## Title

Communication-Optimal Loop Nests

## Permalink

https://escholarship.org/uc/item/5c45s0b6

## Author

Knight, Nicholas Sullender
Publication Date
2015
Peer reviewed|Thesis/dissertation

# Communication-Optimal Loop Nests 

by

Nicholas Sullender Knight

# A dissertation submitted in partial satisfaction of the requirements for the degree of Doctor of Philosophy <br> in <br> Computer Science and the Designated Emphasis <br> in <br> Computational and Data Science and Engineering <br> in the <br> Graduate Division <br> of the <br> University of California, Berkeley 

Committee in charge:
Professor James Demmel, Chair
Professor Michael Christ
Professor Katherine Yelick

Abstract<br>Communication-Optimal Loop Nests<br>by<br>Nicholas Sullender Knight<br>Doctor of Philosophy in Computer Science<br>Designated Emphasis in Computational and Data Science and Engineering<br>University of California, Berkeley<br>Professor James Demmel, Chair

Communication (data movement) often dominates a computation's runtime and energy costs, motivating organizing an algorithm's operations to minimize communication. We study communication costs of a class of algorithms including many-body and matrix/tensor computations and, more generally, loop nests operating on array variables subscripted by linear functions of the loop iteration vector. We use this algebraic relationship between variables and operations to derive communication lower bounds for these algorithms. We also discuss communication-optimal implementations that attain these bounds.

## Contents

Acknowledgements ..... i
1 Introduction ..... 1
2 Modeling Computation and Communication ..... 5
3 Modeling Loop Nests ..... 19
4 Hölder-Brascamp-Lieb-type Inequalities ..... 27
5 Communication Bounds for Loop Nests ..... 67
6 Conclusion ..... 83
References ..... 85

## Acknowledgements

My doctoral studies (August 15, 2009 through August 14, 2015) were supported financially by the Par Lab, funded by Microsoft (024263), Intel (024894), UC Discovery (DIG0710227), National Instruments, Nokia, NVIDIA, Oracle, and Samsung; by ASPIRE, funded by DARPA (HR0011-12-2-0016), C-FAR (a STARnet member, funded by the Semiconductor Research Corporation and DARPA), Intel, Google, Huawei, LG, Nokia, NVIDIA, Oracle, and Samsung; by the MathWorks; by the DOE (DE-SC0003959, DE-SC0004938, DE-SC0005136, DE-SC0008700, DE-AC02-05CH11231, DE-SC0010200); and by the NSF (ACI-1339676).

This dissertation focuses on an ongoing research project with collaborators Michael Christ, James Demmel, Thomas Scanlon, and Katherine Yelick, who I joined in January 2012. This project has benefitted from helpful technical discussions with Grey Ballard, Erin Carson, Benjamin Lipshitz, Oded Schwartz, Harsha Vardhan Simhadri, and Bernd Sturmfels.

## Chapter 1

## Introduction

Computers take time and energy to execute programs. The costs of moving data (communication) frequently dominate the costs of performing operations on that data (computation). This work studies organizing computations to minimize communication.

### 1.1 Avoiding Communication

Computers are ubiquitous in our lives, and have had a profound impact on the scientific method. Computer-based simulation enables scientists to study phenomena that are infeasible for lab experimentation, as well as collect and process data beyond human capabilities. However, computers have physical limitations too, and their users will only wait so long (or pay for so much electricity) before they expect a result. Thus it is important to ask how much computational problems cost and to develop algorithms and implementations to minimize costs.

In practice, communication costs more than computation (in terms of both time and energy) for technological reasons: improvements in processing speed and efficiency consistently exceed improvements in accessing memory, leading to an exponentially growing performance gap. So, even if the costs of a computer program are not dominated by communication on current hardware, they will likely be in the future. Additionally, as a program's working set grows, it must move data further, incurring a greater communication cost. Reducing communication costs will continue to be an important part of improving algorithmic performance.

### 1.2 Communication Lower Bounds

In this work, we derive communication lower bounds for algorithms that apply to all implementations, i.e., "any implementation of the given algorithm must move at least this much data". This means that if an implementation attains a communication lower bound, then to further reduce communication, we must seek a new algorithm. On the other hand, if there is a gap between the best known implementation and the lower bound, then it may be worth trying to find a better implementation. Communication lower bounds have practical applications in program optimization and can give useful insights into algorithm design.

The seminal work of Hong-Kung [21] provided a framework for reasoning about the communication within a sequential machine, in particular, the data movement between a fast memory with bounded capacity and a slow memory with unbounded capacity. HongKung modeled an algorithm by the directed acyclic graph ( $D A G$ ) of dependences between the operations performed in any execution of that algorithm, and modeled an implementation by playing a pebble game on the DAG's vertices, wherein placing/removing a pebble represented writing/erasing a memory cell. Hong-Kung's key result [21, Theorem 3.1] takes an algorithm and returns a lower bound on the communication in any implementation of that algorithm. Of practical importance, this lower bound is computable by solving a graph-partitioning problem, despite the fact that there are infinitely many possible implementations.

Hong-Kung's work inspired a large literature, including the present work. Our main connection with Hong-Kung's work is through the papers of Irony-Toledo-Tiskin [22] and Ballard-Demmel-Holtz-Schwartz [5], who specialized Hong-Kung's algorithm model to target a class of matrix computations whose associated DAGs have a three-dimensional geometric structure.

