

On the Parallel I/O Optimality of Linear Algebra Kernels: Near-Optimal Matrix Factorizations

Grzegorz Kwasniewski, Marko Kabic, Tal Ben-Nun, Alexandros Nikolaos Ziogas, Jens Eirik Saethre, André Gaillard, Timo Schneider, Maciej Besta, Anton Kozhevnikov, Joost VandeVondele, Torsten Hoefer
 Department of Computer Science, ETH Zurich
 Swiss National Computing Center
 Switzerland

ABSTRACT

Matrix factorizations are among the most important building blocks of scientific computing. However, state-of-the-art libraries are not communication-optimal, underutilizing current parallel architectures. We present novel algorithms for Cholesky and LU factorizations that utilize an asymptotically communication-optimal 2.5D decomposition. We first establish a theoretical framework for deriving parallel I/O lower bounds for linear algebra kernels, and then utilize its insights to derive Cholesky and LU schedules, both communicating $N^3/(P\sqrt{M})$ elements per processor, where M is the local memory size. The empirical results match our theoretical analysis: our implementations communicate significantly less than Intel MKL, SLATE, and the asymptotically communication-optimal CANDMC and CAPITAL libraries. Our code outperforms these state-of-the-art libraries in almost all tested scenarios, with matrix sizes ranging from 2,048 to 524,288 on up to 512 CPU nodes of the Piz Daint supercomputer, decreasing the time-to-solution by up to three times. Our code is ScaLAPACK-compatible and available as an open-source library.

CCS CONCEPTS

• **Mathematics of computing** → **Mathematical software performance**; **Computations on matrices**; • **Theory of computation** → **Communication complexity**; *Distributed computing models*; • **Computing methodologies** → **Linear algebra algorithms**;

KEYWORDS

Distributed linear algebra algorithms, communication complexity, matrix factorization

ACM Reference Format:

Grzegorz Kwasniewski, Marko Kabic, Tal Ben-Nun, Alexandros Nikolaos Ziogas, Jens Eirik Saethre, André Gaillard, Timo Schneider, Maciej Besta, Anton Kozhevnikov, Joost VandeVondele, Torsten Hoefer. 2021. On the Parallel I/O Optimality of Linear Algebra Kernels: Near-Optimal Matrix Factorizations. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '21)*, November 14–19, 2021,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
 SC '21, November 14–19, 2021, St. Louis, MO, USA
 © 2021 Association for Computing Machinery.
 ACM ISBN 978-1-4503-8442-1/21/11...\$15.00
<https://doi.org/10.1145/3458817.3476167>

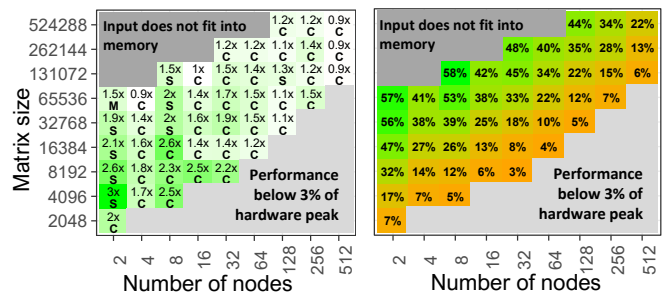


Figure 1: Left: measured runtime speedup of COntrol LUX vs. fastest state-of-the-art library (S=SLATE [28], C=CANDMC [57], M=MKL [34]). Right: COntrol LUX’s achieved % of machine peak performance.

St. Louis, MO, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3458817.3476167>

1 INTRODUCTION

Matrix factorizations, such as LU and Cholesky decompositions, play a crucial role in many scientific computations [42, 49, 65], and their performance can dominate the overall runtime of entire applications [19]. Therefore, accelerating these routines is of great significance for numerous domains [18, 44]. The ubiquity and importance of LU factorization is even reflected by the fact that it is used to rank top supercomputers worldwide [25].

Since the arithmetic complexity of matrix factorizations is $O(N^3)$ while the input size is $O(N^2)$, these kernels are traditionally considered compute-bound. However, the end of Dennard scaling [22] puts increasing pressure on data movement minimization, as the cost of moving data far exceeds its computation cost, both in terms of power and time [40, 63]. Thus, deriving algorithmic I/O lower bounds is a subject of both theoretical analysis [15, 36, 37] and practical value for developing I/O-efficient schedules [33, 59, 60].

While asymptotically optimal matrix factorizations were proposed, among others, by Ballard et al. [7] and Solomonik et al. [33, 61], we observe two major challenges with the existing approaches: First, the presented algorithms are only asymptotically optimal: the I/O cost of these proposed parallel algorithms can be as high as 7 times the lower bound for LU [61] and up to 16 times for Cholesky [33]. This means that they communicate less than “standard” 2D algorithms like ScaLAPACK [14] only for almost prohibitively large numbers of processors — e.g., according to the LU cost

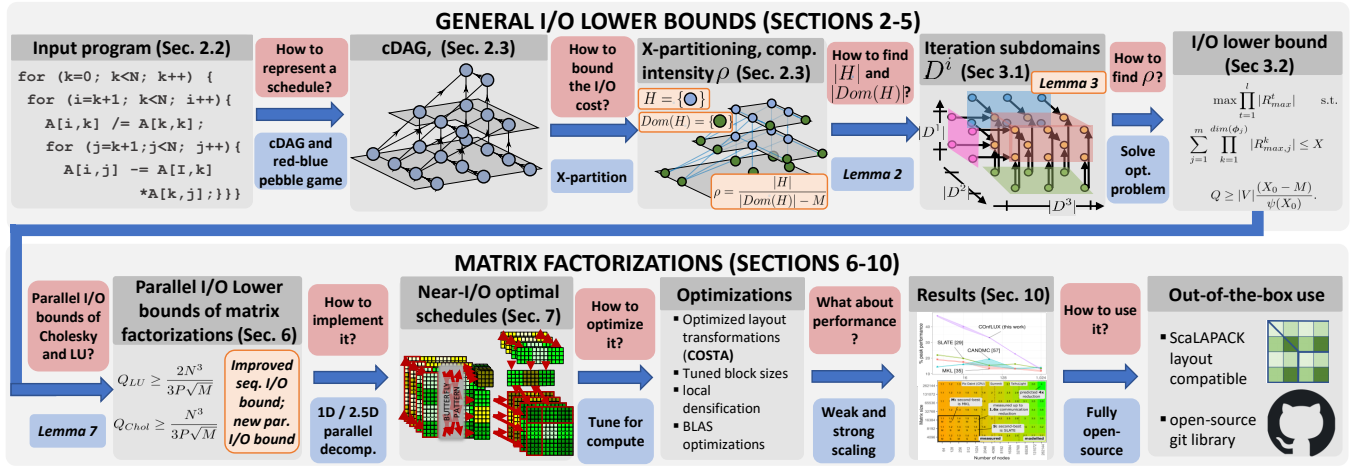


Figure 2: From the input program through the I/O lower bounds to communication-minimizing parallel schedules and high performing implementations. In this paper, we mainly focus on the Cholesky and LU factorizations. The proofs of the lemmas presented in this work can be found in the extended version of this paper (<https://arxiv.org/abs/2108.09337>).

model [61], it requires more than 15,000 processors to communicate less than an optimized 2D algorithm. Second, their time-to-solution performance can be worse than highly-optimized, existing 2D-parallel libraries [33].

To tackle these challenges, we first provide a *general* method for deriving *precise* I/O lower bounds of Disjoint Array Access Programs (DAAP) — a broad range of programs composed of a sequence of statements enclosed in an arbitrary number of nested loops. We then illustrate the applicability of our framework to derive parallel I/O lower bounds of Cholesky and LU factorizations: $\frac{1}{3} \frac{N^3}{P\sqrt{M}}$ and $\frac{2}{3} \frac{N^3}{P\sqrt{M}}$ elements, respectively, where N is the matrix size, P is the number of processors, and M is the local memory size.

Moreover, we use the insights from deriving the above lower bounds to develop *ConfLUX* and *ConfCHOX*, near communication-optimal parallel LU and Cholesky factorization algorithms that minimize data movement across the 2.5D processor decomposition. For LU factorization, to further reduce the latency and bandwidth cost, we use a row-masking tournament pivoting strategy resulting in a communication requirement of $\frac{N^3}{P\sqrt{M}} + O(\frac{N^2}{P})$ elements per processor, where the leading order term is only 1.5 times the lower bound. Furthermore, to secure high performance, we carefully tune block sizes and communication routines to maximize the efficiency of local computations such as *trsm* (triangular solve) and *gemm* (matrix multiplication).

We measure both communication volume and achieved performance of *ConfLUX* and *ConfCHOX* and compare them to state-of-the-art libraries: a vendor-optimized Intel MKL [34], SLATE [28] (a recent library targeting exascale systems), as well as CANDMC [57, 58] and CAPITAL [32, 33] (codes based on the asymptotically optimal 2.5D decomposition). In our experiments on the Piz Daint supercomputer, we measure up to 1.6x communication reduction compared to the second-best implementation. Furthermore, our 2.5D decomposition communicates asymptotically less than SLATE and MKL, with even greater expected benefits on exascale machines.

Compared to the communication-avoiding CANDMC library with I/O cost of $5N^3 / (P\sqrt{M})$ elements [61], *ConfLUX* communicates five times less. Most importantly, *our implementations outperform all compared libraries in almost all scenarios*, both for strong and weak scaling, reducing the time-to-solution by up to three times compared to the second best performing library (Figure 1).

In this work, we make the following contributions:

- A general method for deriving parallel I/O lower bounds of a broad range of linear algebra kernels.
- *ConfLUX* and *ConfCHOX*, provably near-I/O-optimal parallel algorithms for LU and Cholesky factorizations, with their full communication volume analysis.
- Open-source and fully ScaLAPACK-compatible implementations of our algorithms that outperform existing state-of-the-art libraries in almost all scenarios.

A bird’s eye view of our work is presented in Figure 2.

2 BACKGROUND

We now establish the background for our theoretical model (Sections 3-5). We use it to derive parallel I/O lower bounds for Cholesky and LU factorizations (Section 6) that will guide the design of our communication-minimizing implementations (Section 7).

2.1 Machine Model

To model algorithmic I/O complexity, we start with a model of a sequential machine equipped with a two-level deep memory hierarchy. We then outline the parallel machine model.

Sequential machine. A computation is performed on a sequential machine with a fast memory of limited size and unlimited slow memory. The fast memory can hold up to M elements at any given time. To perform any computation, all input elements must reside in fast memory, and the result is stored in fast memory.

