



Tightening I/O Lower Bounds through the Hourglass Dependency Pattern

Lionel Eyraud-Dubois

Inria, Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR
5800
Talence, France
lionel.eyraud-dubois@inria.fr

Julien Langou

University of Colorado Denver, USA
Centre Inria de Lyon, France
julien.langou@ucdenver.edu

Guillaume Iooss

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG
Grenoble, France
guillaume.iooss@inria.fr

Fabrice Rastello

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG
Grenoble, France
fabrice.rastello@inria.fr

ABSTRACT

When designing an algorithm, one cares about arithmetic/computational complexity, but data movement (I/O) complexity plays an increasingly important role that highly impacts performance and energy consumption. For a given algorithm and a given I/O model, scheduling strategies such as loop tiling can reduce the required I/O down to a limit, called the I/O complexity, inherent to the algorithm itself.

The objective of I/O complexity analysis is to compute, for a given program, its minimal I/O requirement among all valid schedules. We consider a sequential execution model with two memories, an infinite one, and a small one of size S on which the computations retrieve and produce data. The I/O is the number of reads and writes between the two memories.

We identify a common “hourglass pattern” in the dependency graphs of several common linear algebra kernels. Using the properties of this pattern, we mathematically prove tighter lower bounds on their I/O complexity, which improves the previous state-of-the-art bound by a parametric ratio. This proof was integrated inside the IOLB automatic lower bound derivation tool.

CCS CONCEPTS

- Theory of computation → Design and analysis of algorithms;
- Software and its engineering → Compilers.

KEYWORDS

I/O complexity, Data Movement Lower Bound

ACM Reference Format:

Lionel Eyraud-Dubois, Guillaume Iooss, Julien Langou, and Fabrice Rastello. 2024. Tightening I/O Lower Bounds through the Hourglass Dependency Pattern. In *Proceedings of the 36th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2024)*, June 17–21, 2024, Nantes, France. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3626183.3659986>

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

SPAA, June 17 - 21, 2024, Nantes, France

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0416-1/24/06...

<https://doi.org/10.1145/3626183.3659986>

1 INTRODUCTION

When designing an algorithm, we usually reason about its computational complexity, to estimate the increase in its execution time, when its problem sizes increase. However, the amount of computation is not the only factor when estimating the performance of an algorithm. The data movement (I/O) also plays an increasingly important role in both performance and the energy consumed by an algorithm.

In order to minimize the amount of data movement of an algorithm, we modify its schedule by using program transformations, such as the tiling transformation [13]. However, finding an optimal schedule is not trivial, due to the size and complexity of the optimization space. The notion of *I/O complexity* of an algorithm indicates the minimal amount of data movement required for any valid schedule of an algorithm. It provides an indication on how far it is possible to optimize the I/O of an algorithm.

However, due to the number of possible valid schedules, it is not possible to directly compute the I/O complexity. Instead, we search for a lower and upper bound on the minimal amount of data movement. To find an upper bound, one only needs to exhibit a valid schedule and compute its I/O cost. However, finding a lower bound requires us to reason over all possible schedules, which requires a mathematical proof based on the data dependency graph of the program.

Several proof techniques exist for deriving lower bounds, such as the *wavefront* [10] or the *K*-partitioning technique [12]. Depending on the shape of the dependencies, some of these techniques might work better than the others: for example, the *wavefront* technique is usually the most effective for stencil-like computation, while the *K*-partitioning technique is usually better for linear-algebraic computation. Other algorithms might require a specialized proof to obtain an asymptotically tight bound (e.g., for the *SYRK* kernel [4]). However, there are still some kernels for which the ratio between their proven lower bound and their best upper bound is parametric. For example, such a ratio is described in Table 1 of [18] for the kernels of the *Polybench* benchmark suite [19].

Some of the proof techniques mentioned above have been automated and implemented in automatic data movement lower bound derivation tools, such as IOLB [18], so that they can be applied to any kernel given as an input.

Contributions. In this paper, we consider several important linear algebra kernels whose dependence graph exhibits a dependency pattern called the *hourglass pattern*. We propose a new proof technique that uses the properties of this pattern to improve the lower bound on the minimal data movement required by these kernels.

In more detail:

- We define the *hourglass pattern*, a pattern of the dependencies of a program, and present its properties.
- We provide a lower bound derivation proof, based on an adaptation of the K -partitioning technique, that tighten the lower bound of a program exhibiting an hourglass pattern. This proof has been fully automatized, inside the tool IOLB [17, 18].
- We present new data movement lower bounds for several linear algebra kernels: Modified GramSchmidt (MGS), QR Householder (GEQR2 and ORG2R in the LAPACK [16] library; also called A2V and V2Q), bidiagonal matrix reduction (GEBD2) and Hessenberg matrix reduction (GEHD2). For all of these kernels, their asymptotic bound was improved by a parametric factor, compared to the bound obtainable by the classical K -partitioning technique. Figure 4 and Figure 5 summarize the new lower bounds found for several linear algebra kernels.
- We also provide tiled orderings for MGS and Householder, resulting in upper bounds that asymptotically match these new lower bounds. This proves the optimality of the new I/O lower bounds.

Outline. In Section 2, we provide some background on the I/O complexity and introduce the K -partitioning method for deriving a lower bound on the minimal amount of data movement of a computation. In Section 3, we present the *hourglass pattern*, a pattern of dependencies whose properties can be used to improve the derived lower bound. In Section 4, we show how to exploit the hourglass pattern to adapt the K -partitioning method, in order to obtain a tighter bound. In Section 5, we list different linear algebra kernels exhibiting an hourglass pattern, and their associated improved lower bound. Additionally, the appendix of the full version of this paper [11] contains the tiled algorithm for two of these kernels. The data movement of these algorithms provides an upper bound to the minimal amount of data movement, which matches asymptotically their new lower bound.

2 BACKGROUND - I/O COMPLEXITY AND THE K-PARTITIONING METHOD

In this section, we present a state-of-the-art method of proof – the K -partitioning method – which infers a lower bound on the amount of data movement needed by a computation.

Memory model and I/O complexity. We consider a simple two-level memory model, composed of (1) a slow memory of unbounded size, and (2) a fast memory of size S . Both memories can transmit data from one to another, as long as the constraint on the size of the small memory is satisfied. When we perform an operation, the data used must be present in the small memory, and the data produced must be committed in the small memory.

We consider a program, which performs a collection of operations organized in *statements*. Each statement SX has multiple *instances* $SX[\vec{i}]$, where \vec{i} is a vector of the surrounding loop indexes.

In this paper, we consider a subclass of programs called *polyhedron* (or *affine*) *programs*. These programs are combinations of nested loops and statements such as: (i) the loop bounds are affine constraints using the surrounding loop indexes and the program parameters (e.g., the sizes of an input array); (ii) the array accesses are affine expression of the surrounding loop indexes and program parameters. All the programs presented in this paper, such as Figure 3, are polyhedral programs. Also, we will call a quantity “*parametric*” when it is a function of the parameters of the program, which are considered as symbolic constants.

A statement instance might depend on the data produced by another instance, for example when the first instance uses a value produced by the second instance. We call this a *dependency* between these two statement instances. These dependencies impose constraints on the order of execution of the program. This order of execution is called the *schedule* or the *ordering*. Frequently, the data consumed/produced by the statements of a program are too big to fit all at once in the small memory, it is thus necessary to *spill*, i.e., to transfer some data back into the slow memory and retrieve it later when needed. However, doing so increases the amount of data movement between both memories.

Given a valid schedule for a program, the *I/O cost* for this program and for this ordering is the amount of data movement required, i.e. the number of data transfers between both memories. The *I/O complexity* of a program is the minimal I/O cost that can be reached by any valid schedule. This quantity is interesting, in particular in the context of an architecture where the transfer of data is the limiting factor for performance. Knowing the minimal amount of data transfers is a good algorithmic indicator to know if it could be theoretically optimized further. However, because we need to find the minimal I/O cost for all possible orderings, the exact I/O complexity is hard to evaluate. Instead, we rely on bounds on the I/O complexity of a program: we can provide a mathematical proof for the lower bound, and exhibit an ordering (i.e., an implementation of a program) that reaches an I/O cost and provides an upper bound.