For example, consider the following program to multiply $N$-by- $N$ matrices, $C=A \cdot B$, assuming $C$ is zero-initialized:

$$
\begin{gathered}
\text { for } i=1: N, \quad \text { for } j=1: N, \quad \text { for } k=1: N, \\
C(i, j)=C(i, j)+A(i, k) * B(k, j) .
\end{gathered}
$$

The inner-loop statement is performed $N^{3}$ times with different $i, j, k$, each time accessing array variables $A(i, k), B(k, j)$, and $C(i, j)$. The associated DAG can be embedded in threedimensional Euclidean space, identifying each inner-loop statement as a point $(i, j, k)$ and identifying its operands (entries of $A, B, C)$ by its projections $(i, k),(k, j),(i, j)$ onto the three coordinate planes.

Irony-Toledo-Tiskin used this geometric interpretation of matrix multiplication to apply an inequality of Loomis-Whitney [24] to bound above the number of inner-loop statements doable with a fixed number of operands, enabling a simplification of Hong-Kung's graphpartitioning problem. Ballard-Demmel-Holtz-Schwartz extended these lower bounds to a more general class of matrix computations like LU-factorization that can be reduced to matrix multiplication. Of practical importance, these communication lower bounds indicated that a number of algorithms in numerical linear algebra had suboptimal communication costs, motivating a number of efforts to develop new algorithms and implementations (see [3] for a survey).

Our main contribution is to model a more general class of nested-loop programs than the ones considered by Ballard-Demmel-Holtz-Schwartz. In particular, we will allow any number of nested loops (instead of three) and any number of arrays (instead of three) each of any dimension (instead of two), and we will allow the arrays to be indexed by arbitrary linear combinations of the loop indices (instead of projections onto coordinate hyperplanes). The key to this extension is replacing Loomis-Whitney's inequality with a generalization proposed by Brascamp-Lieb [12] and analyzed by Bennett-Carbery-Christ-Tao [9].

### 1.3 Contributions

The main contributions of this work are as follows:

- New shared-memory parallel variant of Hong-Kung's sequential communication lower bound approach [21] - see Section 2.3 .
- New communication lower bounds for a class of algorithms whose dependence structures model nested-loop programs with linear array index expressions, generalizing the 'matrix-multiplication-like' class studied by Irony-Toledo-Tiskin 22] and Ballard-Demmel-Holtz-Schwartz [5] - see Section 3.2;
- Sharper constants in Hölder-Brascamp-Lieb-type inequalities, extending work by Bennett-Carbery-Christ-Tao [9] and enabling our new communication lower bounds - see Section 4.2;
- New algorithm to compute constraints for Hölder-Brascamp-Lieb-type inequalities, extending work by Valdimarsson [34] and showing that communication lower bounds are computable - see Section 4.2, and
- New communication-optimal implementations, which attain communication lower bounds by tiling, in many cases of practical interest - see Section 5.2 .


### 1.4 Outline

A chapter-by-chapter outline of this work is as follows.
In the remainder of Chapter 1 we review the mathematical notation used subsequently.
In Chapter 2 we model the execution of an idealized shared-memory parallel computer, where each processor has a private local memory and all processors share a global memory, and we define communication as data movement between local and global memory. We use input- and output-paths, structural properties of algorithms, to derive lower bounds on the amount of communication in any implementation of a given algorithm.

In Chapter 3 we define Hölder-Brascamp-Lieb (HBL) interpretations of algorithms which provide an intuitive way to model a class of nested-loop programs with certain input-/outputpath properties. Recall the earlier example of multiplying $N$-by- $N$ matrices, $C=A \cdot B$,

$$
\begin{gathered}
\text { for } i=1: N, \quad \text { for } j=1: N, \quad \text { for } k=1: N, \\
C(i, j)=C(i, j)+A(i, k) * B(k, j)
\end{gathered}
$$

modeling this program as a DAG, we will see that any inner-loop statement $(i, j, k)$ has incoming paths from input array variables $A(i, j)$ and $B(k, j)$ and an outgoing path to output array variable $C(i, j)$. These input-/output-paths model data dependences, which we use to derive lower bounds on data movement (communication).

In Chapter 4, we study HBL-type inequalities, for later use analyzing algorithms with HBL interpretations. We extend previous work [9] on HBL-type inequalities by demonstrating sharper constants, as well as extend previous work 34 on computability issues related to these inequalities.

Finally, in Chapter 5, we apply the HBL theory to complete our communication lower bound analysis of algorithms with HBL interpretations. We also show, in some important special cases, that these communication lower bounds are asymptotically attainable via tiling.

Conclusions are drawn and future work is reviewed in Chapter 6.

### 1.5 Notation

We employ standard mathematical notations with a few non-standard extensions.
The set of natural numbers is denoted by $\mathbb{N}=\{0,1,2, \ldots\}$; for any $n \in \mathbb{N}$, we let $\mathbb{N}_{\geq n}=\{n, n+1, \ldots$,$\} , so \mathbb{N}_{\geq 0}=\mathbb{N}$, and we let $[n]=\{1, \ldots, n\}$, so $[0]=\emptyset$.

For any set $X$, its cardinality $|X| \in \mathbb{N} \cup\{\infty\}$. For any $n \in \mathbb{N}, X^{n}$ denotes the $n$-th Cartesian power of $X$; more generally, if $Y$ is another set, then $X^{Y}$ denotes the set of all functions from $Y$ to $X$.

An Abelian group is a set $G$ that is closed under addition + , an associative and commutative binary operation, with identity element 0 and an inverse element $-x$ for each $x \in G$. A group homomorphism is a function between Abelian groups that is compatible with the group structure. The rank of an Abelian group is the maximum cardinality of any linearly independent subset of $G$.