Parallel machine. The sequential model is extended to a machine with P processors, each equipped with a private fast memory of size

M . There is no global memory of unlimited size — instead, elements are transferred between processors' fast memories.

2.2 Input Programs

We consider a general class of programs that operate on multidimensional arrays. Array elements can be loaded from slow to fast memory, stored from fast to slow memory, and computed inside fast memory. These elements have *versions* that are incremented every time they are updated. We model the program execution as a computational directed acyclic graph (cDAG, details in Section 2.3), where each vertex corresponds to a different version of an array element. Thus, for a statement $A[i, j] \leftarrow f(A[i, j])$, a vertex corresponding to $A[i, j]$ after applying f is different from a vertex corresponding to $A[i, j]$ before applying f . In a cDAG, this is expressed as an edge from vertex $A[i, j]$ before f to vertex $A[i, j]$ after f . Initial versions of each element do not have any incoming edges and thus form the cDAG inputs. *The distinction between elements and vertices* is important for our I/O lower bounds analysis, as we will investigate how many vertices are computed for a given number of loaded vertices.

An input program is a collection of statements S enclosed in loop nests, each of the following form (we use the loop nest notation introduced by Dinh and Demmel [23]):

$$\text{for } \psi^1 \in \mathcal{D}^1, \text{ for } \psi^2 \in \mathcal{D}^2(\psi^1), \dots, \text{ for } \psi^l \in \mathcal{D}^l(\psi^1, \dots, \psi^{l-1}) : \\ S : A_0[\phi_0(\psi)] \leftarrow f(A_1[\phi_1(\psi)], A_2[\phi_2(\psi)], \dots, A_m[\phi_m(\psi)]),$$

where (cf. Figure 3 for a summary) for each innermost loop iteration, statement S is an evaluation of some function f on m inputs, where every input is an element of array $A_j, j = 1, \dots, m$, and the result of f is stored to the output array A_0 .

Each loop has an associated *iteration variable* ψ^t that iterates over its domain $\mathcal{D}^t \in \mathcal{D}^t$. All l iteration variables form the *iteration vector* $\psi = [\psi^1, \dots, \psi^l]$. Array elements are accessed by an *access function vector* $\phi_j = [\phi_j^1, \dots, \phi_j^{dim(A_j)}]$ that maps $dim(A_j)$ iteration variables to a *unique* element in array A_j (note that the access function vector is injective). Only vertices associated with the newest element versions can be referenced. Furthermore, a given vertex can be referenced by only one access function vector per statement. We refer to this as the *disjoint access property*. The *access dimension* of $A_j(\phi_j)$, denoted $dim(A_j(\phi_j))$, is the number of different iteration variables present in ϕ_j . We call such programs Disjoint Access Array Programs.

Example: Consider statement $S1$ of LU factorization (Figure 3). The loop nest depth is $l = 2$, with two iteration variables $\psi^1 = k$ and $\psi^2 = i$ forming the iteration vector $\psi = [k, i]$. For access $A[k, k]$, the access function vector $\phi_j = [k, k]$ is a function of only one iteration variable k . Therefore, $dim(A_j) = 2$, but $dim(A_j(\phi_j)) = 1$.

2.3 I/O Complexity and Pebble Games

We now establish the relationship between DAAP and the red-blue pebble game - a powerful abstraction for deriving lower bounds and optimal schedules of cDAG evaluation.

2.3.1 cDAG and red-blue pebble game. We base our computation model on the red-blue pebble game, played on the computational directed acyclic graph $G = (V, E)$, as introduced by Hong and Kung [37]. Every vertex $v \in V$ represents the result of a unique

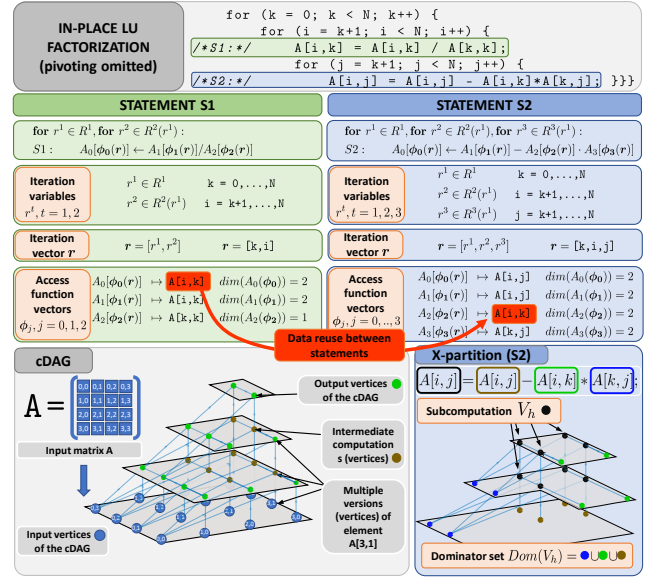


Figure 3: In-place LU factorization (for simplicity, no pivoting is performed). The algorithm contains two statements (S1 and S2), for which we provide key components of our program representation together with the corresponding cDAG for $N = 4$. For statement S2, we also provide a graphical visualization of a single subcomputation H in its X -partition.

computation stored in some memory, and a directed edge $(u, v) \in E$ represents a data dependency. Vertices without any incoming (outgoing) edges are called *inputs* (*outputs*). To perform a computation, i.e., to evaluate the value corresponding to vertex v , all vertices that are direct predecessors of v must be loaded into fast memory. The vertices that are currently in fast memory are marked by a red pebble on the corresponding vertex of the cDAG. Since the size of fast memory is limited, we can never have more than M red pebbles on the cDAG at any moment. Analogously, the contents of the slow memory (of unlimited size) is represented by an unlimited number of blue pebbles.

2.3.2 Dominator and Minimum Sets [37]. For any subset of vertices $H \subset V$, a *dominator set* $Dom(H)$ is a set such that every path in the cDAG from an input vertex to any vertex in H must contain at least one vertex in $Dom(H)$. In general, for a given H , its $Dom(H)$ is not uniquely defined. The *minimum set* $Min(H)$ is the set of all vertices in H that do not have any immediate successors in H . In this work, to avoid the ambiguity of non-uniqueness of dominator set size (in principle, for any subset, its valid dominator set is always the whole V), we will refer to $Dom_{min}(H)$ as a minimum dominator set, i.e. a dominator set with the smallest size.

Intuition. One can think of H 's dominator set as a set of inputs required to execute subcomputation H , and of H 's minimum set as the output of H . We use the notions of $Dom_{min}(H)$ and $Min(H)$ when proving I/O lower bounds. Intuitively, we bound computation "volume" (number of vertices in H) by its communication "surface", comprised of its inputs - vertices in $Dom_{min}(H)$ - and outputs - vertices in $Min(H)$.

2.3.3 X-Partitioning. Introduced by Kwasniewski et al. [45], X -Partitioning generalizes the S -partitioning abstraction [37]. An

X -partition of a cDAG is a collection of s mutually disjoint subsets (referred to as *subcomputations*) $\mathcal{P}(X) = \{H_1, \dots, H_s\}$, $\bigcup_{i=1}^s H_i = V$ with two additional properties:

- $\mathcal{P}(X)$ has no cyclic dependencies between subcomputations.
- $\forall H, |Dom_{min}(H)| \leq X$ and $|Min(H)| \leq X$.

For a given cDAG and for any given $X > M$, let $\Pi(X)$ denote a set of all its valid X -partitions, $\mathcal{P}(X) \in \Pi(X)$. Kwasniewski et al. prove that an I/O optimal schedule of G that performs Q load and store operations has an associated X -partition $\mathcal{P}_{opt}(X) \in \Pi(X)$ with size $|\mathcal{P}_{opt}(X)| \leq \frac{Q+X-M}{X-M}$ for any $X > M$ ([45], Lemma 2).

2.3.4 Deriving lower bounds. To bound the I/O cost, we further need to introduce the *computational intensity* ρ . For each subcomputation H_i , ρ_i is defined as a ratio of the number of computations (vertices) in H_i to the number of I/O operations required to pebble H_i , where the latter is bounded by the size of the dominator set $Dom(H_i)$ [45]. Then, the following lemma bounds the number of I/O operations required to pebble a given cDAG:

Lemma 1. (Lemma 4 in [45]) *For any constant X_c , the number of I/O operations Q required to pebble a cDAG $G = (V, E)$ with $|V| = n$ vertices using M red pebbles is bounded by $Q \geq n/\rho$, where $\rho = \frac{H_{max}}{X_c - M}$ is the maximal computational intensity and $H_{max} = \arg \max_{H \in \mathcal{P}(X_c)} |H|$ is the largest subcomputation among all valid X_c -partitions.*

3 GENERAL SEQUENTIAL I/O LOWER BOUNDS

We now present our method for deriving the I/O lower bounds of a sequential execution of programs in the form defined in Section 2.2. Specifically, in Section 3.2 we derive I/O bounds for programs that contain only a single statement. In Section 4 we extend our analysis to capture interactions and reuse between multiple statements.

In this paper, we present only the key lemmas required to establish the lower bounds of parallel Cholesky and LU factorizations. However, the method covers a much wider spectrum of algorithms. For curious readers, we present all proofs of provided lemmas in the extended version of this paper¹.

We start by stating our key lemma:

Lemma 2. *If $|H_{max}|$ can be expressed as a closed-form function of X , that is if there exists some function χ such that $|H_{max}| = \chi(X)$, then the lower bound on Q can be expressed as*

$$Q \geq n \frac{(X_0 - M)}{\chi(X_0)},$$

where $X_0 = \arg \min_X \rho = \arg \min_X \frac{\chi(X)}{X-M}$.

To find $\chi(X)$, we take advantage of the DAAP structure. Observe that every computation (and therefore, every compute vertex $v \in V$ in the cDAG $G = (V, E)$) is executed in a different iteration of the loop nest, and thus, there is a one-to-one mapping from a value of the iteration vector ψ to the compute vertex v . Moreover, each vertex accessed from any of the input arrays A_i is also associated with some iteration vector value - however, if $dim(A_i) < l$, this is a one-to-many relation, as the same input vertex may be used to evaluate multiple compute vertices v . This is, in fact, the source of the data reuse, and exploiting this relation is a key to minimizing the I/O cost. If for all input arrays A_i we have that $dim(A_i) = l$, then




Figure 4: Lemma 3 bounds the set sizes (both the subcomputation's H and input access sets' $A_j(D)$) with the number of values $|D^t|$ each iteration variable ψ^t takes during the subcomputation.

for each compute vertex v , m different, unique input vertices are required, there is no data reuse and it implies a trivial computational intensity $\rho = \frac{1}{m}$.

