$ - $\text{BLR}_\ell$  triangular solve (Algorithm A.5) is described in its upper  
705 triangular version; the lower triangular version is similar and is omitted for the sake of concise-  
706 ness.

707 To avoid a profusion of different test cases, we use the compact notation  $\tilde{A} \otimes \tilde{B}$  and  $\tilde{A} \oplus \tilde{B}$  to  
708 denote the product and addition of two matrices  $\tilde{A}$  and  $\tilde{B}$  that may be either LR or MBLR. If at  
709 least one of the two is MBLR, then  $\oplus$  corresponds to a recursive call to Algorithms A.1 or A.2, while  
710  $\otimes$  corresponds to a recursive call to Algorithms A.3 or A.4. If both  $\tilde{A}$  and  $\tilde{B}$  are LR, then  $\otimes$  simply  
711 corresponds to a regular LR-LR product, while  $\oplus$  corresponds to accumulating the two LR matrices  
712 together and recompressing the result.

713 To prove that the strong admissibility case does not increase the asymptotic complexity of the  
714  $\ell$ -level BLR factorization, we thus need to compute the cost of these new kernels. To do so, two

**Algorithm A.1** BLR $_{\ell}$ -LR-Addition kernel

---

Input: a  $p \times p$  BLR $_{\ell}$  matrix  $\tilde{A}$  and a LR matrix  $\tilde{B} = \Phi\Psi^T$ .
Output: a BLR $_{\ell}$  matrix  $\tilde{C} = \tilde{A} + \tilde{B}$ 

```

1: for  $i = 1$  to  $p$  do
2:   for  $j = 1$  to  $p$  do
3:     if  $\tilde{A}_{ij}$  is LR then
4:        $\tilde{C}_{ij} \leftarrow \text{Recompress}(\tilde{A}_{ij} + \Phi_i\Psi_j^T)$ 
5:     else
6:        $\tilde{C}_{ij} \leftarrow \text{BLR}_{(\ell-1)\text{-LR-Addition}}(\tilde{A}_{ij}, \Phi_i\Psi_j^T)$ 
7:     end if
8:   end for
9: end for

```

---

**Algorithm A.2** BLR $_{\ell}$ -BLR $_{\ell}$ -Addition kernel

---

Input: two  $p \times p$  BLR $_{\ell}$  matrices  $\tilde{A}$  and  $\tilde{B}$ .
Output: a  $p \times p$  BLR $_{\ell}$  matrix  $\tilde{C} = \tilde{A} + \tilde{B}$ .

```

1: for  $i = 1$  to  $p$  do
2:   for  $j = 1$  to  $p$  do
3:      $\tilde{C}_{ij} \leftarrow \tilde{A}_{ij} \oplus \tilde{B}_{ij}$ 
4:   end for
5: end for

```

---

715 ingredients are necessary. First, we prove in Section A.2 that the number of non-admissible blocks  
716 at any level of the block hierarchy is bounded by a constant (Lemma A.1). Then, in Section A.3, we  
717 compute their cost by induction (Lemma A.2).

718 **A.2. Boundedness of the number of non-admissible blocks.** In this section, we seek to  
719 compute a bound on  $N_{na}$ , the maximal number of non-admissible blocks per row and column at any  
720 level of the block-hierarchy.

721 The result is directly derived from Lemma 2 in [4], whose result on BLR matrices we reproduce  
722 here.

723 **LEMMA A.1.** *For any BLR matrix, it is possible to build a partition such that the number of*  
724 *non-admissible blocks per row and column is bounded by  $N_{na} = O(1)$ .*

725 *Proof:* See [4], Lemma 2. □

726 Since this result holds for any BLR matrix, it can be recursively applied to the non-admissible  
727 blocks of a  $\ell$ -level matrix, and therefore the result trivially holds at any level of the block hierarchy.

728 As a consequence, the  $\ell$ -level BLR factorization of a dense matrix clustered into  $p = m/b$  blocks  
729 requires  $O(p^2)$  BLR $_{(\ell-1)}$ -LR products and additions and  $O(p)$  BLR $_{(\ell-1)}$ -BLR $_{(\ell-1)}$  products, addi-  
730 tions, and triangular solves to be performed.

731 Now that we have bounded the number of times these kernels are performed, let us bound the  
732 cost of performing them once.