## Chapter 2

## Modeling Computation and Communication

In this chapter we formalize our computation model in order to reason mathematically about the communication costs of algorithms. A section-by-section outline is as follows.

- In Section 2.1, we model a computation as an execution of an idealized parallel machine with a two-level memory hierarchy, where each processor has a private local memory and all processors share a global memory. Communication is data movement between the local memories and the global memory.
- In Section 2.2, we introduce our notion of algorithms, which correspond to sets of machine executions called implementations.
- Then in Section 2.3, we use input- and output-paths, structural properties of algorithms, to derive lower bounds on the amount of communication in any implementation of a given algorithm.
- We discuss related work in Section [2.4 our model and communication lower bounds approach was inspired by Hong-Kung's seminal work 21.


### 2.1 Model Preliminaries

Machine Model A machine has a nonempty finite set of processors, each connected to a local memory and all connected to a global memory (see Figure 2.1). A memory contains storage elements called cells, each of which stores a single value; local memories have finitely many cells, while global memory can have infinitely many. Local memory is private, i.e., each processor can access its own, and only its own. On the other hand, any processor can access global memory, i.e., global memory is shared. Processors can only perform (arithmetic/logical) operations on local data, and it is often necessary to communicate (move data) between local and global memory. For example, communication is necessary when a processor's working set becomes sufficiently large relative to its (bounded) local memory size - our main results concern this particular computation/communication relationship.


Figure 2.1: Model machine.

While values and operations will remain tangential to our analysis, it helps to have concrete definitions of these terms: a value is an element of some given nonempty set $D$ and an operation is a function $f: D^{m} \rightarrow D^{n}$ for some $m, n \in \mathbb{N}_{\geq 1}$, an element of some given nonempty set $B$ of such functions. Our machine model is thus implicitly parameterized by $D$ and $B$.

Execution Model An execution comprises a finite ordered set of instructions and an assignment of (a function from) instructions to processors. Each processor can be instructed to load a value from global to local memory (Load), to store a value from local to global memory (Store), and to perform an operation on values in local memory (Operation). (An execution gives an a posteriori perspective, so there is no need to model branches, only operations and data movement.)

The execution order, a partial order, defines a temporal precedence relation: if two instructions are comparable then they were not issued concurrently. In particular, the execution order enforces two model assumptions regarding parallel execution:

1. "Sequential processors": each processor performs one instruction at a time in a given total order. That is, each processor's instructions compose a chain (totally ordered subset) in the execution order - see Figure 2.2 (left).
2. "Concurrent reads, exclusive writes, and no data races": each memory cell is associated with a total order that contains all pairs of reads and writes in which at least one element in the pair is a write. That is, each global cell's associated Stores compose a chain in the execution order which remains a chain when including any one of that cell's associated Loads - see Figure 2.2 (right).

Note that assumption 2 neglects local memory accesses because assumption 1 implies the stronger constraint that local reads (in addition to local writes) are exclusive.

A schedule of an execution defines start and end times (real numbers) for each instruction such that each instruction's start time follows its predecessors' end times. Thus, an execution


Figure 2.2: Execution assumptions. Left: some processor's instructions ("sequential processors"); right: some cell's reads/writes ("concurrent reads, exclusive writes, and no data races").
order enforces the two preceding model assumptions by constraining the set of permissible schedules (notably, an execution itself has no notion of time).

Cost Model Suppose each instruction in an execution is assigned a nonnegative cost. We then define the cost of a chain in the execution as the sum of its instructions' costs; as a special case, we define the cost of a processor as the cost of that processor's chain. We also define the critical-path cost as the maximum cost over all chains (not necessarily a single processor's) in the execution order. Lastly, we define the total cost as the sum of all instructions' costs.

In this work we consider simple 0/1-cost assignments to count different types of instructions: for each chain in the execution order, we define its computation cost to be its number of Operations and its communication cost to be its number of Loads and Stores. The computation/communication tradeoffs mentioned above will be quantified in Section 2.3 in terms of these two costs.

We detour to show how three physical costs - time, energy, and power - can also modeled in this framework.

First, suppose that the instructions' costs are their time-costs (durations): we can model runtime with the critical-path time-cost. To see this connection, consider the greedy schedule, constructed by recursively assigning each instruction a start time equal to the maximum of its predecessors' end times (zero if none exist) and an end time equal to its start time plus its time-cost. This approach may introduce idle time - e.g., if processor $p$ stores a value and processor $q$ subsequently loads it, then $q$ may be kept idle while the Store completes but there always exists a critical path with no idle time. For any time-cost assignment of an execution, according to its greedy schedule, the length of the interval between zero and the maximum end time equals the critical-path time-cost.

Second, suppose that the instructions' costs are their energy-costs (amounts of energy consumed): we can model total energy consumption (or some processor $p$ 's energy consump-
tion) by the total energy-cost (or $p$ 's energy-cost).
Third, if we are given both time- and energy-cost assignments, then we can model average power of the machine (or processor $p$ ) by dividing the total energy-cost (or $p$ 's energy-cost) by the critical-path time-cost.

The point of this digression was to demonstrate that our simple cost model, counting instructions by type, can in some cases yield quantitative insights into physical costs: e.g., if Operations have a fixed time-cost and Loads/Stores have a fixed time-cost, then (modeled) runtime is bounded above/below by the sum/maximum of the critical-path computation and communication costs (this lower bound assumes complete overlap of computation and communication).