The high-level idea of our method is to *count how many different iteration vector values ϕ can be formed if we know how many different values each iteration variable ϕ^1, \dots, ϕ^l takes*. We now formalize this in Lemmas 3-6.

3.1 Iteration vector, iteration domain, access set

Each execution of statement S is associated with the *iteration vector* value $\psi = [\psi^1, \dots, \psi^l] \in \mathbb{N}^l$ representing the current iteration, that is, the values of iteration variables ψ^1, \dots, ψ^l . Each subcomputation H is uniquely defined by all iteration vectors' values associated with vertices pebbled in $H = \{\psi_1, \dots, \psi_{|H|}\}$. For each iteration variable ψ^t , $t = 1, \dots, l$, denote the set of all values that ψ^t takes during H as D^t . We define $D = [D^1, \dots, D^l] \subseteq \mathcal{D}$ as the *iteration domain* of subcomputation H .

Furthermore, recall that each input access $A_j[\phi_j(\psi)]$ is uniquely defined by $dim(\phi_j)$ iteration variables $\psi_j^1, \dots, \psi_j^{dim(\phi_j)}$. Denote the set of all values each of ψ_j^k takes during H as D_j^k . Given D , we also denote the number of different vertices that are accessed from each input array A_j as $|A_j(D)|$.

We now state the lemma which bounds $|H|$ by the iteration sets' sizes $|D^t|$ (the intuition behind the lemma is depicted in Figure 4):

Lemma 3. *Given the ranges of all iteration variables D^t , $t = 1, \dots, l$ during subcomputation H , if $|H| = \prod_{t=1}^l |D^t|$, then $\forall j = 1, \dots, m$: $|A_j(D)| = \prod_{k=1}^{dim(\phi_j)} |D_j^k|$ and $|H|$ is maximized among all valid subcomputations that iterate over $D = [D^1, \dots, D^l]$.*

Intuition. Lemma 3 states that if each iteration variable ψ^t , $t = 1, \dots, l$ takes $|D^t|$ different values, then there are at most $\prod_{t=1}^l |D^t|$ different iteration vector values ψ that can be formed in H . Thus, to maximize $|H|$ all combinations of values of ψ^t should be evaluated. On the other hand, this also implies the maximization of all access sizes $|A_j(D)| = \prod_{k=1}^{dim(\phi_j)} |D_j^k|$. This result is more general than, e.g., polyhedral techniques [11, 15, 51] as it does not require loop nests to be affine. Instead, it solely relies on set algebra and combinatorial methods.

3.2 Finding the I/O Lower Bound

Denoting $H_{max} = \arg \max_{H \in \mathcal{P}(X)} |H|$ as the largest subcomputation among all valid X -partitions, we use Lemma 3 and combine it with the dominator set constraint from Section 2.3.3. Note that all access set sizes are strictly positive integers $|D^t| \in \mathbb{N}_+$, $t = 1, \dots, l$.

¹<https://arxiv.org/abs/2108.09337>

Otherwise, if any of the sets is empty, no computation can be performed. However, as we only want to find the bound on the number of I/O operations, we relax the integer constraints and replace them with $|D^t| \geq 1$. Then, we formulate finding $\chi(X)$ (Lemma 2) as the following optimization problem:

$$\begin{aligned} & \max \prod_{t=1}^l |D^t| \quad \text{s.t.} \\ & \sum_{j=1}^m \prod_{k=1}^{\dim(\phi_j)} |D_j^k| \leq X \\ & \forall 1 \geq t \geq l : |D^t| \geq 1 \end{aligned} \quad (1)$$

We then find $|H_{max}| = \chi(X)$ as a function of X using Karush–Kuhn–Tucker (KKT) conditions [43]. Next, we solve

$$\frac{d \chi(X)}{dX} = 0. \quad (2)$$

Denoting X_0 as the solution to Equation (2), we finally obtain

$$Q \geq |V| \frac{(X_0 - M)}{\chi(X_0)}. \quad (3)$$

Computational intensity and out-degree-one vertices. There exist cDAGs where every non-input vertex has at least $u \geq 0$ direct predecessors that are input vertices with out-degree 1. We can use this fact to put an additional bound on the computational intensity.

Lemma 4. *If in a cDAG $G = (V, E)$ every non-input vertex has at least u direct predecessors with out-degree one that are graph inputs, then the maximum computational intensity ρ of this cDAG is bounded by $\rho \leq \frac{1}{u}$.*

4 DATA REUSE ACROSS MULTIPLE STATEMENTS

Until now, we have analyzed each statement separately. However, almost all computational kernels contain multiple statements connected by data dependencies – e.g., a column update (S1) and a trailing matrix update (S2) in LU factorization (Figure 3). The challenge here is that, in general, I/O cost Q is not composable: due to the data reuse, the total I/O cost of the program may be smaller than the sum of I/O costs of its constituent kernels. In this section we examine how these dependencies influence the total I/O cost of a program.

We derive I/O lower bounds for programs with w statements S_1, \dots, S_w in two steps. First, we analyze each statement S_i separately, as in Section 3. Then, we derive how many loads could be avoided at most during one statement if another statement owned shared data. There are two cases where data reuse can occur: **I**) input overlap, where shared arrays are inputs for multiple statements, and **II**) output overlap, where the output array of one statement is the input array of another.

Case I. Assume there are w statements in the program, and there are k arrays $A_j, j = 1, \dots, k$ that are shared between at least two statements. We still evaluate each statement separately, but we will subtract the upper bound on shared loads $Q_{tot} \geq \sum_{i=1}^w Q_i - \sum_{j=1}^k |Reuse(A_j)|$, where $|Reuse(A_j)|$ is the reuse bound on array A_j (Section 4.1). **Case II.** Consider each pair of “producer-consumer” statements S and T , that is, the output of S is the input of T . The

I/O lower bound Q_S of statement S does not change due to the reuse, as the same number of loads has to be performed to evaluate S . On the other hand, it may invalidate Q_T , as the dominator set of T formulated in Section 3.1 may not be minimal – inputs of a statement may not be graph inputs anymore. For each “consumer” statement T , we reevaluate $Q'_T \leq Q_T$ using Lemma 6. Finally, for a program consisting of w statements in total, connected by the output overlap, we have $Q_{tot} \geq \sum_{i=1}^w Q'_i$. Note that for each “producer” statement i , $Q'_i = Q_i$, i.e. output overlap does not change their I/O lower bound.

4.1 Case I: Input Reuse and Reuse Size

Consider two statements S and T that share one input array A_i . Let $|A_i(\mathbf{R}_S)|$ denote the total number of accesses to A_i during the I/O optimal execution of a program that contains only statement S . Naturally, $|A_i(\mathbf{R}_T)|$ denotes the same for a program containing only T . Define $Reuse(A_i) := \min\{|A_i(\mathbf{R}_S)|, |A_i(\mathbf{R}_T)|\}$. We then have:

Lemma 5. *The I/O cost of a program containing statements S and T that share the input array A_i is bounded by*

$$Q_{tot} \geq Q_S + Q_T - Reuse(A_i),$$

where Q_S, Q_T are the I/O costs of a program containing only statement S or T , respectively.

Note that $Reuse(A_i)$ is an overapproximation of the actual reuse. Since finding the optimal schedule is PSPACE-complete [46], we conservatively assume that only the minimum number of loads from A_i is performed. Thus, Lemma 5 generalizes to any number of statements S_1, \dots, S_w sharing array A_i – the total number of loads from A_i is lower-bounded by a maximum number of loads from A_i among $S_j, \max_{j=1, \dots, w} |A_i(\mathbf{R}_{S_j})|$.

4.2 Case II: Output Reuse and Access Sizes

Consider the case where the *output* A_0 of statement S is also the *input* B_j of statement T . Furthermore, consider subcomputation H of statement T (and its associated iteration domain D). Any path from the graph inputs to vertices in $B_0(D)$ must pass through vertices in $B_j(D)$. The following question arises: Is there a smaller set of vertices $B'_j(D), |B'_j(D)| < |B_j(D)|$ that every path from graph inputs to $B_j(D)$ must pass through?

Let ρ_S denote computational intensity of statement S . With that, we can state the following lemma:

Lemma 6. *Any dominator set of set $B_j(D)$ must be of size at least $|Dom(B_j(D))| \geq \frac{|B_j(D)|}{\rho_S}$.*

Corollary 1. *Combining Lemmas 6 and 3, the data access size of $|B_j(D)|$ during subcomputation H is*

$$|Dom(B_j(D))| \geq \frac{\prod_{k=1}^{\dim(\phi_j)} |D_j^k|}{\rho_S}. \quad (4)$$

Similarly to **case I**, this result also generalizes to multiple “consumer” statements that reuse the same output array of a “producer” statement, as well as any combination of input and output reuse for multiple arrays and statements. Since the actual I/O optimal schedule is unknown, the general strategy to ensure correctness of our lower bound is to consider each pair of interacting statements separately as one of these two cases. Since both Lemma 5 and 6

overapproximate the reuse, the final bound may not be tight - the more inter-statement reuse, the more overapproximation is needed. Still, this method can be successfully applied to derive *tight* I/O lower bounds for many linear algebra kernels, such as matrix factorizations, tensor products, or solvers.

5 GENERAL PARALLEL I/O LOWER BOUNDS

We now establish how our method applies to a parallel machine with P processors (Section 2.1). Since we target large-scale distributed systems, our parallel pebbling model differs from the one introduced e.g. by Alwen and Serbinenko [5], which is inspired by shared-memory models like PRAM [39]. Instead, we disallow sharing memory (pebbles) between the processors, and enforce explicit communication – analogous to the load/store operations – using red and blue pebbles. This allows us to better match the behavior of real, distributed applications that use, e.g., MPI.

Each processor p_i owns its private fast memory that can hold up to M words, represented in the cDAG as M vertices of color p_i . Vertices with different colors (belonging to different processors) cannot be shared between these processors, but any number of different pebbles may be placed on one vertex.

All the standard red-blue pebble game rules apply with the following modifications:

- (1) **compute.** If all direct predecessors of vertex v have pebbles of p_i 's color placed on them, one can place a pebble of color p_i on v (no sharing of pebbles between processors),
- (2) **communication.** If a vertex v has *any* pebble placed on it, a pebble of any other color may be placed on this vertex.