CDAG and red-white pebble game. When trying to prove a lower bound, one should decide whether redundant computation is allowed or not. The *red-white pebble game*, a variation of the red-blue pebble game of Hong and Kung [12], was introduced by Olivry et al. [18] in order to model the state of the memories during the execution of a program *without* recomputation, which matched the assumptions we make in this paper.

This game is played on the *Computational Directed Acyclic Graph* (CDAG) of the program. This is a directed graph G , where

- the nodes V represent the computation (statement instance) of a program, and
- the edges represent the flow dependencies between the computations of the program.

Notice that the inputs of a program are nodes that do not have incoming edges. The outputs of a program are a subset of nodes $O \subset V$; they might have outgoing edges.

During a *red-white pebble game*, red and white pebbles are placed on the nodes of a CDAG. A *white pebble* represents a computation that was performed, and a *red pebble* represents a computation whose output is currently stored in the small memory. A game follows this set of rules:

- At the start, the only pebbles in the CDAG are white pebbles, placed on the inputs of the program.
- At most S red pebbles can be simultaneously present on the nodes of the CDAG.
- **Spill:** a red pebble can be removed from a node.
- **Compute:** When a node does not have a white pebble, but all its predecessors have red pebbles, then we can place both a white and a red pebble on it.
- **Load:** A red pebble can be added to nodes with a white pebble.
- The game ends when each node has a white pebble on it.

Notice that once a white pebble is placed on a node, it cannot be removed. This prevents recomputation. In order to compute the amount of data movements, we focus on the number of red pebbles added with the **Load** rule during a game. This means that we only focus on the “Load” portion of the data movements and ignore its “Store” part. The resulting bounds are still valid, and because the number of “Load” often dominates the number of “Store”, their tightness should not be strongly impacted. This assumption is identical to the one made in [18].

K-partitioning method. The *K-partitioning* method introduced in the seminal paper of Hong and Kunk [12] is a proof technique that allows to derive a lower bound.

The first idea is to consider a partition of the CDAG and games that play on each set of the partition one by one.

Then, we consider the notion of *K-bounded set*. An *inset* of a set E of nodes of the CDAG, noted $\text{InSet}(E)$, is the set of data used by E but not produced by a computation of E . A *K-bounded set* is a set of nodes E of the CDAG whose inset has a size at most K : $|\text{InSet}(E)| \leq K$. This notion is interesting because an $(S + T)$ -bounded set requires at least T additional data to fit in the small memory (of size S). Thus, even if the small memory is filled with interesting data, we will need at least T load operations to perform the computations of this set. In addition, we assume that our *K*-bounded sets are *convex*: if there is a dependency chain between two points of a *K*-bounded set E , then all the intermediate points must belong to E .

Finally, a *K-partition* is a partition into convex *K*-bounded sets. The idea is to consider all *K*-partitions of a CDAG and to count how many sets are in this partition. By choosing $K = S + T$, we know that there will be at least T loads per set of the partition. Therefore, a lower bound on the number of loads is T times the minimal number of sets in such a partition.

THEOREM 2.1 ((S + T)-PARTITIONING I/O LOWER BOUND [9]). *Let S be the size of the small memory, and for any $T > 0$ let U be the maximal size of a $(S + T)$ -partition. Let V be the set of nodes of the CDAG of the program. Then, a lower bound on the number Q of data movement of the program is:*

$$T \cdot \left\lfloor \frac{|V|}{U} \right\rfloor \leq Q$$

Then, we pick a value of T (which means a value of K) that leads to the tightest lower bound.

To estimate the minimal number of sets in a *K*-partition, we can estimate the maximum size of a set inside this partition. In other words, an upper bound on the size of a *K*-bounded set can be transformed into a lower bound on the amount of data movement required.

*Upper bound on the size of a *K*-bounded set.* An upper bound on the size of a *K*-bounded set E can be obtained by analyzing the dependencies of the program. Indeed, for a polyhedral program, dependencies between its statement instances are associated with affine relations, matching the loop indices of the data producing instance with the data consuming instance. When examining the path of affine dependencies starting from any node of E to a node of the inset of E , we can either obtain a projection or a translation. In both cases, the image of E through these affine functions ϕ can be mapped to geometrical borders or projections of E , and can be associated with parts of $\text{InSet}(E)$. This is the key geometrical intuition that leads us to use the Brascamp-Lieb theorem.

The Brascamp-Lieb theorem is a geometrical way to bound the volume of a set by the volume of its projections, which can be bounded by K .

THEOREM 2.2 (BRASCAMP-LIEB THEOREM [6]). *Let d and d_j be non-negative integers and $\phi_j : \mathbb{Z}^d \mapsto \mathbb{Z}^{d_j}$ be a collection of group homomorphisms for all $1 \leq j \leq m$.*

If we have a collection of coefficients $s_j \in [0, 1]$ such that, for any subgroup $\mathcal{H} \subset \mathbb{Z}^d$:

$$\text{rank}(\mathcal{H}) \leq \sum_{j=1}^m s_j \times \text{rank}(\phi_j(\mathcal{H})).$$

Then, for any non-empty finite set $E \subset \mathbb{Z}^d$:

$$|E| \leq \prod_{j=1}^m |\phi_j(E)|^{s_j}.$$

For example, if we consider a 3D set and the 3 canonical projections $\phi_{j,k}(i, j, k) = (j, k)$, $\phi_{i,k}(i, j, k) = (i, k)$ and $\phi_{i,j}(i, j, k) = (i, j)$, this theorem gives us the following inequality between the volume of E and the area of its faces $\phi_x(E)$:

$$|E| \leq |\phi_{j,k}(E)|^{1/2} \times |\phi_{i,k}(E)|^{1/2} \times |\phi_{i,j}(E)|^{1/2}.$$

As another example, we could consider instead the projections $\phi_i(i, j, k) = (i)$, $\phi_j(i, j, k) = (j)$ and $\phi_k(i, j, k) = (k)$, to obtain the following inequality:

$$|E| \leq |\phi_i(E)| \times |\phi_j(E)| \times |\phi_k(E)|.$$

In the case of a *K*-bounded set, we consider the path of dependencies to automatically derive these projections $\phi \in \Phi$. So, $\phi(E)$ can be mapped to one part of the inset of E , and its size is bounded by K .

```

1  for (k = 0; k < N; k += 1) {
2    nrm = 0.0;
3    for (i = 0; i < M; i += 1)
4      nrm += A[i][k] * A[i][k];
5    R[k][k] = sqrt(nrm);
6
7    for (i = 0; i < M; i += 1)
8      Q[i][k] = A[i][k] / R[k][k];
9
10   for (j = k + 1; j < N; j += 1) {
11     R[k][j] = 0.0;
12     for (i = 0; i < M; i += 1)
13       SR: R[k][j] += Q[i][k] * A[i][j];
14     for (i = 0; i < M; i += 1)
15       SU: A[i][j] = A[i][j] - Q[i][k] * R[k][j];
16   }
17 }
```

Figure 1: Modified Gram-Schmidt - Right-Looking (from Polybench [19]). The input matrix A is of size $M \times N$, and the output of the algorithm are matrices Q (the orthonormalized column vector basis) and R such that $A = QR$. The usual right-looking Gram-Schmidt reuses the matrix A , instead of defining a new matrix Q . SR and SU are labels of two statements, updating R and A .

3 THE HOURGLASS PATTERN

In this section, we describe the intuition of our core contribution. We consider a specific pattern of dependencies, called the *hourglass pattern*, that forces a convex K -bounded set to have a specific shape. We can exploit this property to significantly improve the derived I/O complexity lower bound of programs that exhibit such a pattern.