### 2.2 Analyzing Executions

We now introduce some notation and bookkeeping tools that will aid our subsequent analyses. Our analytical approach is independent of the particular values arising during execution, as well as the particular operations performed on them. Rather, our analysis depends on the particular cells each instruction accesses (reads from and writes to), and the order in which they do so.

Instruction Encoding As mentioned, for each instruction $i$ in a given execution, we need only observe the cells it accesses. Each instruction $i$ reads values from its input cells $a_{1}, \ldots, a_{m}$ (not necessarily distinct) and writes values to its output cells $b_{1}, \ldots, b_{n}$ (necessarily distinct). We define $\operatorname{In}(i)=\left\{a_{1}, \ldots, a_{m}\right\}$ and $\operatorname{Out}(i)=\left\{b_{1}, \ldots, b_{n}\right\}$. If $i$ is a Load (resp., Store), then $m=n=1$, and $a_{1}$ is a global (resp., local) cell and $b_{1}$ is a local (resp., global) cell. If $i$ is an Operation, then $m, n \in \mathbb{N}_{\geq 1}$ and all cells are local.

Data Dependence An execution implies the manner in which individual values propagate: instructions depend on values created/copied by their predecessors in the execution order. Since new values can only be introduced by Operations - Loads and Stores only make copies of existing values - a value's lifetime either begins with the Operation that created it, or otherwise the value resided in memory since the start of execution.

We now formalize the notion of a lifetime as a chain in the execution order; keep in mind that this definition will depend only on the cells associated with the instructions and not the particular values they store: for example, the same lifetimes are constructed even if only one distinct value arises in the execution. In this sense, data dependence is a data-independent notion.

Consider any instruction $i$ : we first define the lifetimes of $i$ 's output values and second the lifetimes of $i$ 's input values.

- The lifetime of the value written to some (local or global) cell $c \in$ Out $(i)$ by some instruction $i$ is a chain in the execution order that ends with $i$. If $i$ is an Operation, then we return the singleton chain (i). Otherwise, $i$ was a Load or Store, thus has a unique $b \in \operatorname{In}(i)$; we recursively compute the lifetime $\left(i_{1}, \ldots, i_{n}\right)$ of the value read from
$b \in \operatorname{In}(i)$ by $i$ according to the procedure in the following bullet, and then return the chain $\left(i_{1}, \ldots, i_{n}, i\right)$.
- The lifetime of the value read from some (local or global) cell $c \in \operatorname{In}(i)$ by some instruction $i$ is a chain in the execution order, possibly empty. There are two possibilities: either some ancestor $j$ of $i$ wrote this value to $c$ (i.e., $c \in \operatorname{Out}(j)$ ), or no such ancestor $j$ of $i$ exists. In the latter case, the value initially stored in $c$ was never altered prior to being read by $i$, which we indicate by returning the empty chain (). In the former case, due to the "concurrent reads, exclusive writes, and no data races" assumption defined in Section 2.1, we can identify the unique maximal (most recent) ancestor $j$ of $i$ such that $c \in \operatorname{Out}(j)$. If $j$ is an Operation, then we return the singleton chain $(j)$. Otherwise, $j$ was a Load or Store, thus has a unique $b \in \operatorname{In}(j)$; we recursively compute the lifetime $\left(j_{1}, \ldots, j_{n}\right)$ of the value read from $b \in \operatorname{In}(j)$ by $j$ according to the procedure in the preceding bullet, and then return the chain $\left(j_{1}, \ldots, j_{n}, j\right)$.

Since there are finitely many instructions in the execution, this (downwards) recursion terminates with a finite chain, possibly empty, consisting entirely of Loads and Stores except possibly for its initial element. In this manner, we can trace the lifetime of every value that is read or written by an instruction.

When additionally given a schedule of the execution, we can extend this procedure to define the lifetime of the value stored in any cell $c$ at any particular point $t$ in time, provided that $c$ is not being (over-) written at time $t$. That is, since our model implies that $c$ stores a consistent value between the (totally ordered) instructions that write to $c$, we can apply the procedure to the most recent (before $t$ ) instruction $i$ with $c \in \operatorname{Out}(i)$, or returning () if none exists. This observation also emphasizes that lifetimes are independent of the schedule.

Algorithms Each execution on a given machine implements an algorithm, an abstraction that ignores mechanical details to focus on the dependences between values. In Section 2.3. we will derive lower bounds on communication cost in terms of algorithms: these lower bounds apply to all executions which implement that algorithm. We will first define algorithms without reference to executions and then demonstrate how executions implement algorithms. Note that the following (non-standard) definition of 'algorithm' is more abstract than the standard definition (a computable function, Turing machine, etc.). We use the standard definition in Chapter 4 in the context of decision problems (see Theorems 4.3 and 4.4), but use the non-standard definition exclusively in all other chapters.