From this game definition it follows that from a perspective of a single processor p_i , any data is either local (the corresponding vertex has p_i 's pebble placed on it) or remote, without a distinction on the remote location (remote access cost is uniform).

Lemma 7. *The minimum number of I/O operations in a parallel pebble game, played on a cDAG with $|V|$ vertices with P processors each equipped with M pebbles, is $Q \geq \frac{|V|}{P \cdot \rho}$, where ρ is the maximum computational intensity, which is independent of P (Lemma 1).*

6 I/O LOWER BOUNDS OF PARALLEL FACTORIZATION ALGORITHMS

We gather all the insights from Sections 2 to 5 and use them to obtain the parallel I/O lower bounds of LU and Cholesky factorization algorithms, which we use to develop our communication-avoiding implementations.

6.1 LU Factorization

In our I/O lower bound analysis we omit the row pivoting, since swapping rows can increase the I/O cost by at most N^2 , which is the cost of permuting the entire matrix. However, the total I/O cost of the LU factorization is $O(N^3/\sqrt{M})$ [61].

LU factorization (without pivoting) contains two statements (Figure 3). Observe that we can use Lemma 4 (out-degree one vertices) for statement $S1 : A[i, k] = A[i, k] / A[k, k]$. The loop nest depth is $l_{S1} = 2$, with iteration variables $\psi^1 = k$ and $\psi^2 = i$. The dimension of the access function vector (k, k) is 1, revealing potential for data reuse: every input vertex $A[k, k]$ is accessed $N - k$ times and used to compute vertices $A[i, k]$, $k + 1 \leq i < N$. However,

the access function vector (i, k) has dimension 2; therefore, every compute vertex has one direct predecessor with out-degree one, which is the previous version of element $A[i, k]$. Using Lemma 4, we therefore have $\rho_{S1} \leq 1$.

We now proceed to statement $S2 : A[i, j] = A[i, j] - A[i, k] * A[k, j]$. Let $|D^k| = K$, $|D^i| = I$, $|D^j| = J$. Observe that there is an output reuse (Section 4.2 and Figure 3, red arrow) of $A[i, k]$ between statements $S1$ and $S2$. We therefore have the following access size in statement $S2$: $|A_2(D_{S2})| = |A[i, k]| = \frac{IK}{\rho_{S1}} \geq IK$ (Equation 4). Note that in this case where the computational intensity is $\rho_{S1} \leq 1$, the output reuse does not change the access size $|A_2(D_{S2})|$ of statement $S2$. This follows the intuition that it is not beneficial to recompute vertices if the recomputation cost is higher than loading it from the memory. Denoting H_{S2} as the maximal subcomputation for statement $S2$ over the subcomputation domain D , we have the following (Lemma 3):

- $|H_{S2}| = KIJ$
- $|A_1(D)| = |A[i, j]| = IJ$
- $|A_2(D)| = |A[i, k]| = IK$
- $|A_3(D)| = |A[k, j]| = KJ$
- $|Dom(H_{S2})| = |A_1(D)| + |A_2(D)| + |A_3(D)| = IJ + IK + KJ$

We then solve the optimization problem from Section 3.2:

$$\begin{aligned} & \max \quad KIJ, & \text{s.t.} \\ & IJ + IK + KJ \leq X \\ & I \geq 1, \quad J \geq 1, \quad K \geq 1 \end{aligned}$$

Which gives $|H_{S2}| = \chi(X) = \left(\frac{X}{3}\right)^{\frac{3}{2}}$ for $K = I = J = \left(\frac{X}{3}\right)^{\frac{1}{2}}$.

Then, we find X_0 that minimizes the expression $\rho_{S2}(X) = \frac{|H_{max}|}{X - M}$ (Equation 2), yielding $X_0 = 3M$. Plugging it into $\rho_{S2}(X)$, we conclude that the maximum computational intensity of $S2$ is bounded by $\rho_{S2} \leq \sqrt{M}/2$.

We bounded the maximum computational intensities ρ_{S1} and ρ_{S2} , that is, the minimum number of I/O operations to compute vertices belonging to statements $S1$ and $S2$. As the last step, we find the total number of compute vertices for each statement: $|V_1| = \sum_{k=1}^N (N - k - 1) = N(N - 1)/2$, and $|V_2| = \sum_{k=1}^N \sum_{i=k+1}^N (N - k - 1) = N(N - 1)(N - 2)/3$. Using Lemmas 1 (bounding I/O cost with the computational intensity) and 7 (I/O cost of the parallel machine), the parallel I/O lower bound for LU factorization is therefore

$$Q_{P, LU} \geq \frac{2N^3 - 6N^2 + 4N}{3P\sqrt{M}} + \frac{N(N - 1)}{2P} = \frac{2N^3}{3P\sqrt{M}} + O\left(\frac{N^2}{P}\right).$$

Previously, Solomonik et al. [61] established the asymptotic I/O bound for sequential execution $Q = O(N^3/\sqrt{M})$. Recently, Olivry et al. [51] derived a tight leading term constant $Q \geq 2N^3/(3\sqrt{M})$. To the best of our knowledge, our result is the first non-asymptotic bound for parallel execution. The generalization from the sequential to the parallel bound is straightforward. Note, however, that this is only the case due to our pebble-based execution model, and it may thus not apply to other parallel machine models.

6.2 Cholesky Factorization

We proceed analogously to our derivation of the LU I/O bound – here we just briefly outline the steps. The algorithm contains three statements (Listing 1). For statements $S1$ and $S2$, we can

again use Lemma 4 (out-degree-one vertices). For $S1 : L(k, k) = \text{sqrt}(L(k, k))$, the loop nest depth is $l_1 = 1$, we have a single iteration variable $\psi^1 = k$, and a single input array $A_1 = L$ with the access function $\phi_1(\psi) = (k, k)$. Since there is only one iteration variable present in ϕ_1 , we have $\text{dim}(\phi_1) = 1 = l_1$. Therefore, for every compute vertex v we have one direct predecessor, which is the previous version of element $L(k, k)$. We conclude that $\rho_{S1} \leq 1$ and $|V_{S1}| = N$.

```

1  for k = 1:N
2  S1:   L(k, k) = sqrt(L(k, k));
3      for i = k+1:N
4  S2:   L(i, k) = (L(i, k)) / L(k, k);
5      for j = k+1:i
6  S3:   L(i, j) = L(i, j) - L(i, k) * L(j, k);
7  end; end; end;

```

Listing 1: Cholesky Factorization

For statement $S2 : L(i, k) = (L(i, k)) / L(k, k)$, we also have output reuse of $L(k, k)$ between statements $S2$ and $S1$. However, as with the output reuse considered in the LU analysis, the computational intensity is $\rho_{S1} \leq 1$. Therefore, it does not change the dominator set size of $S2$. We then use the same reasoning as for the corresponding statement $S1$ in LU factorization, yielding $\rho_{S2} \leq 1$.

For statement $S3$, we derive its bound similarly to $S2$ of LU, with $\rho_{S3} = \sqrt{M}/2$ and $|V_{S3}| = \sum_{k=1}^N \sum_{i=k+1}^N (i - k - 1) = N(N-1)(N-2)/6$. Note that compared to LU, the only significant difference is the iteration domain $|V_3|$. Even though Cholesky has one statement more – the diagonal element update $L(k, k)$ – its impact on the final I/O bound is negligible for large N .

Again, using Lemmas 1 and 7 we establish the Cholesky factorization’s parallel I/O lower bound:

$$Q_{Chol} \geq Q_1 + Q_2 + Q_3 = \frac{|V_1|}{P\rho_1} + \frac{|V_2|}{P\rho_2} + \frac{|V_3|}{P\rho_3} \approx \frac{N^3}{3P\sqrt{M}} + \frac{N^2}{2P} + \frac{N}{P}$$

The derived I/O lower bound for a sequential machine ($P = 1$) improves the previous bound $Q_{Chol} \geq N^3/(6\sqrt{M})$ derived by Olivry et al. [51]. Furthermore, to the best of our knowledge, this is the first parallel bound for this kernel.

7 NEAR-I/O OPTIMAL PARALLEL MATRIX FACTORIZATION ALGORITHMS

We now present our parallel LU and Cholesky factorization algorithms. We start with the former, more complex algorithm, i.e. LU factorization. Pivoting in LU poses several performance challenges. First, since pivots are not known upfront, additional communication and synchronization is required to determine them in each step. Second, the nondeterministic pivot distribution between the ranks may introduce load imbalance of computation routines. Third, to minimize the communication a 2.5D parallel decomposition must be used, i.e. parallelization along the reduction dimension. We address all these challenges with *CONfLUX* – a near *Communication Optimal LU factorization using X-Partitioning*.

7.1 LU Dependencies and Parallelization

Due to the dependency structure of LU, the input matrix is often divided recursively into four submatrices A_{00} , A_{10} , A_{01} , and

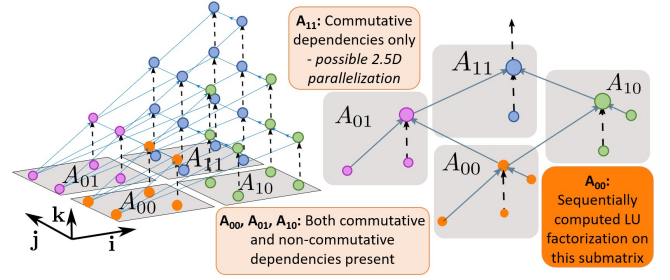


Figure 5: LU Factorization cDAG for $n = 4$ with the logical decomposition into A_{00} , A_{10} , A_{01} , and A_{11} . Dashed arrows represent commutative dependencies (reduction of a value). Solid arrows represent non-commutative operations, meaning that any parallel pebbling has to respect the induced order (e.g., no vertex in A_{11} can be pebbled before A_{00} is pebbled).

A_{11} [24, 61]. Arithmetic operations performed in LU create non-commutative dependencies (Figure 5) between vertices in A_{00} (LU factorization of the top-left corner of the matrix), A_{10} , and A_{01} (triangular solve of left and top panels of the matrix). Only A_{11} (Schur complement update) has no such dependencies, and may therefore be efficiently parallelized in the reduction dimension. A high-level summary is presented in Algorithm 1.