In the whole section, we use the Modified Gram-Schmidt algorithm as an illustrative example, whose right-looking variant is provided in Figure 1. Using the automatic tool IOLB [18] to apply the K -partitioning method (described in Section 2) to the MGS computation results in a lower bound in $\Omega\left(\frac{MN^2}{\sqrt{S}}\right)$. By using the hourglass pattern, we obtain a more precise lower bound:

$$\frac{M^2N(N-1)}{8(S+M)} \leq Q(MGS)$$

3.1 Intuition of the hourglass pattern

Intuition. Figure 2 presents the main idea of the hourglass pattern. It is a repeating succession of reduction and broadcast statements, such that the number of elements reduced/broadcasted is parametric, thus greater than the cache size S . There are 3 categories of dimensions in this pattern: (a) the dimensions over which the reduction and the broadcast are performed (horizontal axis of Figure 2), (b) the “temporal” dimensions over which the hourglass pattern is repeated (vertical axis of Figure 2), and (c) the neutral dimensions that do not interact with the hourglass pattern.

Running example. The hourglass pattern appears on several linear algebra kernels, including the Modified Gram-Schmidt kernel (Figure 1). The pattern appears between the last two statements: statement SR which updates $R[k][j]$, and statement SU which updates $A[i][j]$. The statement SR is a reduction along the i dimension and uses, in particular, all the values of $A[\cdot][j]$ produced during the

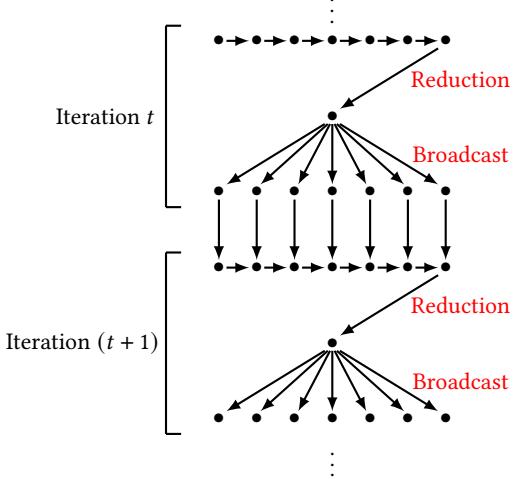


Figure 2: Shape of an hourglass pattern, inside the dependence graph. A node is an instance of a statement of the program, and an edge is a data dependency between two nodes. The t dimension is an external loop surrounding the hourglass.

previous iteration of k . The statement SU broadcasts $R[k][j]$ across the i dimension to update all the $A[\cdot][j]$ of the current iteration of k . Therefore, dimension k is a temporal dimension, dimension i is the reduction/broadcast dimension and dimension j is a neutral dimension. There is exactly one dimension in each category in this example, but in general, there might be several.

Consequences of the hourglass pattern. When considering a K -bounded set over this pattern, we notice that if it spans over several iterations of the temporal dimension t , then the set **must** include all the nodes of the broadcast/reduction in between, due to the convexity property of the set. Therefore, we have two situations:

- Either the K -bounded set spans over several iterations of t and includes all the nodes over the reduction/broadcast dimension. Notice that this is not always possible, depending on the size of the broadcast/reduction dimension and the value of K .
- Or, the K -bounded set is “flat” along the t dimension.

For both situations, we can deduce much stronger constraints on the sizes of the projection $|\phi(E)|$ used in the Brascamp-Lieb theorem (Theorem 2.2). This provides the derivation of an improved lower bound compared to the classical methodology.

3.2 Hourglass pattern - formal definition

In this section, we provide a formal definition of the hourglass pattern.

Preliminary notations. Given an instance for a statement SX , we call an *iteration vector* the tuple of the (integral) values of its surrounding loop indices. The *iteration domain* \mathcal{D}_{SX} of the statement SX is the set of iteration vectors that respect the conditions on the indices of the surrounding loops.

```

1  for (k = 0; k < N; k += 1){
2      norma2 = 0.0;
3      for (i = k + 1; i < M; i += 1){
4          norma2 += A[i][k] * A[i][k];
5      }
6      norma = sqrt(A[k][k] * A[k][k] + norma2);
7      A[k][k]=(A[k][k] > 0)?
8          (A[k][k]+norma):(A[k][k]-norma);
9
10     tau[k] = 2.0 /
11         (1.0 + norma2 / (A[k][k] * A[k][k]));
12
13     for (i = k + 1; i < M; i += 1){
14         A[i][k] /= A[k][k];
15     }
16     A[k][k] = (A[k][k] > 0)?(-norma):(norma);
17
18     for (j = k + 1; j < N; j += 1){
19         tau[j] = A[k][j];
20         for (i = k + 1; i < M; i += 1){
21             SR:   tau[j] += A[i][k] * A[i][j];
22         }
23         tau[j] = tau[k] * tau[j];
24         A[k][j] = A[k][j] - tau[j];
25         for (i = k + 1; i < M; i += 1){
26             SU:   A[i][j] = A[i][j] - A[i][k] * tau[j];
27         }
28     }
29 }
```

Figure 3: QR Householder computation - Part A2V (LAPACK routine GEQR2).

As mentioned above, in general, there might be several reduction or temporal dimensions. Thus, we will consider sets of dimensions, and represent their iteration as a vector $\vec{i} = (i_1, i_2, \dots)$. Given an iteration \vec{k} , we write $\vec{k} + \vec{1}$ to represent the next valid lexicographic value of \vec{k} . We extend this notation to $\vec{k} + \vec{n}$ where n is an integer.

The hourglass pattern. Considering a statement S of the CDAG of a program, the *hourglass pattern* is a pattern of dependencies with the following properties:

- *Partitioning of the dimensions.* The dimensions of the statement SX can be partitioned into 3 groups: (i) the temporal dimensions \vec{k} , (ii) the reduction/broadcast dimensions \vec{i} , and (iii) the neutral dimensions \vec{j} . For simplicity of the presentation, we assume that \vec{k} are the first/outer dimensions and that \vec{i} are the last/inner dimensions.
- *Path in the dependence graph.* For any valid value of \vec{i} and \vec{i}' , there is a dependency chain between the instances $SX[\vec{k}, \vec{j}, \vec{i}]$ and $SX[\vec{k} + \vec{1}, \vec{j}, \vec{i}']$.
- *Large width of the hourglass.* Let us consider W , the number of statement instances on all the dependency chains between $SX[\vec{k}, \vec{j}, \vec{i}]$ and $SX[\vec{k} + \vec{2}, \vec{j}, \vec{i}]$. This expression depends on the parameters of the program, and cannot be bounded by a constant value.

Notice that, to have such a chain of dependencies, it must include a reduction and a broadcast. The dimensions \vec{i} are the dimensions which are reduced on and broadcasted over. The dimensions \vec{k} are the dimensions that are incremented by a constant factor when

looping along this loop. So, once a path is found, partitioning the dimensions should be unambiguous, if this condition is also satisfied.

The automatic detection of such an hourglass pattern has been implemented inside the IOLB tool, using a polyhedral library [24, 25].

Examples. For the MGS computation (Figure 1), we consider the statement SU and the cycle of dependencies going through the SR . We confirm that there are $2M$ statement instances inside a dependency chain between two instances $SU[k, j, i]$ and $SU[k + 2, j, i]$: $SR[k + 1, j, \cdot]$ and $SU[k + 1, j, \cdot]$. The same reasoning would hold if we considered the statement SR instead of SU .

For the A2V QR Householder computation (Figure 3), we consider the statement SU , and the cycle of dependencies going through the SR statement. There are $(M - k)$ statement instances $SR[k, j, \cdot]$, inside a dependency chain between two instances $SU[k, j, i]$ and $SU[k + 1, j, i]$.

4 LOWER BOUND PROOF USING THE HOURGLASS PATTERN

In this section, we show how to exploit the hourglass pattern of a program to derive a tighter lower bound on the data movement. We will use the Modified Gram-Schmidt program (Figure 1) as a running example to illustrate our proof.

Once again, this proof has been integrated inside the automatic data movement lower bound derivation tool, IOLB [18], and is applied when an hourglass pattern is detected.

Preliminary notations. Let us denote with E a set of integral iteration vectors. The cardinality of such a set $|E|$ is the number of integer points inside this set.