An algorithm is a finite set of (function) evaluations, symbolic expressions of the form

$$
\left(y_{1}, \ldots, y_{n}\right)=f\left(x_{1}, \ldots, x_{m}\right)
$$

for various $m, n \in \mathbb{N}_{\geq 1}$. The symbol $f$ represents an operation with input arity $m$ and output arity $n$; we suppose that $f$ is distinct in each function evaluation and let $F$ be the set of these symbols. The tuples $\left(x_{1}, \ldots, x_{m}\right)$ and $\left(y_{1}, \ldots, y_{n}\right)$ contain symbols called variables; we let $V$ be the set of these symbols and suppose that $V \cap F=\emptyset$. Additionally, a variable can only appear once as some output $y_{j}$ ("single assignment form"), but can appear more than once as an input $x_{i_{1}}, x_{i_{2}}, \ldots$. To formalize this constraint, we introduce $f$-subscripts to
disambiguate between evaluations, $x_{f, 1}, \ldots, x_{f, m_{f}}$ and $y_{f, 1}, \ldots, y_{f, n_{f}}$, i.e., we define the two functions

$$
\begin{aligned}
& x:\left\{(f, i) \mid f \in F \wedge i \in\left[m_{f}\right]\right\} \rightarrow V \\
& y:\left\{(f, j) \mid f \in F \wedge j \in\left[n_{f}\right]\right\} \rightarrow V
\end{aligned}
$$

With this notation, the 'single assignment form' assumption means that $y$ is an injection while $x$ need not be. The elements of $\operatorname{dom}(x)$ and $\operatorname{dom}(y)$ are called input and output arguments, resp., and the elements of $\operatorname{im}(x)$ and $\operatorname{im}(y)$ are called input and output argument variables, resp. (by definition, $V=\operatorname{im}(x) \cup \operatorname{im}(y)$ ). By injectivity of $y$ (and not of $x$ ), each $v \in V$ is an output argument variable $y_{f, j}$ for at most one output argument $(f, j)$, but can be an input argument variable $x_{f_{1}, i_{1}}, x_{f_{2}, i_{2}}, \ldots$ for multiple input arguments $\left(f_{1}, i_{1}\right),\left(f_{2}, i_{2}\right), \ldots$.

Besides $V, F, x, y$, an algorithm has designated input and output variables $I, O \subseteq V$, where $I=\operatorname{im}(x) \backslash \operatorname{im}(y)$ and $\operatorname{im}(y) \backslash \operatorname{im}(x) \subseteq O \subseteq \operatorname{im}(y)$. Here we observe an asymmetry in our model concerning inputs and outputs, i.e., that $I$ is completely specified by $x, y$ but $O$ is not necessarily; by convention, we assume $O=\operatorname{im}(y) \backslash \operatorname{im}(x)$ when $O$ is unspecified.

The functions $x$ and $y$ imply a partial order on the set of evaluations. The algorithm $D A G(V, E)$ has the variables as vertices and has edge set $E=\bigcup_{f \in F}\left\{x_{f, 1}, \ldots, x_{f, m_{f}}\right\} \times$ $\left\{y_{f, 1}, \ldots, y_{f, n_{f}}\right\}$; the input variables $I$ are the indegree- 0 vertices and the output variables $O$ include the outdegree-0 vertices.

Implementations Next we show how to construct an algorithm from an execution that implements it. Before doing so, we explain how this approach deviates from Hong-Kung's 21.

Hong-Kung start with algorithms modeled by DAGs (Hong-Kung's DAGs are equivalent to our algorithm DAGs in the special case $n_{f}=1$ for all $f \in F$, except for the definition of input variables $I$ ) and then define an algorithm's implementations as the different gameplays of the 'red/blue pebble game' on that DAG. (In the red/blue pebble game, defined in 21, Section 1], red and blue pebbles, which model local and global memory cells, resp., are placed on the vertices of the DAG, modeling storing the associated values in the associated cells.) Our approach starts with executions and then defines the mapping from executions to the algorithms they implement: in Hong-Kung's framework, this is like starting with a gameplay and reconstructing the DAG by studying the gameplay's moves. Our approach was motivated by complications we encountered while extending Hong-Kung's approach to the parallel setting (Hong-Kung modeled sequential machines). We note that our approach is more general than Hong-Kung's by allowing certain 'wasteful' implementations, e.g., those including Operations whose outputs are subsequently overwritten - however, we have not seen any theoretical benefit from this generality.

Roughly speaking, we construct an algorithm from an execution by identifying Operations and values in the execution with function evaluations and variables in the algorithm.

We consider each cell $c$ (local and global) in turn. Compute the lifetime of the value stored in $c$ at the end of execution: that is, identify the last instruction $i$ in the execution order with $c \in \operatorname{Out}(i)$ and compute the lifetime of the value $i$ wrote to $c$, or otherwise if no instructions wrote to $c$, return the empty chain () as the lifetime. Now examine the lifetime $\left(i_{1}, \ldots, i_{t}\right)$ just computed and consider two cases; we encapsulate our processing of
these cases in a subprocedure which will be reentered recursively with additional parameters denoted $g$ and $h$.

- If the lifetime is the empty chain or if $i_{1}$ is a Load or Store, then we do nothing unless this is a recursive call. If this is a recursive call, we augment $V$ with the variable $c$ (identified with the cell $c$ ) and define $x_{g, h}=c$.
- Otherwise, $i_{1}$ is an Operation. If $i_{1}$ has been encountered already, let $f \in F$ denote its index; otherwise, augment $F$ with a new symbol $f \in F$ that indexes $i_{1}$. Let $m_{f}$ and $n_{f}$ denote the number of $i_{1}$ 's input and output cells. We can identify exactly one $k \in\left[n_{f}\right]$ such that the lifetime began when $i_{1}$ wrote to its $k$-th output cell $b_{k} \in \operatorname{Out}\left(i_{1}\right)$. If the variable $(f, k)$ does not yet exist in $V$, augment $V$ with the variable $(f, k)$ and define $y_{f, k}=(f, k)$. Additionally, if this is a recursive call, define $x_{g, h}=(f, k)$; otherwise, if it is not, designate $(f, k)$ as an output variable, $(f, k) \in O$.
Now, for each $i \in\left[m_{f}\right]$, trace the lifetime $\left(j_{1}, \ldots, j_{q}\right)$ of the value that $i_{1}$ read from its $i$-th input cell $a_{i} \in \operatorname{In}\left(i_{1}\right)$ and recursively evaluate this subprocedure substituting $a_{i}$ for $c$ and $\left(j_{1}, \ldots, j_{q}\right)$ for $\left(i_{1}, \ldots, i_{t}\right)$, and setting the parameters $g=f$ and $h=k$. Each of these recursive calls must terminate because the execution is finite.