Algorithm 1 CONfLUX

Input: Input matrix $A \in \mathbb{R}^{n \times n}$

Output: In-place factored matrix A , permutation matrix P

```

A1 ← A                                     ▷ First step
P ← I                                       ▷ Permutation matrix is initially identity
for t = 1, . . . , N/v do
  1. Reduce next block column                ▷ Cost: (N-t·v)·v·M/N²
  2. [rows, Pt+1] ← TournPivot(At, Pt)   ▷ Find next v pivots. Cost: v² ⌈log(N/v)⌉
  3. Scatter computed A00 and v pivot rows   ▷ Cost: v² + v
  4. Scatter A10                             ▷ Cost: (N-t·v)v/P
  5. Reduce v pivot rows                       ▷ Cost: (N-t·v)·v·M/N²
  6. Scatter A01                             ▷ Cost: (N-t·v)v/P
  7. FactorizeA10(At)                       ▷ 1D parallel, block-row
  8. Send data from panel A10                 ▷ Cost: (N-t·v)N·v/P√M
  9. FactorizeA01(At)                       ▷ 1D parallel, block-column
  10. Send data from panel A01                ▷ Cost: (N-t·v)N·v/P√M
  11. FactorizeA11(At)                       ▷ 2.5D parallel
  At+1 ← At[rows, v : end] ▷ Recursively process remaining rows and columns
end for

```

7.2 LU Computation Routines

The computation is performed in $\frac{N}{v}$ steps, where v is a tunable block size. In each step, only submatrix A_t of input matrix A is updated. Initially, A_t is set to A . A_t can be further viewed as composed of four submatrices A_{00} , A_{10} , A_{01} , and A_{11} (see Figure 6). These submatrices are distributed and updated by routines *TournPivot*, *FactorizeA₁₀*, *FactorizeA₀₁*, and *FactorizeA₁₁*:

- A_{00} . This $v \times v$ submatrix contains the first v elements of the current $v \times v$ submatrix. It is computed during *TournPivot*, and, as it is required to compute A_{10} and A_{01} , it is redundantly copied to all processors.
- A_{10} and A_{01} . Submatrices A_{10} and A_{01} of sizes $(N - t \cdot v) \times v$ and $v \times (N - t \cdot v)$ are distributed using a 1D decomposition




Figure 6: CONfLUX’s parallel decomposition for $P = 8$ processors decomposed into a $[P_x, P_y, P_z] = [2, 2, 2]$ grid, together with the indicated steps of Algorithm 1. In each iteration t , each processor $[pi, pj, pk]$ updates $(2 - \lfloor(t + pi)/Px\rfloor) \times (2 - \lfloor(t + pj)/Py\rfloor)$ tiles of A_{11} . In the presented example, there are $v = 4$ planes in dimension k to be reduced, which are distributed among $Pz = 2$ processor layers (green and yellow tiles).

among all processors. They are updated using a triangular solve. 1D decomposition guarantees that there are no dependencies between processors, so no communication or synchronization is performed during computation, as A_{00} is already owned by every processor.

- A_{11} This $(N - t \cdot v) \times (N - t \cdot v)$ submatrix is distributed using a 2.5D, block-cyclic distribution (Figure 6). First, the updated submatrices A_{10} and A_{01} are broadcast among the processors. Then, A_{11} (Schur complement) is updated. Finally, the first block column and v chosen pivot rows are reduced, which will form A_{10} and A_{01} in the next iteration.

Block size v . The minimum size of each block is the number of processor layers in the reduction dimension ($v \geq c = \frac{PM}{N^2}$). However, to secure high performance, this value should be adjusted to hardware parameters of a given machine (e.g., vector length, prefetch distance of a CPU, or warp size of a GPU). Throughout the analysis, we assume $v = a \cdot \frac{PM}{N^2}$ for some small constant a .

7.3 Pivoting

Our pivoting strategy differs from state-of-the-art block [6], tile [4], or recursive [24] pivoting approaches in two aspects:

- To minimize I/O cost, we do not swap pivot rows. Instead, we keep track of which rows were chosen as pivots and we use masks to update the remaining rows.
- To reduce latency, we take advantage of our derived block decomposition and use tournament pivoting [29].

Tournament Pivoting. This procedure finds v pivot rows in each step that are then used to mask which rows will form the new A_{01} and then filter the non-processed rows in the next step. It is shown to be as stable as partial pivoting [29], which might be an issue for, e.g., incremental pivoting [54]. On the other hand, it reduces the $O(N)$ latency cost of partial pivoting, which requires step-by-step column reduction to find consecutive pivots, to $O(\frac{N}{v})$, where v is the tunable block size parameter.

Row Swapping vs. Row Masking. To achieve close to optimal I/O cost, we use a 2.5D decomposition. This, however, implies that in the presence of extra memory, the matrix data is replicated $\frac{PM}{N^2}$ times. This increases the row swapping cost from $O(\frac{N^2}{P})$ to

$O(\frac{N^3}{P\sqrt{M}})$, which asymptotically matches the I/O lower bound of the entire factorization. Performing row swapping would then increase the constant term of the leading factor of the algorithm from $\frac{N^3}{P\sqrt{M}}$ to $\frac{2N^3}{P\sqrt{M}}$. To keep the I/O cost of our algorithm as low as possible, instead of performing row-swapping, we only propagate pivot row indices. When the tournament pivoting finds the v pivot rows, they are broadcast to all processors with only v cost per step.

Pivoting in CONfLUX. In each step t of the outer loop (line 1 in Algorithm 1), $\frac{N}{\sqrt{M}}$ processors perform a tournament pivoting routine using a butterfly communication pattern [55]. Each processor owns $\sqrt{M} \frac{N-vt}{N}$ rows, among which it chooses v local candidate pivots. Then, final pivots are chosen in $\log(\frac{N}{\sqrt{M}})$ “playoff-like” tournament rounds, after which all $\frac{N}{\sqrt{M}}$ processors own both v pivot row indices and the already factored, new A_{00} . This result is distributed to all remaining processors (line 2). Pivot row indices are then used to determine which processors participate in the reduction of the current A_{01} (line 4). Then, the new A_t is formed by masking currently chosen rows $A_t \leftarrow A_t[\text{rows}, v : \text{end}]$ (Line 12).

7.4 I/O cost of CONfLUX

We now prove the I/O cost of CONfLUX, which is only a factor of $\frac{1}{3}$ higher than the obtained lower bound for large N .

Lemma 8. *The total I/O cost of CONfLUX, presented in Algorithm 1, is $Q_{\text{CONfLUX}} = \frac{N^3}{P\sqrt{M}} + O(\frac{N^2}{P})$.*

PROOF. We assume that the input matrix A is already distributed in the block cyclic layout imposed by the algorithm. Otherwise, data reshuffling imposes only $O(\frac{N^2}{P})$ cost, which does not contribute to the leading order term. We first derive the cost of a single iteration t of the main loop of the algorithm, proving that its cost is $Q_{\text{step}}(t) = \frac{2Nv(N-tv)}{P\sqrt{M}} + O(\frac{Nv}{P})$. The total cost after $\frac{N}{v}$ iterations is:

$$Q_{\text{CONfLUX}} = \sum_{t=1}^{\frac{N}{v}} Q_{\text{step}}(t) = \frac{N^3}{P\sqrt{M}} + O\left(\frac{N^2}{P}\right).$$

We define $P1 = \frac{N^2}{M}$ and $c = \frac{PM}{N^2}$. P processors are decomposed into the 3D grid $[\sqrt{P1}, \sqrt{P1}, c]$. We refer to all processors that share the same second and third coordinate as $[:, j, k]$. We now examine each of 11 steps in Algorithm 1.

Step 2. Processors with coordinates $[:, t \bmod \sqrt{P1}, t \bmod c]$ perform the tournament pivoting. Every processor owns the first v elements of $N - (t - 1)v$ rows, among which they choose the next v pivots. First, they locally perform the LUP decomposition to choose the local v candidate rows. Then, in $\lceil \log_2(\sqrt{P1}) \rceil$ rounds they exchange $v \times v$ blocks to decide on the final pivots. After the exchange, these processors also hold the factorized submatrix A_{00} . *I/O cost per processor:* $v^2 \lceil \log_2(\sqrt{P1}) \rceil$.

Steps 3, 4, 6. Factorized A_{00} and v pivot row indices are broadcast. First v columns and v pivot rows are scattered to all P . *I/O cost per processor:* $v^2 + v + \frac{2(N-tv)v}{P}$.

Steps 1 and 5. v columns and v pivot rows are reduced. With high probability, pivots are evenly distributed among all processors. There are c layers to reduce, each of size $(N - tv)v$. *I/O cost per processor:* $\frac{(N-tv)vc}{P} = \frac{2(N-tv)vM}{N^2}$.

		CONfLUX (LU)		CONfCHOX (Cholesky)		
	description	comm. cost	comp. cost	description	comm. cost	comp. cost
pivoting	<i>TournPivot</i>	$v^2 \lceil \log_2(\sqrt{P1}) \rceil$	$v^3/3 \lceil \log_2(\sqrt{P1}) \rceil$	(no pivoting)	—	—
A_{00}	local getrf	0	0 (done during <i>TournPivot</i>)	potrf	v^2	$v^3/6$
A_{10} and A_{01}	reduction, local trsm	$\frac{2(N-tv)vM}{N^2}$	$\frac{2(N-tv)v^2}{2P}$	(similar to LU)	$\frac{2(N-tv)vM}{N^2}$	$\frac{2(N-tv)v^2}{2P}$
A_{11}	scatter, local gemm	$\frac{2(N-tv)v}{P}$	$\frac{(N-tv)^2v}{P}$	scatter, local gemmt (triangular gemm)	$\frac{2(N-tv)v}{P}$	$\frac{(N-tv)^2v}{2P}$

Table 1: Comparison of the implemented LU and Cholesky factorizations. Even though Cholesky performs half as many computations (the use of gemmt instead of gemm in A_{11}), it communicates the same amount of data, since the number of elements needed to perform gemm and gemmt is the same.

Steps 7, 9, 11. The updates *Factorize* A_{10} , *Factorize* A_{01} , and *Factorize* A_{11} are local and incur no additional I/O cost.

Steps 8 and 10. Factorized A_{10} and A_{01} are scattered among all processors. Each processor requires $\frac{v(N-tv)}{c\sqrt{P1}}$ elements from A_{10} and A_{01} . I/O cost per processor: $\frac{2(N-tv)Nv}{P\sqrt{M}}$.

Summing steps 1 – 11: $Q_{step}(t) = \frac{2Nv(N-tv)}{P\sqrt{M}} + O\left(\frac{Nv}{P}\right)$. \square

Note that this cost is a factor 1/3 over the lower bound established in Section 6.1. This is due to the fact that any processor can only maximally utilize its local memory in the first iteration of the outer loop. In this first iteration, a processor updates a total of $\sqrt{M} \times \sqrt{M}$ elements of A . In subsequent iterations, however, the local domain shrinks as less rows and columns are updated, which leads to an underutilization of the resources. Since the shape of the iteration space is determined by the algorithm, this behavior is unavoidable for $P \geq N^2/M$. Note that the bound is attainable by a sequential machine, however.