**A.3. Recursive complexity analysis of the new kernels.**

734 **LEMMA A.2.** *The costs of Algorithms A.1 through A.5 are*

$$735 \quad \mathcal{F}_{A.1}^{\ell}(m, r) = \mathcal{F}_{A.2}^{\ell}(m, r) = \mathcal{F}_{A.3}^{\ell}(m, r) = O(r) \times \mathcal{S}_{ds}^{\ell}(m, r), \quad (\text{A.1})$$

$$736 \quad \mathcal{F}_{A.4}^{\ell}(m, r) = \mathcal{F}_{A.5}^{\ell}(m, r) = \mathcal{F}_{ds}^{\ell}(m, r), \quad (\text{A.2})$$

738 where  $\mathcal{F}_{ds}^{\ell}(m, r)$  and  $\mathcal{S}_{ds}^{\ell}(m, r)$  are given by Theorem 5.1.

739 *Proof.* We proceed by induction. Let us note  $b = O(m^x)$  the top level block size of the matrices  
740 and  $p = m/b = O(m^{1-x})$ . Let us also assume that  $r = O(m^{\alpha})$ .

**Algorithm A.3** BLR $_{\ell}$ -LR-Product kernel

---

Input: a  $p \times p$  BLR $_{\ell}$  matrix  $\tilde{A}$  and a LR matrix  $\tilde{C} = \Phi\Psi^T$ .
Output: a LR matrix  $\tilde{C}^{out} = \Phi^{out}\Psi^T = \tilde{A}\tilde{C}$ 

- 
- 1: **for**  $i = 1$  **to**  $p$  **do**
  - 2:    $\Phi_{i,:}^{out}\Psi^T \leftarrow$  BLR $_{(\ell-1)}$ -LR-Product( $\tilde{A}_{ii}, \Phi_{i,:}, \Psi^T$ )
  - 3:    $\Phi_{i,:}^{out} \leftarrow \Phi_{i,:}^{out} + \sum_{j=1; j \neq i}^p X_{ij}Y_{ij}^T\Phi_{j,:}$
  - 4: **end for**
- 

**Algorithm A.4** BLR $_{\ell}$ -BLR $_{\ell}$ -Product kernel

---

Input: two  $p \times p$  BLR $_{\ell}$  matrices  $\tilde{A}$  and  $\tilde{B}$ .
Output: a  $p \times p$  BLR $_{\ell}$  matrix  $\tilde{C} = \tilde{A}\tilde{B}$ .

- 
- 1: **for**  $i = 1$  **to**  $p$  **do**
  - 2:   **for**  $j = 1$  **to**  $p$  **do**
  - 3:      $\tilde{C}_{ij} \leftarrow []$  (rank-0 matrix)
  - 4:     **for**  $k = 1$  **to**  $p$  **do**
  - 5:        $\tilde{C}_{ij} \leftarrow \tilde{C}_{ij} \oplus \tilde{A}_{ik} \otimes \tilde{B}_{kj}$
  - 6:     **end for**
  - 7:   **end for**
  - 8: **end for**
- 

741 We begin with the initial case  $\ell = 1$ . Algorithm A.1 requires  $O(p^2)$  LR-LR additions of cost  
742  $O(br^2)$  and  $O(p)$  LR-FR additions of cost  $O(b^2r)$ . Algorithm A.2 requires the same operations,  
743 plus  $O(p)$  FR-FR additions of cost  $O(b^2)$ . Algorithm A.3 requires  $O(p^2)$  LR-LR products of cost  
744  $O(br^2)$  and  $O(p)$  FR-LR products of cost  $O(b^2r)$ . Therefore, all three algorithms have a total cost of  
745  $O(p^2br^2 + pb^2r) = r \times \mathcal{S}_{ds}^1(b, p, r)$ . Algorithm A.4 requires  $O(p^3)$  LR-LR additions and products of  
746 cost  $O(br^2)$ ,  $O(p^2)$  LR-FR additions and products of cost  $O(b^2r)$ , and  $O(p)$  FR-FR products of cost  
747  $O(b^3)$ . Overall, the total cost of Algorithm A.4 is therefore  $O(p^3br^2 + p^2b^2r + pb^3) = \mathcal{F}_{ds}^1(b, p, r)$ .  
748 Algorithm A.5 requires the same computations and therefore has the same asymptotic cost.