We then proceed with the next cell $c$.
After all (potentially infinitely many) cells have been processed in this manner, we have constructed the objects $V, F, x, y, O$, where $V, F$ are finite; recall that the input variables are completely defined by $x, y: I=\operatorname{im}(x) \backslash \operatorname{im}(y)$. The constructed objects $V, F, x, y, I, O$ may fail to define an algorithm for a subtle reason: some function evaluations may not have had variables defined for each of their output arguments. Suppose, for example, that function evaluation $f \in F$ has an undefined variable $y_{f, k}$ for its $k$-th output $\operatorname{argument}(f, k)$ : this means that the associated Operation's $k$-th output cell $b_{k}$ was subsequently overwritten by another instruction without any intervening reads. Due to the top-down recursive approach taken, the fact that $f \in F$ implies that at least one of $f$ 's output argument variables must be defined - Operations whose outputs are all overwritten are never added to the algorithm so $n_{f} \geq 2$. Thus we can simply delete $f$ 's $k$-th output argument, reindex the output argument $(f, j+1) \mapsto(f, j)$ and output argument variables $y_{f, j+1} \mapsto y_{f, j}$ for each $j \in\left\{k, \ldots, n_{f}-1\right\}$, and redefine $n_{f} \mapsto n_{f}-1 \geq 1$ (this alters $F, y$ but not $V, x, I, O$ ).

After performing the above postprocessing step, we obtain an algorithm which we say is the one implemented by the given execution; every execution implements an algorithm, while an algorithm has many implementations. This construction also emphasizes an earlier remark that our model does not capture branching: if a conditional branch caused a different control path to be taken, then the corresponding executions would be mapped to different (branch-free) algorithms.

### 2.3 Communication Lower Bounds

We now explore a connection between an algorithm's DAG and its implementations' communication costs. (Communication costs were defined in Section 2.1 and algorithms/implementations were defined in Section 2.2.) A similar connection was observed by Hong-Kung [21] - we
review related work in Section 2.4. In this section, we assume a fixed algorithm, specified by the parameters $V, F, x, y, I, O$ and inducing the algorithm DAG $G=(V, E)$, as defined in Section 2.2.

We first define two graph-theoretic properties called input- and output-path cutsize. For any $K \subseteq F$, let $\Pi_{\text {in }}(K)$ denote the set of all paths in $G$ from $I$ to $\left\{x_{f, 1}, \ldots, x_{f, m_{f}} \mid f \in K\right\}$, called input-paths, and let $\Pi_{\text {out }}(K)$ denote the set of all paths in $G$ from $\left\{y_{f, 1}, \ldots, y_{f, n_{f}} \mid\right.$ $f \in K\}$ to $O$, called output-paths. (An input or output path can be a single element of $I$ or $O$.) Let $Q_{\text {in }}(K)$ and $Q_{\text {out }}(K)$ denote the maximum cardinalities of subsets of $\Pi_{\text {in }}(K)$ and $\Pi_{\text {out }}(K)$ comprising only disjoint paths: we call $Q_{\text {in }}(K)$ and $Q_{\text {out }}(K)$ the input- and output-path cutsizes of $K$.

We observe two immediate properties of $\Pi_{\mathrm{in}}, \Pi_{\text {out }}, Q_{\mathrm{in}}, Q_{\text {out }}$. First, if $K=\emptyset$, then $Q_{\text {in }}(K)=Q_{\text {out }}(K)=0$; otherwise, if $\emptyset \neq K \subseteq F, 1 \leq Q_{\text {in }}(K) \leq \sum_{f \in K} m_{f}$ and $1 \leq$ $Q_{\text {out }}(K) \leq \sum_{f \in K} n_{f}$. Second, for any $E_{1}, E_{2} \subseteq F, E_{1} \subseteq E_{2}$ implies that $\Pi_{\text {in }}\left(E_{1}\right) \subseteq \Pi_{\text {in }}\left(E_{2}\right)$ and $\Pi_{\mathrm{out}}\left(E_{1}\right) \subseteq \Pi_{\mathrm{out}}\left(E_{2}\right)$, therefore $Q_{\mathrm{in}}\left(E_{1}\right) \leq Q_{\mathrm{in}}\left(E_{2}\right)$ and $Q_{\mathrm{out}}\left(E_{1}\right) \leq Q_{\mathrm{out}}\left(E_{2}\right)$.

Now consider any implementation of the algorithm fixed above. Let $P$ denote the set of processors; for each processor $p \in P$, let $M_{p} \in \mathbb{N}_{\geq 1}$ be processor $p$ 's (finite) local memory size. (These machine parameters are implied by the implementation.) For each $p \in P$, let $F_{p} \subseteq F$ denote the set of evaluations (in the algorithm) whose corresponding Operations (in the execution) are assigned to processor $p$. Continuing notation, we have the following connection between input-/output-path cutsize and communication.

Lemma 2.1. For each $p \in P$ and each nonempty $K \subseteq F_{p}$, processor $p$ performs at least $Q_{\text {in }}(K)-M_{p}$ Loads and at least $Q_{\text {out }}(K)-M_{p}$ Stores between the first and last Operations corresponding to $K$.