7.5 Cholesky Factorization

From a data flow perspective, Cholesky factorization can be viewed as a special case of LU factorization without pivoting for symmetric, positive definite matrices. Therefore, our Cholesky algorithm – CONfCHOX – heavily bases on CONfLUX, using the same 2.5D parallel decomposition, block-cyclic data distribution, and analogous computation routines.

For both algorithms, the dominant cost, both in terms of computation and communication, is the A_{11} update. Due to the Cholesky factorization’s iteration domain, which exploits the symmetry of the input matrix, the compute cost is twice as low, as only one half of the matrix needs to be updated. However, the input size required to perform this update is the same – therefore, the communication cost imposed by A_{11} is similar. We list the key differences between the two factorization algorithms in Table 1.

8 IMPLEMENTATION

Our algorithms are implemented in C++, using MPI for inter-node communication. For static communication patterns (e.g., column reductions) we use dedicated, asynchronous MPI collectives. For runtime-dependent communication (e.g., pivot index distribution) we use MPI one-sided [31]. For intra-node tasks, we use OpenMP

and local BLAS calls (provided by Intel MKL [34]) for computations. Our code is available as an open-source git repository².

Parallel decomposition. Our experiments show that the parallelization in the reduction dimension, while reducing communication volume, does incur performance overheads. This is mainly due to the increased communication latency, as well as smaller buffer sizes used for local BLAS calls. Since formal modeling of the tradeoff between communication volume and performance is outside of the scope of the paper, we keep the depth of parallelization in the third dimension as a tunable parameter, while providing heuristics-based default values.

Data distribution. CONfLUX and CONfCHOX provide ScaLAPACK wrappers by using the highly-optimized COSTA algorithm [38] to transform the matrices between different layouts. In addition, they support the COSTA API for matrix descriptors, which is more general than ScaLAPACK’s layout, as it supports matrices distributed in arbitrary grid-like layouts, processor assignments, and local blocks orderings.

9 EXPERIMENTAL EVALUATION

We compare CONfLUX and CONfCHOX with state-of-the-art implementations of corresponding distributed matrix factorizations.

Measured values. We measure both the I/O cost and total time-to-solution. For I/O, the aggregate communication volume in distributed runs is counted using the Score-P profiler [41]. We provide both measured values and theoretical cost models. Local `std::chrono` calls are used for time measurements and the maximum execution time among all ranks is reported.

Infrastructure and Measurement. We run our experiments on the XC40 partition of the CSCS Piz Daint supercomputer which comprises 1,813 CPU nodes equipped with Intel Xeon E5-2695 v4 processors (2x18 cores, 64 GiB DDR3 RAM), interconnected by the Cray Aries network with a Dragonfly network topology. Since the CPUs are dual-socket, two MPI ranks are allocated per compute node.

Comparison Targets. We use 1) Intel MKL (v19.1.1.217). While the library is proprietary, our measurements reaffirm that, like ScaLAPACK, the implementation uses the suboptimal 2D processor decomposition; 2) SLATE [28] – a state-of-the-art distributed linear algebra framework targeted at exascale supercomputers; 3)

²<https://github.com/eth-cscs/conflux>







	MKL [34]	SLATE [28]	CANDMC [57]	CAPITAL [33]	CONfLUX / CONfCHOX (this work)
Decomposition	2D, panel decomp.	2D, block decomp.	Nested 2.5D, block decomp.	2.5D, block decomp.	1D / 2.5D, block decomp.
Block size	 user-specified	 user-specified, (default 16)	 $\frac{N^3}{P \cdot M}, \frac{N^2}{P \sqrt{M}}$	 user-specified	 optimized, $\geq \frac{P \cdot M}{N^2}$
Program parameters	required from user 🗨	required from user 🗨	provided defaults 🗨	optimized defaults 🗨🗨	optimized defaults 🗨🗨
Parallel I/O cost	$\frac{N^2}{\sqrt{P}} + O\left(\frac{N^2}{P}\right)$	$\frac{N^2}{\sqrt{P}} + O\left(\frac{N^2}{P}\right)$	$\frac{5N^3}{P \sqrt{M}} + O\left(\frac{N^2}{P \sqrt{M}}\right)$ [61]	$\frac{45N^3}{8P \sqrt{M}} + O\left(\frac{N^2}{P \sqrt{M}}\right)$ [33]	$\frac{N^3}{P \sqrt{M}} + O\left(\frac{N^2}{P \sqrt{M}}\right)$

Table 2: Parallelization strategies and I/O cost models of the considered matrix factorization implementations. MKL and SLATE require a user to specify the processor decomposition and the block size. CANDMC provides default values, but our experiments show that the performance was significantly improved when we tuned the parameters. CONfLUX and CONfCHOX outperform all state-of-the-art libraries with out-of-the-box parameters. We validated our parallel I/O cost models: for MKL, SLATE, CONfLUX, and CONfCHOX, the error was within +/-3%. For CANDMC and CAPITAL, we used the models derived by the authors [33, 61], which overapproximated the measured values by approx. 30-40%.



(a) Communication volume per node for varying node counts P and a fixed $N = 16,384$. Only the leading factors of the models are shown. The models are scaled by the element size (8 bytes). (b) Communication volume per node for weak scaling (constant work per node), $N = 3200 \cdot \sqrt[3]{P}$. 2.5D algorithms (CANDMC and CONfLUX) retain constant communication volume per processor. (c) Communication reduction vs. second-best algorithm (M=MKL, S=SLATE), for varying P, N , for both measured and predicted scenarios.

Figure 7: Communication volume measurements across different scenarios for MKL, SLATE, CANDMC, and CONfLUX. In all considered scenarios, enough memory $M \geq N^2/P^{2/3}$ was present to allow for the maximum number of replications $c = P^{1/3}$.

the latest version of the CANDMC and CAPITAL libraries [32, 58], which use an asymptotically-optimal 2.5D decomposition. The implementations and their characteristics are listed in Table 2.

Problem Sizes. We evaluate the algorithms starting from 2 compute nodes (4 MPI ranks) up to 512 nodes (1,024 ranks). For each node count, matrix sizes range from $N = 2,048$ to $N = 2^{19} = 524,288$, provided they fit into the allocated memory (e.g., LU or Cholesky factorization on a double-precision input matrix of dimension $262,144 \times 262,144$ cannot be run on less than 32 nodes). Runs in which none of the libraries achieved more than 3% of the hardware peak are discarded since by adding more nodes the performance starts to deteriorate.

Our benchmarks reflect real-world problems in scientific computing. The High-Performance Linpack benchmark uses a maximal size of $N = 16,473,600$ [62]. In quantum physics, matrix size scales with 2^{qubits}. In physical chemistry or density functional theory (DFT), simulations require factorizing matrices of atom interactions, yielding sizes ranging from $N = 1,024$ up to $N = 131,072$ [18, 66]. In machine learning, matrix factorizations are used for inverting Kronecker factors [52] whose sizes are usually around $N = 4,096$. This motivates us to focus not only on exascale problems, but also improve performance for relatively small matrices ($N \leq 100,000$).

Communication Models. Together with empirical measurements, we put significant effort into understanding the underlying communication patterns of the compared LU factorization implementations. Both MKL and SLATE base on the standard partial pivoting algorithm using the 2D decomposition [10]. For CANDMC and CAPITAL, the models provided by the authors [33, 61] are used. For CONfLUX and CONfCHOX, we use the results from Section 7. These models are summarized in Table 2.

10 RESULTS

Our experiments confirm advantages of CONfLUX and CONfCHOX in terms of both communication volume and time-to-solution over all other implementations tested. A significant communication reduction can be observed (up to 1.42 times for CONfLUX compared with the second-best implementation for $P = 1,024$). Moreover, the performance models predict even greater benefits for larger runs (expected 2.1 times communication reduction for a full-machine run on the Summit supercomputer – Figure 7c). Most importantly, our implementations consistently outperform existing implementations (up to three times – Figures 1 and 8).

Communication volume. Fig. 7a presents the measured communication volume per node, as well as our derived cost models




Figure 8: Achieved % of peak performance for LU factorization. We show median and 95% confidence intervals.




Figure 9: Achieved % of peak performance for Cholesky factorization. We show median and 95% confidence intervals.

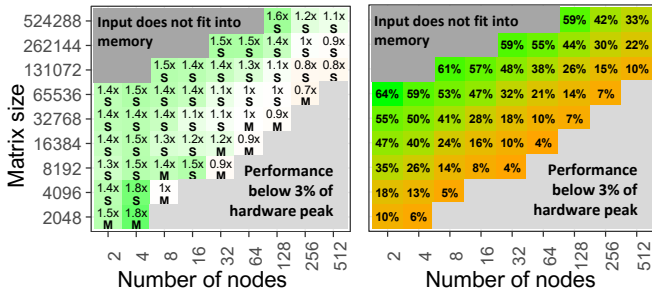


Figure 10: Left: measured runtime speedup of *ConfCHOX* vs. fastest state-of-the-art library (S=SLATE [28], C=CAPITAL [33], M=MKL [34]). Right: *ConfCHOX*'s achieved % of machine peak performance.

(Table 2) presented with solid lines, for $N = 16,384$. Observe that *ConflLUX* communicates the least for all values of P . Note that since both MKL and SLATE use similar 2D decompositions, their communication volumes are mostly equal, with a slight advantage for SLATE. In Fig. 7b, we show the weak scaling characteristics of the analyzed implementations. Observe that for a fixed amount of work per node, the 2D algorithms - MKL and SLATE - scale

sub-optimally. Figure 7c summarizes the communication volume reduction of *ConflLUX* compared with the second-best implementation, both for measurements and theoretical predictions. It can be seen that for all combinations of P and N , *ConflLUX* always communicates the least. For all measured data points, the asymptotically optimal CANDMC performed worse than MKL or SLATE. The figure also presents the predicted communication cost of all considered implementations for up to $P = 262,144$ based on our theoretical models.