Given a dimension k , $\phi_k : (i, j, k) \mapsto (k)$ is the projection to the dimension k . This notation can be extended to several dimensions. For example, $\phi_{j,k} : (i, j, k) \mapsto (j, k)$.

Given a set E , a dimension k and a value k_0 , $E_{k=k_0}$ is a slice of E , i.e., the set of points of E whose value along the dimension k is k_0 : $E_{k=k_0} = \{(i, j, k) \in E \mid k = k_0\}$. We extend this notation to several dimensions, as in $E_{k=k_0, j=j_0}$. In the rest of the paper, we use a compact notation for slices, by omitting the dimension when it is not ambiguous, like E_{k_0, j_0} .

Starting point and intuition. The goal of the proof is to find an upper bound on the size of a K -bounded set, to transform it into a lower bound on the data volume required by a portion of the program. So, we start the proof with an arbitrary convex K -bounded set E containing instances of the broadcast statement $SX[\vec{k}, \vec{j}, \vec{i}]$, which we assume is part of an hourglass pattern. Just like in the classical proof, an analysis of the dependencies of the program yields a set Φ of projections $\phi_{\vec{x}}$ for some dimensions \vec{x} , which project on the inset of E .

The classical proof involves applying the Brascamp-Lieb theorem on the K -bounded set E , by bounding the size $|\phi_{\vec{x}}(E)|$ of each of these projections by K .

In our proof, we split E into two parts (Section 4.1) and we adapt the set of projections when we apply the Brascamp-Lieb theorem on each of these parts (Sections 4.2 and 4.3) to obtain more precise bounds.

Running example. In the case of MGS, by following the chain of dependencies associated with the accesses $A[i][j]$, $Q[i][k]$ and $R[k][j]$ of statement SU , we infer that the projections are $\phi_{i,j}$, $\phi_{i,k}$ and $\phi_{k,j}$. When following the classical proof, the application of the Brascamp-Lieb theorem results in the following inequality:

$$|E| \leq |\phi_{i,j}(E)|^{\frac{1}{2}} \cdot |\phi_{i,k}(E)|^{\frac{1}{2}} \cdot |\phi_{k,j}(E)|^{\frac{1}{2}} \leq K^{3/2}$$

We prove in Section 4.4 a tighter upper bound: $|E| \leq \frac{K^2}{M} + 2K$.

4.1 Part 1 - Decomposition of the K-bounded set

In the first part of the proof, we decompose E into the union of two fragments: (i) I' which has volume along the \vec{k} dimensions, and (ii) F which is flat along the \vec{k} dimensions. Both parts have their own upper bound, obtained with different reasoning, that is described in Section 4.2 and Section 4.3.

Decomposition of E . We consider the number of different values of \vec{k} for a given value of \vec{j} :

$$Tick_{\vec{j}} = \{\vec{k} \mid \exists \vec{i}, SX[\vec{k}, \vec{j}, \vec{i}] \in E_{\vec{j}}\}.$$

We split E into two sets:

- $E' = \bigcup_{\vec{j} \in J_{3+}} E_{\vec{j}}$ where $J_{3+} = \{\vec{j} \mid 3 \leq Tick_{\vec{j}}\}$, and
- $E'' = \bigcup_{\vec{j} \in J_{12}} E_{\vec{j}}$ where $J_{12} = \{\vec{j} \mid 1 \leq Tick_{\vec{j}} \leq 2\}$.

The intuition between this separation is that (i) E' contains the connected components which need to include an entire line of statement instances along the \vec{i} dimensions, and (ii) E'' contains the connected components which are “flat” along the \vec{k} dimensions.

We start by focusing on E' , to show that these components need to include a parametric number of iterations along the \vec{i} dimension, using the hourglass pattern.

LEMMA 4.1 (STRUCTURE OF E'). *Given some $\vec{j} \in \phi_{\vec{j}}(E')$, let us consider the slice $E'_{\vec{j}}$ along the dimensions \vec{j} . Let us consider:*

- $\overrightarrow{k_{\min}(\vec{j})} = \min_{\vec{k}} \{\vec{k} \mid SX[\vec{k}, \vec{j}, \vec{i}] \in E'_{\vec{j}}\}$,
- $\overrightarrow{k_{\max}(\vec{j})} = \max_{\vec{k}} \{\vec{k} \mid SX[\vec{k}, \vec{j}, \vec{i}] \in E'_{\vec{j}}\}$,
- \vec{a} , any subset of the indices of \vec{i} ,
- $|\phi_{\vec{a}}(\mathcal{D}_S)| \geq W_{\vec{a}}$, a lower bound on the size of the projection of \mathcal{D}_S , the iteration domain of the statement SX .

Then:

- (1) $E'_{\vec{j}}$ is a connected component.
- (2) For all $\overrightarrow{k_{\min}(\vec{j})} <_{lex} \vec{k} <_{lex} \overrightarrow{k_{\max}(\vec{j})}$, $|\phi_{\vec{a}}(E'_{\vec{j}, \vec{k}})| \geq W_{\vec{a}}$.

PROOF. (1) Because SX satisfies the hourglass pattern properties, for any two instances $SX[\vec{k}, \vec{j}, \vec{i}]$ and $SX[\vec{k}', \vec{j}, \vec{i}']$ in $E'_{\vec{j}}$ where $\vec{k} <_{lex} \vec{k}'$, we can prove that there is a chain of dependencies from one statement instance to another. Therefore, using the convexity property of E , we conclude that $E'_{\vec{j}}$ is a connected component.

(2) By considering one instance of index $\vec{k} = \overrightarrow{k_{\min}}$ and another of index $\vec{k} = \overrightarrow{k_{\max}}$, we can show that all the $SX[\vec{i}, \vec{j}, \vec{k}]$ are in the

middle of a dependency chain between the two of them. Therefore, by projecting the \vec{i} on the subset of dimension \vec{a} , we conclude:

$$\forall \overrightarrow{k_{\min}(\vec{j})} <_{lex} \vec{k} <_{lex} \overrightarrow{k_{\max}(\vec{j})}, |\phi_{\vec{a}}(E'_{\vec{j}, \vec{k}})| \geq |\phi_{\vec{a}}(\mathcal{D}_S)| \geq W_{\vec{a}}.$$

□

We consider $E' = I' \uplus B'$ defined by:

- $I' = \bigcup_{\vec{j} \in J_{3+}} (E_{\vec{j}} - E_{\overrightarrow{j, k_{\min}(\vec{j})}} - E_{\overrightarrow{j, k_{\max}(\vec{j})}})$: the “inside” of the connected components of E' , according to dimensions \vec{k} .
- $B' = \bigcup_{\vec{j} \in J_{12}} (E_{\overrightarrow{j, k_{\min}(\vec{j})}} \cup E_{\overrightarrow{j, k_{\max}(\vec{j})}})$: the “boundaries” of the connected components of E' , according to dimensions \vec{k} .

We define $F = B' \uplus E''$ the flat parts of E , and we adapt our previous decomposition of E to obtain the desired decomposition:

$$E = E' \uplus E'' = I' \uplus F.$$

4.2 Part 2 - Bound on the size of I'

Let us focus on the upper bound of the size of I' , using Lemma 4.1. The goal is to have tighter bounds on the projections involving some of the reduce/broadcast dimensions \vec{i} , to use them with the Brascamp-Lieb theorem, instead of the classical “ $\leq K$ ” bound.

LEMMA 4.2 (BOUNDS ON THE SIZE OF SOME OF THE PROJECTIONS OF I'). *Let us consider a projection $\phi_{\vec{x}, \vec{a}} \in \Phi$, where \vec{a} is a subset of \vec{i} , and \vec{x} is a subset of (\vec{j}, \vec{k}) . Assume that we have a lower bound $|\phi_{\vec{a}}(\mathcal{D}_S)| \geq W_{\vec{a}}$ on the size of the projection of the iteration domain \mathcal{D}_S of statement S . Then:*

$$|\phi_{\vec{x}}(I')| \leq \frac{K}{W_{\vec{a}}}.$$

PROOF. Because I' is a subset of a K -partition, then we have $|\phi_{\vec{a}, \vec{x}}(I')| \leq K$.