Proof. See Section 2.3.1.
The conclusion of Lemma 2.1 is a set of lower bounds on $W_{p}$, processor $p$ 's communication cost - the total number of Loads and Stores it performs - for each $p \in P$. Any of these perprocessor communication costs implies a lower bound on $W$, the critical-path communication cost; in particular,

$$
W \geq \max _{p \in P} W_{p}
$$

When applying Lemma 2.1 to derive lower bounds on $W_{p}$ for any $p \in P$, by monotonicity of $Q_{\text {in }}, Q_{\text {out }}$, there is no loss of generality to consider only the case $K=F_{p}$, i.e.,

$$
\begin{aligned}
W_{p} & \geq \max _{K \subseteq F_{p}}\left(\max \left(0, Q_{\mathrm{in}}(K)-M_{p}\right)+\max \left(0, Q_{\mathrm{out}}(K)-M_{p}\right)\right) \\
& =\max \left(0, Q_{\mathrm{in}}\left(F_{p}\right)-M_{p}\right)+\max \left(0, Q_{\mathrm{out}}\left(F_{p}\right)-M_{p}\right) .
\end{aligned}
$$

Additionally, we can tighten this lower bound given more information about the contents of the processor $p$ 's local memory: for example, following 21 and requiring that all input/output values reside in global memory before/after execution, we can show that $W_{p} \geq$ $Q_{\text {in }}\left(F_{p}\right)+Q_{\text {out }}\left(F_{p}\right)$ for all $p \in P$. Lower bounds of this form have been called memoryindependent (see, e.g., [4]). Our goal is to derive lower bounds on each $W_{p}$ that are memorydependent, i.e., parameterized by $M_{p}$.

The following result, the main result of Section 2.3, applies Lemma 2.1 over a sequence of subsets of $F_{p}$ to obtain a lower bound on $W_{p}$. We also adopt a stronger hypothesis regarding the input- and output-path cutsizes of these subsets.

Theorem 2.1. For any $p \in P$, if there exists $K \subseteq F_{p}$ such that $N=|K|>0$ and $a$ nondecreasing $q:[N] \rightarrow \mathbb{R}$ such that $q(|J|) \leq Q_{\text {in }}(J)+Q_{\text {out }}(J)$ for all nonempty $J \subseteq K$, then processor $p$ performs

$$
\begin{equation*}
W_{p} \geq \max _{l \in \mathbb{N} \geq 1} l \cdot\left(\left\lceil\frac{N}{\max \left\{i \in[N] \mid\lceil q(i)\rceil \leq 2 M_{p}+l\right\}}\right]-1\right) \tag{2.1}
\end{equation*}
$$

Loads and Stores between the first and last Operations corresponding to $K$.
Proof. See Section 2.3.1.
In particular, the two maxima in (2.1) are well-defined. We show after proving Theorem 2.1 how its proof can be adjusted to exploit individual lower bounds on $Q_{\text {in }}, Q_{\text {out }}$ to derive similar conclusions for Loads and Stores independently (as in Lemma 2.1).

In this work, we will use the following specialization and simplification of Theorem 2.1.
Corollary 2.1. If the hypotheses of Theorem 2.1 are satisfied by $q(i)=i^{1 / \sigma}$ for some $\sigma \in$ $[1, \infty)$, then the conclusion of Theorem 2.1 can be simplified: processor $p$ performs at least

$$
W_{p} \geq \max _{l \in \mathbb{N} \geq 1} l\left(\left\lceil\frac{N}{\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor}\right]-1\right) \geq \begin{cases}\frac{1}{2 \cdot 3^{\sigma}} \frac{N}{M_{p}^{\sigma-1}} & \text { if } N \geq\left(4 M_{p}\right)^{\sigma}  \tag{2.2}\\ 0 & \text { otherwise }\end{cases}
$$

Loads and Stores between the first and last Operations corresponding to $K$.
Proof. See Section 2.3.1.
Corollary 2.1 addresses implementations of algorithms where, if processor $p$ performs sufficiently many Operations $N$ with respect to their local memory size $M_{p}$, then they must communicate as indicated by (2.2). Moreover, in this regime, there exists a lower-bound tradeoff between communication and memory, which we express in asymptotic notation,

$$
W_{p}=\Omega\left(\frac{N}{M_{p}^{\sigma-1}}\right) \quad \Rightarrow \quad W_{p} \cdot M_{p}^{\sigma-1}=\Omega(N)
$$

In Chapter 3, we will define and give examples of a class of algorithms to which Corollary 2.1 applies (the proof of this follows later, in Section 5.1). Then, in Chapter 5, we determine the parameter $\sigma$ for the example algorithms in Chapter 3 as well as study the attainability of the communication lower bounds.

### 2.3.1 Proofs

Before proving Lemma 2.1, we give an informal alternative argument which ultimately fails for reasons related to an obstacle called 'R2/D2 arguments' by Ballard-Demmel-HoltzSchwartz [5], also called 'paired Allocate/Frees statements of array variables' in our earlier work [15]. A contribution of the present work is a graph-theoretic reinterpretation of techniques both prior works employed to overcome this obstacle.