Performance. Our measurements show that both *ConflLUX* and *ConfCHOX* outperform all considered state-of-the-art libraries in almost all scenarios (Figures 1 and 10). Thanks to the optimized block data decomposition and efficient overlap of computation and communication, our implementations achieve high performance already on relatively small matrices (approx. 40% of hardware peak for cases where $N^2/P > 2^{27}$). In cases where the local domain per processor becomes very small ($N^2/P < 2^{27}$) our block decomposition does not add that much benefit, since the performance is mostly latency-bound, and not bandwidth-bound. This is visible not only in strong scaling (Figures 8 and 9, a) and b)), but also in weak scaling (c)), where the input size per processor N^2/P is constant. This is again caused by latency overheads of scattering data between 1D and 2.5D layouts.

	Pebbling [13, 26, 37, 45, 56]	Projection-based [8, 15, 20, 21, 23, 51]	Problem specific [1, 9, 17, 48, 66]
Scope	👉👉 General cDAGs	👉 Programs Geometric structure of iteration space	👊 Individually tailored for given problem
Key Features	<ul style="list-style-type: none"> 👉 General scope 👉 Expresses complex data dependencies 👉 Directly exposes schedules 👉 Intuitive 👊 PSPACE-complete in general case 👊 No guarantees that a solution exists 👊 No well-established method how to automatically translate code to cDAGs 	<ul style="list-style-type: none"> 👉 Well-developed theory and tools 👉 Guaranteed to find solution for given class of programs 👊 Bounds are often not tight 👊 Fails to capture dependencies between statements 👊 Limited scope 	<ul style="list-style-type: none"> 👉 Takes advantage of problem-specific features 👉 Tends to provide best practical results 👊 Requires large manual effort for each algorithm separately 👊 Difficult to generalize 👊 Often based on heuristics with no guarantees on optimality

Table 3: Overview of different approaches to modeling data movement.

However, as the local domains become larger and may be more efficiently pipelined and overlapped using asynchronous MPI routines and intra-node OpenMP parallelism, the advantage becomes significant (Figures 8 and 9). *ConfLUX* outperforms existing libraries up to three times (for $P = 4, N = 4096$, second-best library is SLATE – Figure 1) and *ConfCHOX* achieves up to 1.8 times speedup (e.g., $P = 4, N = 4,096$, second-best is again SLATE).

Implications for Exascale. Both the communication models’ predictions (Figure 7c) and measured speedups (Figures 1 and 10) allow us to predict that when running our implementations on exascale machines, we can expect to see further performance improvements over state-of-the-art libraries. Furthermore, throughput-oriented hardware, such as GPUs and FPGAs, may benefit even more from the communication reduction of our schedules. Thus, *ConfLUX* and *ConfCHOX* not only outperform the state-of-the-art libraries at relatively small scales – which are most common use cases in practice [18, 52, 66] – but also promise speedups on full-scale performance runs on modern supercomputers.

11 RELATED WORK

Previous work on I/O analysis can be categorized into three classes (see Table 3): work based on **direct pebbling** or variants of it, such as Vitter’s block-based model [64]; works using **geometric arguments of projections** based on the Loomis-Whitney inequality [47]; and works applying optimizations limited to specific structural properties such as **affine loops** [27], and more generally, **the polyhedral model program representation** [9, 48, 51]. Although the scopes of those approaches significantly overlap – for example, kernels like matrix multiplication can be captured by most of the models – there are important differences both in methodology and the end-results they provide, as summarized in Table 3.

Dense linear algebra operators are among the standard core kernels in scientific applications. Ballard et al. [8] present a comprehensive overview of their asymptotic I/O lower bounds and I/O minimizing schedules, both for sparse and dense matrices. Recently, Olivry et al. introduced IOLB [51] – a framework for assessing sequential lower bounds for polyhedral programs. However, their computational model disallows recomputation (cf. Section 4.2).

Matrix factorizations are included in most of linear solvers’ libraries. With regard to the parallelization strategy, these libraries may be categorized into three groups: **task-based**: SLATE [28] (OpenMP tasks), DLAF [35] (HPX tasks), DPLASMA [12] (DaGuE

scheduler), or CHAMELEON [3] (StarPU tasks); **static 2D parallel**: MKL [34], Elemental [53], or Cray LibSci [16]; **communication-minimizing 2.5D parallel**: CANDMC [57] and CAPITAL [33]. In the last decade, heavy focus was placed on heterogeneous architectures. Most GPU vendors offer hardware-customized BLAS solvers [50]. Agullo et. al [2] accelerated LU factorization using up to 4 GPUs. Azzam et. al [30] utilize NVIDIA’s GPU tensor cores to compute low-precision LU factorization and then iteratively refine the linear problem’s solution. Moreover, some of the distributed memory libraries support GPU offloading for local computations [28].

12 CONCLUSIONS

In this work, we present a method of analyzing I/O cost of DAAP – a general class of programs that covers many fundamental computational motifs. We show, both theoretically and in practice, that our pebbling-based approach for deriving the I/O lower bounds is **more general**: programs with disjoint array accesses cover a wide variety of applications, **more powerful**: it can explicitly capture inter-statement dependencies, **more precise**: it derives tighter I/O bounds, and **more constructive**: X -partition provides powerful hints for obtaining parallel schedules.

When applying the approach to LU and Cholesky factorizations, we are able to derive new lower bounds, as well as new, communication-avoiding schedules. Not only do they communicate less than state-of-the-art 2D and 3D decompositions – by a factor of up to 1.6× – but most importantly, they outperform existing commercial libraries in a wide range of problem parameters (up to 3× for LU, up to 1.8× for Cholesky). Finally, our code is openly available, offering full ScaLAPACK layout compatibility.

13 ACKNOWLEDGEMENTS

This project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 programme (grant agreement DAPP, no. 678880), EPIGRAM-HS project (grant agreement no. 801039). Tal Ben-Nun is supported by the Swiss National Science Foundation (Ambizione Project #185778). The authors wish to acknowledge the support from the PASC program (Platform for Advanced Scientific Computing), as well as the Swiss National Supercomputing Center (CSCS) for providing computing infrastructure.