By slicing I' along the dimensions of \vec{x} , we also have:

$$|\phi_{\vec{a}, \vec{x}}(I')| = |\phi_{\vec{a}, \vec{x}}(\bigcup_{\vec{x}''} I'_{\vec{x}''})| = \sum_{\vec{x}''} |\phi_{\vec{a}, \vec{x}}(I'_{\vec{x}''})|.$$

Furthermore:

$$\begin{aligned} |\phi_{\vec{a}, \vec{x}}(I'_{\vec{x}''})| &\geq |\phi_{\vec{a}}(I'_{\vec{x}''})| && \text{(a slice along } \vec{x} \text{ only has a single point to be projected along } \vec{x}) \\ &\geq |\phi_{\vec{a}}(I'_{\overrightarrow{j'', k''}})| && (I'_{\overrightarrow{j'', k''}} \subset I'_{\vec{x}''}, \text{ for some } (\overrightarrow{j''}, \overrightarrow{k''}) \\ &&& \text{matching } \vec{x}'' \text{ on its dimensions}) \\ &\geq W_{\vec{a}}. && \text{(by Lemma 4.1)} \end{aligned}$$

Therefore: $|\phi_{\vec{a}, \vec{x}}(I')| \geq W_{\vec{a}} \times |\phi_{\vec{x}}(I')|$. In other words:

$$|\phi_{\vec{x}}(I')| \leq \frac{|\phi_{\vec{a}, \vec{x}}(I')|}{W_{\vec{a}}} \leq \frac{K}{W_{\vec{a}}}.$$

□

To obtain a bound on $|I'|$, we use the set of projections Φ , modified in the following way:

- We add a projection on \vec{i} whose bound is: $|\phi_{\vec{i}}(I')| \leq W$, where W is the width of the hourglass.

- When a projection on (\vec{x}, \vec{a}) shares some of its dimensions \vec{a} with \vec{i} , we use the projection on \vec{x} instead, and the bound given by Lemma 4.2.
- The rest of the projections, which do not involve dimensions of \vec{i} , are unchanged and associated with a classical upper bound $|\phi_{\vec{x}}(I')| \leq K$.

Then, we apply the Brascamp-Lieb theorem, while optimizing the values of the power s of the size of the projections $|\phi(E)|$. Notice that some projections can be filtered out through this optimization process ($s = 0$), to obtain a tighter bound on I' .

Running example. We notice that Φ contains two projections $\phi_{i,j}$ and $\phi_{i,k}$, both of them using the dimension i and another dimension. Therefore, from Lemma 4.2, we have the following bounds:

$$|\phi_j(I')| \leq \frac{K}{M} \quad \text{and} \quad |\phi_k(I')| \leq \frac{K}{M}.$$

We apply the Brascamp-Lieb theorem on I' , using the projections on i , on j instead of a projection on (i, j) , and on k instead of a projection on (i, k) , each with coefficient 1:

$$|I'| \leq |\phi_i(I')| \times |\phi_j(I')| \times |\phi_k(I')|.$$

By using our specialized bounds, we have:

$$|I'| \leq M \times \frac{K}{M} \times \frac{K}{M} = \frac{K^2}{M}.$$

4.3 Part 3 - Bound on the size of F

We now focus on the bound of the remaining part of E , i.e., $F = (B' \uplus E'')$. Our goal is to exploit the fact that this set is “flat” along the \vec{k} dimensions, to improve the upper bounds on the projections used in the Brascamp-Lieb application to F . In particular, our proof will consider each $F_{\vec{j}}$ separately, the slice of F for a given value \vec{j} . Then, by picking a list of well-chosen projections for the Brascamp-Lieb theorem, we obtain interesting bounds on the size of $F_{\vec{j}}$, that are finally summed to obtain our bound on F .

We recall that E' and E'' contains the \vec{j} -slices of E for the values of \vec{j} such that $\text{Tick}_{\vec{j}}$ are respectively at least 3 and at most 2. Since $B' \subset E'$, B' and E'' do not share the same set of \vec{j} . So we get a “flatness bound”: $\forall \vec{j} \in \phi_{\vec{j}}(F), |\phi_{\vec{k}}(F_{\vec{j}})| \leq 2$.

In addition, we notice that because $F_{\vec{j}} \subset E$, for any projection $\phi \in \Phi$, $|\phi(F_{\vec{j}})| \leq |\phi(E)| \leq K$.

In this part of the proof, instead of focusing on F , we apply the Brascamp-Lieb theorem to the slice $F_{\vec{j}}$.

As in Section 4.2, we start with the same list of projections Φ to the in-set of E , obtained by inspecting the chain of dependencies of the considered statement. Then, we customize this list of projections to exploit the properties of $F_{\vec{j}}$:

- We add a projection on \vec{k} and the flatness bound $|\phi_{\vec{k}}(F_{\vec{j}})| \leq 2$.
- We identify a projection that involves a non-empty subset of the dimensions of \vec{j} : after applying the Brascamp-Lieb theorem, we will not try to immediately use an upper bound for this projection. Let us call $\phi_{\vec{w}}$ this projection.
- The remaining projections are left alone, and associated with the classical upper bound $|\phi_{\vec{x}}(F_{\vec{j}})| \leq |\phi_{\vec{x}}(F)| \leq K$.

At that point, we have obtained an upper bound on $|F_{\vec{j}}|$ of the following form:

$$|F_{\vec{j}}| \leq e \times |\phi_{\vec{w}}(F_{\vec{j}})|.$$

where e is a parametric expression using the parameter K and independent of the value of \vec{j} . Notice that the Brascamp-Lieb power above $|\phi_{\vec{w}}|$ must be equal to 1, due to the fact that this is the only projection involving the \vec{j} dimensions.

Let us consider the collection of projected sets $\phi_{\vec{w}}(F_{\vec{j}})$. Two of these projected sets are either (i) identical, along the dimensions of \vec{j} which are not present in \vec{w} , or (ii) disjoint. Because the distinct values of these projections are always a subset of the inset of E , of size K , the sum of the size of the disjoint union of these distinct values is also bounded by K . Let us call R the number of values that can be taken by the dimensions of \vec{j} which are not in \vec{w} . This is also the maximum number of times a $\phi_{\vec{w}}(E)$ projects on the same value. Notice that $R = 1$ when \vec{w} covers all the dimensions of \vec{j} , and that R can be a parametric expression in general. So, we have:

$$\sum_{\vec{j} \in \phi_{\vec{j}}(F)} |\phi_{\vec{w}}(F_{\vec{j}})| \leq R \times K.$$

Combining all these observations, we obtain the following upper bound on the size of F :

$$\begin{aligned} |F| &= \sum_{\vec{j} \in \phi_{\vec{j}}(F)} |F_{\vec{j}}| \leq \sum_{\vec{j} \in \phi_{\vec{j}}(F)} e \times |\phi_{\vec{w}}(F_{\vec{j}})| \\ &\leq e \times \sum_{\vec{j} \in \phi_{\vec{j}}(F)} |\phi_{\vec{w}}(F_{\vec{j}})| \leq e \times R \times K. \end{aligned}$$

Running example. In the MGS case, the flatness bound yields: $\forall j, |\phi_k(F_j)| \leq 2$. We can apply the Brascamp-Lieb theorem on F_j , with the projection $\phi_{i,j}$, together with the projection ϕ_k :

$$|F_j| \leq |\phi_k(F_j)| \times |\phi_{i,j}(F_j)| \leq 2 \times |\phi_{i,j}(F_j)|.$$

And since dimension j is included in the dimensions (i, j) of the projection, we have $R = 1$ and therefore:

$$|F| = \sum_{j \in \phi_j(F)} |F_j| \leq 2 \times \sum_{j \in \phi_j(F)} |\phi_{i,j}(F_j)| \leq 2K.$$

4.4 Part 4 - Wrapping things up

We have found in Section 4.2 and Section 4.3 an upper bound of the size of the two parts of E . We simply sum them together to obtain an upper bound of the size of E , then apply Theorem 2.1 to deduce a lower bound on the data movement.