749 We now assume that the formulas hold for  $(\ell - 1)$ -level BLR matrices and prove them for  $\ell$ -  
750 level ones. The previous analysis in the case  $\ell = 1$  still holds by replacing every “FR” occurrence  
751 by “BLR $_{(\ell-1)}$ ”. Thus, Algorithms A.1, A.2, and A.3 require  $O(p^2)$  LR-LR additions and prod-  
752 ucts,  $O(p)$  BLR $_{(\ell-1)}$ -LR additions and products, and  $O(p)$  BLR $_{(\ell-1)}$ -BLR $_{(\ell-1)}$  additions. By in-  
753 duction, their cost is therefore  $O(p^2br^2) + O(p) \times \mathcal{S}_{ds}^{\ell-1}(b, r)$ , which is equal to  $r \times \mathcal{S}_{ds}^{\ell}(m, r)$  for  
754  $x = x^* = (\ell + \alpha)/(\ell + 1)$ . Similarly, Algorithm A.4 requires  $O(p^3)$  LR-LR additions and products,  $O(p^2)$   
755 BLR $_{(\ell-1)}$ -LR additions and products and  $O(p)$  BLR $_{(\ell-1)}$ -BLR $_{(\ell-1)}$  products. By induction, its cost is  
756 therefore  $O(p^3br^2) + O(p^2r) \times \mathcal{S}_{ds}^{\ell-1}(b, r) + O(p) \times \mathcal{F}_{ds}^{\ell-1}(b, r)$ , which is equal to  $\mathcal{F}_{ds}^{\ell}(m, r)$  for  $x = x^*$ .  
757 Algorithm A.5 requires the same computations and therefore has the same asymptotic cost.  $\square$

758 **THEOREM A.3.** *The complexity formulas of Theorem 5.1 established in the weakly-admissible*  
759 *case also hold in the strongly-admissible case.*

760 *Proof.* Lemma A.1 states that there are only  $O(1)$  non-admissible blocks per row and column  
761 on any level of the block hierarchy. From this we can already deduce that the storage complexity  
762  $\mathcal{S}_{ds}^{\ell}$  is asymptotically the same in the weakly- and strongly-admissible cases. Moreover, using the  
763 costs computed in Lemma A.2, we can compute the overall cost associated with the new kernels  
764 required to be performed in the strongly-admissible case as

$$765 \quad \mathcal{F}_{New\ Kernels}^{\ell}(b, p, r) = O(p) \times \mathcal{F}_{ds}^{\ell-1}(b, r) + O(p^2r) \times \mathcal{S}_{ds}^{\ell-1}(b, r) \quad (\text{A.3})$$

$$766 \quad = \mathcal{F}_{Factor}^{\ell}(b, p, r) + \mathcal{F}_{Solve}^{\ell}(b, p, r) \quad (\text{A.4})$$

768 which clearly does not asymptotically increase the total, thus proving that  $\mathcal{F}_{ds}^{\ell}$  is also unchanged  
769 by the strong admissibility condition.  $\square$

**Algorithm A.5** BLR $_{\ell}$ -BLR $_{\ell}$ -Solve kernel (upper triangular case)Input: two  $p \times p$  BLR $_{\ell}$  matrices  $\tilde{U}$  and  $\tilde{A}$ ;  $\tilde{U}$  is upper triangular.Output: overwritten  $\tilde{A}$  (modified in-place) corresponding to the operation  $\tilde{A} \leftarrow \tilde{A}\tilde{U}^{-1}$ .

```

1: for  $k = 1$  to  $p$  do
2:   for  $i = 1$  to  $p$  do
3:     if  $\tilde{A}_{ik}$  is LR then
4:        $\tilde{A}_{ik} \leftarrow \text{BLR}_{(\ell-1)\text{-LR-Solve}}(\tilde{U}_{kk}, \tilde{A}_{ik})$ 
5:     else
6:        $\tilde{A}_{ik} \leftarrow \text{BLR}_{(\ell-1)\text{-BLR}_{(\ell-1)\text{-Solve}}(\tilde{U}_{kk}, \tilde{A}_{ik})$ 
7:     end if
8:     for  $j = k + 1$  to  $p$  do
9:        $\tilde{A}_{ij} \leftarrow \tilde{A}_{ij} \oplus (-\tilde{A}_{ik} \otimes \tilde{U}_{kj})$ 
10:    end for
11:  end for
12: end for

```

770 **A.4. Preliminary storage complexity experiments.** The practical implementation of the  
771 kernels presented in this appendix are outside our scope and therefore, so is the experimental study  
772 of the flop complexity of the strongly admissible MBLR factorization. However, we can easily esti-  
773 mate its storage complexity by first factorizing the matrix with a weakly admissible partitioning,  
774 and then compressing it with a strongly admissible partitioning. Preliminary results are reported  
775 in Figure A.1, where the blocks are compressed at accuracy  $\varepsilon = 10^{-10}$  and with no limit on the  
776 rank  $r$ . This leads to an error  $\|F - LU\|$  of the order of  $10^{-9}$ , regardless of the number of levels  $\ell$ .  
777 As shown in the figure, the asymptotic storage complexity results lead to the same conclusions as  
778 those with the weak admissibility.




Fig. A.1: Storage complexity of the strongly admissible MBLR format, on the root separator of size  $m = N^2$  of a Poisson problem of size  $N^3$ , and with  $\varepsilon = 10^{-10}$ .