REFERENCES

- [1] Alok Aggarwal and S Vitter, Jeffrey. 1988. The input/output complexity of sorting and related problems. *Commun. ACM* 31, 9 (1988), 1116–1127.
- [2] Emmanuel Agullo, Cédric Augonnet, Jack Dongarra, Mathieu Faverge, Julien Langou, Hatem Ltaief, and Stanimire Tomov. 2011. LU factorization for accelerator-based systems. In *2011 9th IEEE/ACS International Conference on Computer Systems and Applications (AICCSA)*. IEEE, 217–224.
- [3] Emmanuel Agullo, Cédric Augonnet, Jack Dongarra, Hatem Ltaief, Raymond Namyst, Samuel Thibault, and Stanimire Tomov. 2010. Faster, Cheaper, Better – a Hybridization Methodology to Develop Linear Algebra Software for GPUs. In *GPU Computing Gems*, Wen mei W. Hwu (Ed.). Vol. 2. Morgan Kaufmann. <https://hal.inria.fr/inria-00547847>
- [4] Emmanuel Agullo, Jack Dongarra, Bilel Hadri, Jakub Kurzak, Julie Langou, Julien Langou, Hatem Ltaief, Piotr Luszczek, and Asim YarKhan. 2011. PLASMA Users' Guide. Parallel Linear Algebra Software for Multicore Architectures. *Rapport technique, Innovative Computing Laboratory, University of Tennessee* (2011).
- [5] Joël Alwen and Vladimir Serbinenko. 2015. High parallel complexity graphs and memory-hard functions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 595–603.
- [6] Edward Anderson, Zhaojun Bai, Christian Bischof, Susan Blackford, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, Alan McKenney, and Danny Sorensen. 1999. *LAPACK Users' guide*. Vol. 9. Siam.
- [7] Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. 2010. Communication-optimal parallel and sequential Cholesky decomposition. *SIAM Journal on Scientific Computing* 32, 6 (2010), 3495–3523.
- [8] Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. 2011. Minimizing communication in numerical linear algebra. *SIAM J. Matrix Anal. Appl.* 32, 3 (2011), 866–901.
- [9] Mohamed-Walid Benabderrahmane, Louis-Noël Pouchet, Albert Cohen, and Cédric Bastoul. 2010. The polyhedral model is more widely applicable than you think. In *International Conference on Compiler Construction*. Springer, 283–303.
- [10] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. 1997. *ScalAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- [11] Uday Bondhugula, Muthu Baskaran, Sriram Krishnamoorthy, J. Ramanujam, Atanas Rountev, and P. Sadayappan. 2008. *Automatic Transformations for Communication-Minimized Parallelization and Locality Optimization in the Polyhedral Model*. Springer Berlin Heidelberg, Berlin, Heidelberg, 132–146. https://doi.org/10.1007/978-3-540-78791-4_9
- [12] G. Bosilca, A. Bouteiller, A. Danalis, M. Faverge, A. Haidar, T. Herault, J. Kurzak, J. Langou, P. Lemarinier, H. Ltaief, P. Luszczek, A. YarKhan, and J. Dongarra. 2011. Flexible Development of Dense Linear Algebra Algorithms on Massively Parallel Architectures with DPLASMA. In *2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum*. 1432–1441.
- [13] John Bruno and Ravi Sethi. 1976. Code generation for a one-register machine. *Journal of the ACM (JACM)* 23, 3 (1976), 502–510.
- [14] J. Choi et al. 1996. ScaLAPACK: a portable linear algebra library for distributed memory computers – design issues and performance. *Comp. Phys. Comm.* (1996).
- [15] Michael Christ, James Demmel, Nicholas Knight, Thomas Scanlon, and Katherine Yelick. 2013. Communication lower bounds and optimal algorithms for programs that reference arrays–Part 1. *arXiv preprint arXiv:1308.0068* (2013).
- [16] Cray. 2020. LibSci: Cray Scientific Libraries. (2020). https://olcf.ornl.gov/software_package/libsci/
- [17] Alain Darté. 1999. On the complexity of loop fusion. In *1999 International Conference on Parallel Architectures and Compilation Techniques (Cat. No. PR00425)*. IEEE, 149–157.
- [18] Mauro Del Ben et al. 2015. Enabling simulation at the fifth rung of DFT: Large scale RPA calculations with excellent time to solution. *Comp. Phys. Comm.* (2015).
- [19] Mauro Del Ben, Jurg Hutter, and Joost VandeVondele. 2013. Electron correlation in the condensed phase from a resolution of identity approach based on the Gaussian and plane waves scheme. *Journal of chemical theory and computation* 9, 6 (2013), 2654–2671.
- [20] James Demmel and Grace Dinh. 2018. Communication-optimal convolutional neural nets. *arXiv preprint arXiv:1802.06905* (2018).
- [21] James Demmel and Alex Rusciano. 2016. Parallelepiped obtaining HBL lower bounds. *arXiv preprint arXiv:1611.05944* (2016).
- [22] Robert H Dennard, Fritz H Gaensslen, Hwa-Nien Yu, V Leo Rideout, Ernest Bassous, and Andre R LeBlanc. 1974. Design of ion-implanted MOSFET's with very small physical dimensions. *IEEE Journal of Solid-State Circuits* 9, 5 (1974), 256–268.
- [23] Grace Dinh and James Demmel. 2020. Communication-Optimal Tilings for Projective Nested Loops with Arbitrary Bounds. *arXiv preprint arXiv:2003.00119* (2020).
- [24] Jack Dongarra, Mathieu Faverge, Hatem Ltaief, and Piotr Luszczek. 2014. Achieving numerical accuracy and high performance using recursive tile LU factorization with partial pivoting. *Concurrency and Computation: Practice and Experience* 26, 7 (2014), 1408–1431.
- [25] Jack Dongarra and Piotr Luszczek. 2011. *TOP500*. Springer US, Boston, MA, 2055–2057. https://doi.org/10.1007/978-0-387-09766-4_157
- [26] V. Elango et al. 2013. *Data access complexity: The red/blue pebble game revisited*. Technical Report.
- [27] Paul Feautrier. 1992. Some efficient solutions to the affine scheduling problem. I. One-dimensional time. *International journal of parallel programming* 21, 5 (1992), 313–347.
- [28] Mark Gates, Jakub Kurzak, Ali Charara, Asim YarKhan, and Jack Dongarra. 2019. SLATE: design of a modern distributed and accelerated linear algebra library. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–18.
- [29] Laura Grigori, James W Demmel, and Hua Xiang. 2008. Communication avoiding Gaussian elimination. In *SC'08: Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*. IEEE, 1–12.
- [30] Azzam Haidar, Stanimire Tomov, Jack Dongarra, and Nicholas J Higham. 2018. Harnessing GPU tensor cores for fast FP16 arithmetic to speed up mixed-precision iterative refinement solvers. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 603–613.
- [31] T. Hoefler et al. 2015. Remote Memory Access Programming in MPI-3. *TOPC* (2015).
- [32] Edward Hutter. [n. d.]. Communication-Avoiding Parallelism-Increasing maTriX fActorization Library. ([n. d.]). <https://github.com/huttered40/capital>
- [33] Edward Hutter and Edgar Solomonik. 2019. Communication-avoiding Cholesky-QR2 for rectangular matrices. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 89–100.
- [34] Intel. 2020. Math Kernel Library. (2020). <https://software.intel.com/en-us/mkl>
- [35] Alberto Invernizzi, Teodor Nikolov, Lara Querciagrossa, and Raffaele Solcà. 2021. Distributed Linear Algebra with (HPX) Futures (forthcoming). In *Proceedings of the Platform for Advanced Scientific Computing Conference*.
- [36] Dror Irony et al. 2004. Communication Lower Bounds for Distributed-memory Matrix Multiplication. *JPDC* (2004).
- [37] Hong Jia-Wei and Hsiang-Tsung Kung. 1981. I/O complexity: The red-blue pebble game. In *STOC*.
- [38] Marko Kabić, Simon Pintarelli, Anton Kozhevnikov, and Joost VandeVondele. 2021. COSTA: Communication-Optimal Shuffle and Transpose Algorithm with Process Relabeling. In *International Conference on High Performance Computing*. Springer, 217–236.
- [39] Richard M Karp. 1988. A survey of parallel algorithms for shared-memory machines. (1988).
- [40] Gokcen Kestor, Roberto Gioiosa, Darren J Kerbyson, and Adolfo Hoisie. 2013. Quantifying the energy cost of data movement in scientific applications. In *2013 IEEE international symposium on workload characterization (IISWC)*. IEEE, 56–65.
- [41] Andreas Knüpfer, Christian Rössel, Dieter an Mey, Scott Biersdorff, Kai Diethelm, Dominic Eschweiler, Markus Geimer, Michael Gerndt, Daniel Lorenz, Allen Malony, Wolfgang E. Nagel, Yury Oleynik, Peter Philippen, Pavel Saviankou, Dirk Schmidl, Sameer Shende, Ronny Tschüter, Michael Wagner, Bert Wesarg, and Felix Wolf. 2012. Score-P: A Joint Performance Measurement Run-Time Infrastructure for Periscope, Scalasca, TAU, and Vampir. In *Tools for High Performance Computing 2011*, Holger Brunst, Matthias S. Müller, Wolfgang E. Nagel, and Michael M. Resch (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 79–91.
- [42] Aravindh Krishnamoorthy and Deepak Menon. 2013. Matrix inversion using Cholesky decomposition. In *2013 signal processing: Algorithms, architectures, arrangements, and applications (SPA)*. IEEE, 70–72.
- [43] Harold W Kuhn and Albert W Tucker. 2014. Nonlinear programming. In *Traces and emergence of nonlinear programming*. Springer, 247–258.
- [44] Thomas D Kühne, Marcella Iannuzzi, Mauro Del Ben, Vladimir V Rybkin, Patrick Seewald, Frederick Stein, Teodoro Laino, Rustam Z Khaliullin, Ole Schütt, Florian Schiffmann, et al. 2020. CP2K: An electronic structure and molecular dynamics software package-Quickstep: Efficient and accurate electronic structure calculations. *The Journal of Chemical Physics* 152, 19 (2020), 194103.
- [45] Grzegorz Kwasniewski, Marko Kabić, Maciej Besta, Joost VandeVondele, Raffaele Solcà, and Torsten Hoefler. 2019. Red-Blue Pebbling Revisited: Near Optimal Parallel Matrix-Matrix Multiplication. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC19)*. Extended technical report available at <https://arxiv.org/abs/1908.09606>.
- [46] Quanquan Liu. 2018. Red-Blue and Standard Pebble Games : Complexity and Applications in the Sequential and Parallel Models.
- [47] L. H. Loomis and H. Whitney. 1949. An inequality related to the isoperimetric inequality. *Bull. Amer. Math. Soc.* 55, 10 (10 1949), 961–962.
- [48] Sanyam Mehta, Pei-Hung Lin, and Pen-Chung Yew. 2014. Revisiting loop fusion in the polyhedral framework. In *Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming*, 233–246.
- [49] Carl D Meyer. 2000. *Matrix analysis and applied linear algebra*. SIAM.
- [50] NVIDIA. 2020. CUSOLVER Reference Guide. (2020). <https://docs.nvidia.com/cuda/cusolver>
- [51] Auguste Olivry, Julien Langou, Louis-Noël Pouchet, P Sadayappan, and Fabrice Rastello. 2020. Automated derivation of parametric data movement lower bounds for affine programs. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 808–822.

- [52] Kazuki Osawa, Yohei Tsuji, Yuichiro Ueno, Akira Naruse, Rio Yokota, and Satoshi Matsuoka. 2019. Large-scale distributed second-order optimization using kronecker-factored approximate curvature for deep convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12359–12367.
- [53] Jack Poulson, Bryan Marker, Robert A Van de Geijn, Jeff R Hammond, and Nichols A Romero. 2013. Elemental: A new framework for distributed memory dense matrix computations. *ACM Transactions on Mathematical Software (TOMS)* 39, 2 (2013), 1–24.
- [54] Gregorio Quintana-Ortí, Enrique S Quintana-Ortí, Robert A Van De Geijn, Field G Van Zee, and Ernie Chan. 2009. Programming matrix algorithms-by-blocks for thread-level parallelism. *ACM Transactions on Mathematical Software (TOMS)* 36, 3 (2009), 1–26.
- [55] Rolf Rabenseifner and Jesper Larsson Träff. 2004. More efficient reduction algorithms for non-power-of-two number of processors in message-passing parallel systems. In *European Parallel Virtual Machine/Message Passing Interface Users' Group Meeting*. Springer, 36–46.
- [56] Ravi Sethi. 1975. Complete register allocation problems. *SIAM journal on Computing* 4, 3 (1975), 226–248.
- [57] Edgar Solomonik. 2014. *Provably efficient algorithms for numerical tensor algebra*. Ph.D. Dissertation. UC Berkeley.
- [58] Edgar Solomonik. 2021. Communication Avoiding Numerical Dense Matrix Computations. (2021). <https://github.com/solomonik/CANDMC>
- [59] Edgar Solomonik et al. 2016. Trade-Offs Between Synchronization, Communication, and Computation in Parallel Linear Algebra computations. *TOPC* (2016).
- [60] E. Solomonik et al. 2017. Scaling Betweenness Centrality using Communication-Efficient Sparse Matrix Multiplication. In *SC*.
- [61] Edgar Solomonik and James Demmel. 2011. Communication-Optimal Parallel 2.5D Matrix Multiplication and LU Factorization Algorithms. In *Euro-Par 2011 Parallel Processing*, Emmanuel Jeannot, Raymond Namyst, and Jean Roman (Eds.). Lecture Notes in Computer Science, Vol. 6853. Springer Berlin Heidelberg, 90–109. https://doi.org/10.1007/978-3-642-23397-5_10
- [62] TOP500 list. 2020. November 2019 TOP500 list. <https://www.top500.org/lists/2019/11/> (April. 2020). (2020).
- [63] D. Ünat, A. Dubey, T. Hoefler, J. Shalf, M. Abraham, M. Bianco, B. L. Chamberlain, R. Cledat, H. C. Edwards, H. Finkel, K. Fuerlinger, F. Hannig, E. Jeannot, A. Kamil, J. Keasler, P. H. J. Kelly, V. Leung, H. Ltaief, N. Maruyama, C. J. Newburn, and M. Pericás. 2017. Trends in Data Locality Abstractions for HPC Systems. *IEEE Transactions on Parallel and Distributed Systems* 28, 10 (2017), 3007–3020.
- [64] Jeffrey Scott Vitter. 1998. External memory algorithms. In *European Symposium on Algorithms*. Springer, 1–25.
- [65] Qinqing Zheng and John D. Lafferty. 2016. Convergence Analysis for Rectangular Matrix Completion Using Burer-Monteiro Factorization and Gradient Descent. *CoRR* (2016).
- [66] Alexandros Nikolaos Ziogas, Tal Ben-Nun, Guillermo Indalecio Fernández, Timo Schneider, Mathieu Luisier, and Torsten Hoefler. 2019. A data-centric approach to extreme-scale ab initio dissipative quantum transport simulations. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–13.