Running example. Thanks to the above results, we can obtain lower bounds on the data movement of MGS.

THEOREM 4.3 (LOWER BOUNDS FOR MGS). *The communication volume Q for the MGS algorithm on a $M \times N$ matrix can be bounded as follows:*

$$\frac{M^2 N(N-1)}{8(S+M)} \leq Q$$

Furthermore, if $S \leq M$, we also have:

$$\frac{(M-S)N(N-1)}{4} \leq Q$$

PROOF. From the previous results, we have:

$$|E| = |I'| + |F| \leq \frac{K^2}{M} + 2K.$$

Then, by using Theorem 2.1 with $K = 2S$:

$$(K - S) \times \frac{MN(N - 1)}{2 \cdot \left(\frac{K^2}{M} + 2K\right)} = \frac{M^2N(N - 1)}{8(S + M)} \leq Q$$

Due to Lemma 4.1 and because we have at least an input dependency involving the dimension i , $|InSet(E')| > M$. So, if we have $S \leq M$, then E' must be empty. Therefore, $E = F$, and we can use only the second part of the bound: $|E| \leq 2K$.

Using Theorem 2.1 again, but this time with $K = M$, we obtain:

$$(K - S) \cdot \frac{MN(N - 1)}{2 \times 2K} = (M - S) \cdot \frac{N(N - 1)}{4} \leq Q$$

□

5 EXPERIMENTAL RESULTS - NEW LOWER BOUNDS

In this section, we report the data movement lower bounds generated by IOLB for four kernels exhibiting an hourglass pattern. We compare the results using our technique (new bound) with those obtained without it (old bound). These kernels are:

- Modified Gram-Schmidt (Figure 1), already used as a running example in Section 4.
- QR Householder algorithm: both its A2V (Figure 3) and V2Q parts (left-looking variants of respectively the GEQR2 and ORG2R subroutines in LAPACK [16]).
- Reduction to a bidiagonal matrix (GEBD2 subroutine)
- Reduction to a Hessenberg matrix (GEHD2 subroutine)

Figure 5 summarizes all the newly found full lower bounds, and Figure 4 focuses on their leading term, to emphasize their improvements. More precisely:

- Section 5.1 contains an asymptotic analysis of the MGS lower bound.
- Section 5.2 presents the lower bound for both QR Householder parts, and the GEBD2 computation.
- Section 5.3 presents the lower bounds for the GEHD2 computation.

The appendix of the full version of this paper [11] contains a description of a tiled algorithm for MGS and for HH A2V and the computation of their amount of data movement. This provides an upper bound to the minimal amount of data movement required by these algorithms, which matches asymptotically the provided lower bound.

5.1 MGS - Asymptotic analysis

In this section, we analyze the bound obtained in Theorem 4.3 for MGS, by specializing it for different ordering of M and S :

- If $S \leq M/2$, we have $M/2 \leq M - S$, so that the second bound yields:

$$\frac{MN^2}{8} = \Omega(MN^2) \leq Q$$

| Kernel | Old bound [18] | New bound (hourglass) |
|-----------|--|---|
| MGS | $\Omega\left(\frac{MN^2}{\sqrt{S}}\right)$ | $\Omega\left(\frac{M^2N(N-1)}{S+M}\right)$ |
| QR HH A2V | $\Omega\left(\frac{MN^2}{\sqrt{S}}\right)$ | $\Omega\left(\frac{MN^2(N-M)}{N-M-S}\right)$ |
| QR HH V2Q | $\Omega\left(\frac{MN^2}{\sqrt{S}}\right)$ | $\Omega\left(\frac{MN^2(N-M)}{N-M-S}\right)$ |
| GEBD2 | $\Omega\left(\frac{MN^2}{\sqrt{S}}\right)$ | $\Omega\left(\frac{MN^2(M-N+1)}{8(S+M-N+1)}\right)$ |
| GEHD2 | $\Omega\left(\frac{N^3}{\sqrt{S}}\right)$ | $\Omega\left(\frac{N^4}{N+2S}\right)$ |

Figure 4: Summary of the new asymptotic data movement lower bounds.

- If $M/2 \leq S$, we have $S + M \leq 3S$, so that the first bound becomes:

$$\frac{M^2N^2}{24S} = \Omega\left(\frac{M^2N^2}{S}\right) \leq Q$$

The first result matches the amount of data movements obtained by the classical ordering of the MGS algorithm, as presented at the end of Section 3.1. Both the algorithm and the bound are thus asymptotically optimal when S is small.

Demmel et al. [7] propose a tiled ordering for the case $2M \leq S$ (in Section F.2 of their paper), that achieves an amount of data movement $O(M^2N^2/S)$, thus matching asymptotically the second upper bound. Both of these bounds are thus optimal up to a constant factor. For reference, the tiled ordering is detailed in the Appendix of the full version of this paper [11], together with the proof on its amount of data movement.

We can compare these bounds with the lower bound returned by the classical hourglass-less proof, whose asymptotic bound is $\Omega\left(\frac{MN^2}{\sqrt{S}}\right)$. Our first bound for small values of S is stronger by a factor of $\Theta(\sqrt{S})$. By writing our second bound as $\Omega\left(\frac{M}{\sqrt{S}} \cdot \frac{MN^2}{\sqrt{S}}\right)$, we see that our bound is stronger by a factor of $\Theta\left(\frac{M}{\sqrt{S}}\right)$. Since the input matrix has size $M \times N$, we can assume that $S < MN$, otherwise, the whole matrix fits in the cache and there is no need to minimize data transfers. Besides, we know that $N \leq M$, which leads to $S < M^2$. We can conclude that $\frac{M}{\sqrt{S}} > 1$: our bound is asymptotically at least as strong as the previous bound.

The results of Theorem 4.3 can also be presented differently, to improve the constant in front of the dominant term. Indeed:

If $S \ll M$, the first bound yields

$$\frac{MN^2}{4} \leq Q,$$

Similarly, if $M \ll S$, the second result becomes

$$\frac{M^2N^2}{8S} \leq Q.$$

5.2 Householder QR factorization and GEBD2

In this section, we present the bounds to both parts of the Householder QR factorization (LAPACK routines GEQR2 et ORG2R). The computation of the first part (A2V) is given in Figure 3 and the computation of the second part (V2Q) in Figure 6.

By using the hourglass reasoning, we obtain the following new lower bounds for both kernels.

THEOREM 5.1 (LOWER BOUNDS FOR HH - PART A2V). *The communication volume Q for the A2V part of the HH algorithm on a $M \times N$*

| Kernel | Old bound [18] | New bound (hourglass) |
|-----------|---|---|
| MGS | $\frac{2M+3MN+MN^2}{\sqrt{S}} + 5M - MN + \frac{7N-N^2}{2} - S - 6$ | $\frac{N^2M^2+2M^2-3NM^2}{8(M+S)} + 5M - MN + \frac{7N-N^2}{2} - S - 6$ |
| QR HH A2V | $\frac{3MN^2+6M+7N-N^3-9MN-6}{3\sqrt{S}} + 5M - MN + 5N - S - 13$ | $\frac{3MN^2-9MN+7N+6M-6-N^3}{24(1-\frac{S}{N-M})} + 5M - MN + 5N - S - 13$ |
| QR HH V2Q | $\frac{3MN^2-N^3+6M+7N-9MN-6}{3\sqrt{S}} + 2M + 2N + \frac{N-N^2}{2} - S - 4$ | $\frac{3MN^2-N^3+6M+7N-9MN-6}{24(1+\frac{S}{M-N})} + 2M + 2N + \frac{N-N^2}{2} - S - 4$ |
| GEBD2 | $\frac{3MN^2-N^3-9MN+6M+7N-6}{3\sqrt{S}} + 5N + 5M - MN - S - 13$ | $\frac{3MN^2-N^3+3N^2-15MN+4N+18M-12}{24(1+\frac{S}{1+N})} + 5N + 7M - MN - S - 18$ |
| GEHD2 | $\frac{5N^3-30N^2+55N-30}{3\sqrt{S}} + \frac{69N-9N^2}{2} - 3 * S - 56$ | $\frac{N^3-6N^2+11N-6}{12(1+\frac{S}{N-M-1})} - N^2 + 12N - S - 19$ |

Figure 5: Data movement lower-bounds (with constants) automatically derived by IOLB [18] without/with hourglass detection. In GEHD2's new bound, a new parameter M is introduced, corresponding to the place where we split the outer loop. Depending on S and N , it can be instantiated with a different parametric expression (cf Section 5.3).

```

1  for (k = N - 1; k > -1; k = -1) {
2    for (j = k + 1; j < N; j += 1) {
3      tau[j] = 0.0;
4      for (i = k + 1; i < M; i += 1) {
5        SR:   tau[j] += A[i][k] * A[i][j];
6      }
7    }
8    for (j = k + 1; j < N; j += 1) {
9      ST:   tau[j] *= tau[k];
10   }
11   A[k][k] = 1.0 - tau[k];
12   for (j = k + 1; j < N; j += 1) {
13     A[k][j] = -tau[j];
14   }
15   for (j = k + 1; j < N; j += 1) {
16     for (i = k + 1; i < M; i += 1) {
17       SU:   A[i][j] -= A[i][k] * tau[j];
18     }
19   }
20   for (i = k + 1; i < M; i += 1) {
21     A[i][k] = -A[i][k] * tau[k];
22   }
23 }
```

Figure 6: QR Householder computation - Part V2Q (LAPACK subroutine ORG2R). We assume that $M \geq N$.

matrix, with $M > N$, can be bounded as follows:

$$\frac{(3M - N)N^2(M - N)^2}{24(MS + (M - N)^2)} \leq Q$$

If $M \gg N$, then the bound becomes:

$$\frac{M^2N(N - 1)}{8(S + M)} \leq Q$$

THEOREM 5.2 (LOWER BOUND FOR HH - PART V2Q). The communication volume Q for the HH V2Q algorithm on a $M \times N$ algorithm, with $M > N$, can be bounded as follows:

$$\frac{N(N - 1)(3M - N - 1)(M - N)^2}{24((M - N)^2 + SM)} \leq Q$$

When $M \gg N$, this bound becomes:

$$\frac{N(N - 1)M^2}{8(S + M)} \leq Q$$

One interesting detail of the proof concerns the detection of the hourglass, and the criteria on its size (third criterion introduced in Section 3.2). For the MGS computation, the size of its hourglass was constant and equal to M . In the case of both Householder computations, the size of their hourglass is parametrized by the outer loop iteration value k , and is equal to $(M - 1 - k)$. Its smallest value happens for $k = N - 1$, and is $(M - N)$. By using this value in Lemma 4.1, we can derive the announced lower bound.

The lower bound proof of the GEBD2 subroutine is similar to both Householder proofs.

THEOREM 5.3 (LOWER BOUNDS FOR GEBD2). *The communication volume Q for the GEBD2 subroutine on a $M \times N$ matrix, with $M \geq N$, can be bounded as follows:*

$$\frac{MN^2(M - N + 1)}{8(S + M - N + 1)} \leq Q$$

If $M \gg N$, then the bound becomes:

$$\frac{M^2N^2}{8(S + M)} \leq Q$$

5.3 Hessenberg matrix reduction

In this section, we present the bound for the Hessenberg matrix factorization kernel (GEHD2 kernel in LAPACK), whose code is in Figure 7. By using the hourglass reasoning, we obtain the following new lower bound for the GEHD2 kernel.

THEOREM 5.4 (LOWER BOUND FOR GEHD2). *The communication volume Q for the Hessenberg matrix factorization algorithm on a $N \times N$ algorithm can be bounded as follows:*

$$\frac{1}{12} \cdot \frac{N^4}{N + 2S} \leq Q$$

When $N \gg S$,

$$\frac{N^3}{24} \leq Q$$

As in Section 5.2, the size of the hourglass depends on the value of iteration j of the outermost loop (temporal dimension), and is equal to $(N - 2 - j)$, where $0 \leq j < N - 2$. This means that its minimal value is 1, which causes an issue with our proof.

A way to solve this problem is to introduce a new parameter $M < N - 2$ and to consider a loop splitting of the outermost temporal

```

1  for (j = 0; j < n - 2; j+= 1) {
2    norma2 = 0.0;
3    for (i = j + 2; i < n; i+= 1) {
4      norma2 += A[i][j] * A[i][j];
5    }
6    norma = sqrt( A[j + 1][j] * A[j + 1][j] + norma2);
7    A[j + 1][j] = (A[j + 1][j] > 0)?
8      (A[j + 1][j] + norma) : (A[j + 1][j] - norma);
9    tau = 2.0 / (1.0 + norma2 / (A[j + 1][j] * A[j + 1][j]));
10   for (i = j + 2; i < n; i+= 1) {
11     A[i][j] /= A[j + 1][j];
12   }
13  A[j + 1][j] = (A[j + 1][j] > 0)? (-norma) : (norma);
14  for (i = j + 1; i < n; i+= 1) {
15    tmp[i] = A[j + 1][i];
16    for (k = j + 2; k < n; k+= 1) {
17      tmp[i] += A[k][j] * A[k][i];
18    }
19  }
20  for (i = j + 1; i < n; i+= 1) {
21    tmp[i] *= tau;
22  }
23  for (i = j + 1; i < n; i+= 1) {
24    A[j + 1][i] -= tmp[i];
25  }
26  for (i = j + 2; i < n; i+= 1) {
27    for (k = j + 1; k < n; k+= 1) {
28      A[i][k] -= A[i][j] * tmp[k];
29    }
30  }
31  for (i = 0; i < n; i+= 1) {
32    tmp[i] = A[i][j + 1];
33    for (k = j + 2; k < n; k+= 1) {
34      tmp[i] += A[i][k] * A[k][j];
35    }
36  }
37  for (i = 0; i < n; i+= 1) {
38    tmp[i] *= tau;
39  }
40  for (i = 0; i < n; i+= 1) {
41    A[i][j + 1] -= tmp[i];
42  }
43  for (i = 0; i < n; i+= 1) {
44    for (k = j + 2; k < n; k+= 1) {
45      A[i][k] -= tmp[i] * A[k][j];
46    }
47  }
48 }

```

Figure 7: Hessenberg matrix factorization of a $N \times N$ matrix (LAPACK subroutine GEHD2)

dimension j at iteration M . First, notice that a loop splitting does not change the dependencies of a program, thus the lower bound of the split version will apply to the original one.

The first half of the split program satisfies the hourglass pattern and has a minimum size of its hourglass of $(N - M - 1)$. Thus, as long as this quantity is not bounded by a constant, the hourglass reasoning should apply. The second half does not satisfy the last condition of the hourglass pattern, so the classical derivation is used there. Notice that its bound is asymptotically worse than the bound on the first part, so the bound of the first part will dominate. We consider $M = \frac{N}{2} - 1$ to obtain the first bound, and $M = N - S - 2$ to obtain the second bound when $N \gg S$.

6 RELATED WORK

I/O complexity and automation of the proof. The seminal work of Hong and Kung [12] introduced the red-blue pebble game and used it as a formalism to manually prove the I/O lower bound of programs. Following this contribution, many papers [1–3, 5, 8, 14, 15, 20–23] focused on the manual proof of lower bounds for various (classes of) programs. A large portion of these papers consider a formalism that forbid recomputation. This assumption is necessary to decompose complex CDAGs into simpler subregions, and to be able to recombine the bounds of each region.

Irony et al. [14] used the Loomis-Whitney bound, a specialization of the Brascamp-Lieb theorem with canonical projections, on the gemm algorithm. This was later extended by the work of Christ et al. [6] for the class of affine programs, by using the Brascamp-Lieb theorem. Then, Elango et al. [9] added the idea of considering paths of dependencies, completing the K -partitioning method as introduced in Section 2 of our paper.

This leads to the work of Olivry et al. [18] that introduced the first automatic lower bound derivation tool, based on combining bounds obtained from the K -partitioning and the wavefront methods of proof. This tool includes several refinements to the bound derivation, by taking advantage of some specific properties. For example, if the projections use disjoint parts of the inset, the derived bound could be improved by a constant factor. A later extension [17] exploits the fact that some dimensions have very small sizes and applies this reasoning to convolutions.

7 CONCLUSION

In this paper, we introduced a new proof reasoning to derive data movement lower bounds, based on a pattern of dependencies, called the hourglass pattern. This pattern appears on several linear algebra kernels, such as the Modified Gram-Schmidt, the QR Householder decomposition (A2V and V2Q parts), the bidiagonal matrix reduction, and the Hessenberg matrix reduction. We managed to derive asymptotically tighter lower bounds for these kernels, compared to the previous classical methods. We also provided tiled algorithms for MGS and QR A2V, whose amount of data movement matches asymptotically the new lower bounds, up to a constant factor.

We integrated the hourglass detection and its associated proof in the automatic data movement lower bound derivation tool of Olivry et al., IOLB [18], whose implementation is freely available.

Individual contributions. Fabrice identified the original problem on MGS and the intuition on how to attack this problem. The implementation in IOLB and the proof were written by Guillaume. The proof was refined by Lionel, who also provided a detailed study of the bound, and an experimental performance analysis. Julien identified the interesting kernels exhibiting an hourglass pattern and provided valuable insights on the linear algebra side. Lionel and Julien are also at the source of the upper bound proof in the appendix of the full version of this paper [11].

Acknowledgments. Julien was partially supported for this work by NSF award #2004850.

REFERENCES

[1] Alok Aggarwal and Jeffrey S. Vitter. 1988. The Input/Output Complexity of Sorting and Related Problems. *Commun. ACM* 31 (1988), 1116–1127. Issue 9.

<https://doi.org/10.1145/48529.48535>

[2] Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. 2011. Minimizing Communication in Numerical Linear Algebra. *SIAM J. Matrix Analysis Applications* 32, 3 (2011), 866–901. <https://doi.org/10.1137/090769156>

[3] Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. 2013. Graph Expansion and Communication Costs of Fast Matrix Multiplication. *Journal of the ACM* 59, 6, Article 32 (Jan. 2013), 23 pages. <https://doi.org/10.1145/2395116.2395121>

[4] Olivier Beaumont, Lionel Eyraud-Dubois, Julien Langou, and Mathieu Vérité. 2022. I/O-Optimal Algorithms for Symmetric Linear Algebra Kernels. In *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures* (Philadelphia, PA, USA) (SPAA '22). Association for Computing Machinery, New York, NY, USA, 423–433. <https://doi.org/10.1145/3490148.3538587>

[5] Gianfranco Bilardi, Michele Scquizzato, and Francesco Silvestri. 2018. A Lower Bound Technique for Communication in BSP. *ACM Transactions on Parallel Computing* 4, 3, Article 14 (feb 2018), 27 pages. <https://doi.org/10.1145/3181776>

[6] 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:1308.0068v1* [math.CA]

[7] James Demmel, Laura Grigori, Mark Hoemmen, and Julien Langou. 2008. *Communication-optimal Parallel and Sequential QR and LU Factorizations*. Technical Report UCB/EECS-2008-89. Electrical Engineering and Computer Sciences, University of California at Berkeley.

[8] James Demmel, Laura Grigori, Mark Hoemmen, and Julien Langou. 2012. Communication-optimal Parallel and Sequential QR and LU Factorizations. *SIAM Journal on Scientific Computing* 34, 1 (2012), A206–A239. <https://doi.org/10.1137/080731992>

[9] Venmugil Elango, Fabrice Rastello, Louis-Noël Pouchet, J. Ramanujam, and P. Sadayappan. 2015. On Characterizing the Data Access Complexity of Programs. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. Association for Computing Machinery, New York, NY, USA, 567–580. <https://doi.org/10.1145/2676726.2677010>

[10] Venmugil Elango, Fabrice Rastello, Louis-Noël Pouchet, J. Ramanujam, and P. Sadayappan. 2014. On characterizing the data movement complexity of computational DAGs for parallel execution. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures* (Prague, Czech Republic) (SPAA '14). Association for Computing Machinery, New York, NY, USA, 296–306. <https://doi.org/10.1145/2612669.2612694>

[11] Lionel Eyraud-Dubois, Guillaume Iooss, Julien Langou, and Fabrice Rastello. 2024. Tightening I/O Lower Bounds through the Hourglass Dependency Pattern (Full version). <https://inria.hal.science/hal-04555744>

[12] Jia-Wei Hong and H. T. Kung. 1981. I/O complexity: The Red-Blue Pebble Game. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing (STOC)* (Milwaukee, Wisconsin, United States). ACM, 326–333. <https://doi.org/10.1145/800076.802486>

[13] F. Irigoin and R. Triolet. 1988. Supernode partitioning. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (POPL '88). Association for Computing Machinery, New York, NY, USA, 319–329. <https://doi.org/10.1145/73560.73588>

[14] Dror Irony, Sivan Toledo, and Alexandre Tiskin. 2004. Communication Lower Bounds for Distributed-Memory Matrix Multiplication. *J. Parallel and Distrib. Comput.* 64, 9 (2004), 1017–1026. <https://doi.org/10.1016/j.jpdc.2004.03.021>

[15] Grzegorz Kwasniewski, Marko Kabic, Maciej Besta, Joost VandeVondele, Raffaele Solcà, and Torsten Hoefer. 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 (SC)*. 24 pages. <https://doi.org/10.1145/3295500.3356181>

[16] Reference LAPACK. 1992. <https://github.com/Reference-LAPACK/lapack>.

[17] Auguste Olivry, Guillaume Iooss, Nicolas Tollenaere, Atanas Rountev, P. Sadayappan, and Fabrice Rastello. 2021. IOOpt: Automatic Derivation of I/O Complexity Bounds for Affine Programs. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 1187–1202. <https://doi.org/10.1145/3453483.3454103>

[18] 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 (PLDI)*. ACM, 808–822. <https://doi.org/10.1145/3385412.3385989>

[19] Louis-Noël Pouchet and Tomofumi Yuki. 2016. PolyBench/C: The polyhedral benchmark suite, version 4.2. <http://polybench.sf.net>.

[20] Desh Ranjan, John E. Savage, and Mohammad Zubair. 2011. Strong I/O Lower Bounds for Binomial and FFT Computation Graphs. In *Computing and Combinatorics - 17th Annual International Conference, COCOON 2011, Dallas, TX, USA, August 14-16, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6842)*, Bin Fu and Ding-Zhu Du (Eds.). Springer, 134–145. https://doi.org/10.1007/978-3-642-22685-4_12

[21] Desh Ranjan, John E. Savage, and Mohammad Zubair. 2012. Upper and Lower I/O Bounds for Pebbling r-Pyramids. *J. Discrete Algorithms* 14 (2012), 2–12. <https://doi.org/10.1016/j.jda.2011.12.005>

[22] Jacob Scott, Olga Holtz, and Oded Schwartz. 2015. Matrix Multiplication I/O-Complexity by Path Routing. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures* (Portland, Oregon, USA) (SPAA '15). Association for Computing Machinery, New York, NY, USA, 35–45. <https://doi.org/10.1145/2755573.2755594>

[23] Tyler Michael Smith, Bradley Lowery, Julien Langou, and Robert A. van de Geijn. 2019. A Tight I/O Lower Bound for Matrix Multiplication. *arXiv:1702.02017 [cs.CC]*

[24] Sven Verdoolaege. 2010. ISL: An Integer Set Library for the Polyhedral Model. In *Mathematical Software—ICMS 2010*. 299–302.

[25] Sven Verdoolaege, Rachid Seghir, Kristof Beyls, Vincent Loechner, and Maurice Bruynooghe. 2007. Counting Integer Points in Parametric Polytopes Using Barvinok's Rational Functions. *Algorithmica* 48, 1 (2007), 37–66. <https://doi.org/10.1007/s00453-006-1231-0>