We continue the notation introduced earlier in Section 2.3; in particular, we consider some evaluations $K$ assigned to some processor $p$. Let $X_{K}=\left\{x_{f, 1}, \ldots, x_{f, m_{f}} \mid f \in K\right\}$ and $Y_{K}=\left\{y_{f, 1}, \ldots, y_{f, n_{f}} \mid f \in K\right\}$ denote the associated input and output argument variables. Before the first and after the last element of its subchain $K$ of assigned instructions, processor $p$ can have at most $M_{p}$ values in local memory. Since an implementation of an algorithm, by definition, treats variables as if their values were distinct, $\left|X_{K}\right|-M_{p}$ and $\left|Y_{K}\right|-M_{p}$ are perhaps intuitive candidates for lower bounds on the numbers of Loads and Stores processor $p$ must perform. This intuition fails: some values associated with both $X_{K}$ and $Y_{K}$ could be temporary, meaning that processor $p$ computes and discards (overwrites) them without incurring Loads and Stores, thus beating the 'lower bounds'. In addition to being quite common, this scenario is also desirable from the standpoint of minimizing communication (the practical motivation of our work).

To derive a lower bound on Loads and Stores in terms of numbers of input and output argument variables, we need to account for temporaries. Our accounting technique, codified in the definitions of $Q_{\text {in }}, Q_{\text {out }}$, is closely related to Ballard-Demmel-Holtz-Schwartz's usage of 'accumulators' [5, Section 2] and to 'surrogates', a generalization we introduced later in 15 , Section 4.2]. These ideas, like many in the present work, have their roots in Hong-Kung's seminal paper [21], which we discuss in Section 2.4. Lastly, we point out a complementary accounting technique, a reduction approach called imposing reads and writes [5, Section 3.4] (see also [5, Corollary 5.1], [2, Section 4.2], and [15, Section 4.4]); developing a graphtheoretic analogue of this tool for use in the present model remains future work.

Proof of Lemma 2.1. Recall the notation defined in the paragraphs preceding the statement of Lemma 2.1. We consider any processor $p \in P$ with $F_{p} \neq \emptyset$ and any nonempty $K \subseteq F_{p}$. We let $s$ and $t$ denote the first and last instructions (both Operations) corresponding to $K$. Let $X_{K}=\left\{x_{f, 1}, \ldots, x_{f, m_{f}} \mid f \in K\right\}$ and $Y_{K}=\left\{y_{f, 1}, \ldots, y_{f, n_{f}} \mid f \in K\right\}$.

We first show that any path in the algorithm DAG from some $u \in I$ to some $w \in X_{K}$ contains at least one vertex $v$ whose associated value either resided in processor $p$ 's local memory immediately before $s$ or was loaded by processor $p$ between $s$ and $t$. By definition, $w=x_{f, i}$ for some some evaluation (and Operation) $f \in F_{p}$ and $i \in\left[m_{f}\right]$. There are three possibilities for the value that $f$ read from its input cell $a_{i} \in \operatorname{In}(f)$, not necessarily mutually exclusive: it resided in processor $p$ 's local memory immediately before $s$, it was Loaded by processor $p$ after $s$ and before $f$, or it was written by processor $p$ during some Operation $g \in F_{p}$ such that $s \leq g<f$. In the first two cases, we conclude with $v=w$. Otherwise, $w$ 's predecessor along the given path equals $x_{g, k}$ for some $k \in\left[m_{g}\right]$, and we repeat the same reasoning, replacing $x_{f, i}$ by $x_{g, k}$. We have made progress along the given path, and since the algorithm is finite, this recursion must terminate, and it does so having determined the desired vertex $v$.

We second show that any path in the algorithm DAG from some $u \in Y_{K}$ to some $w \in O$ contains at least one vertex $v$ whose associated value either resided in processor $p$ 's local memory immediately after $t$ or was stored by processor $p$ between $s$ and $t$. By definition, $u=y_{f, j}$ for some evaluation (and Operation) $f \in F_{p}$ and $j \in\left[1, \ldots, n_{f}\right]$. There are three possibilities for the value that $f$ wrote to its output cell $b_{j} \in \operatorname{Out}(f)$, not necessarily mutually exclusive: it resided in processor $p$ 's local memory immediately after $t$, it was Stored by processor $p$ after $f$ and before $t$, or it was read by processor $p$ during some Operation $g \in F_{p}$ such that $f<g \leq t$. (It is impossible that this value was simply overwritten without intervening reads due to the definition of an implementation.) In the first two cases, we conclude with $v=u$. Otherwise, $u$ 's successor along the given path equals $y_{g, k}$ for some $k \in\left[n_{g}\right]$, and we repeat the same reasoning, replacing $y_{f, j}$ by $y_{g, k}$. We have made progress along the given path, and since the algorithm is finite, this recursion must terminate, and it does so having determined the desired vertex $v$.

Consider any $\Pi_{\mathrm{in}}^{\prime} \subseteq \Pi_{\mathrm{in}}(K)$ : each element of $\Pi_{\mathrm{in}}^{\prime}$ contributes a variable $w$ whose value either resides in processor $p$ 's local memory immediately before $s$ or is Loaded by processor $p$ between $s$ and $t$, and at most $M_{p}$ variables can satisfy the former case. If $\Pi_{\mathrm{in}}^{\prime}$ is a set of disjoint paths, then these variables are distinct, so the number of processor $p$ 's Loads between $s$ and $t$ must be at least $\max \left(0,\left|\Pi_{\text {in }}^{\prime}\right|-M_{p}\right\}$. Similarly, consider any $\Pi_{\text {out }}^{\prime} \subseteq \Pi_{\text {out }}(K)$ : each element of $\Pi_{\text {out }}^{\prime}$ contributes a variable $w$ whose value either resides in processor $p$ 's local memory immediately after $t$ or is Stored by processor $p$ between $s$ and $t$, and at most $M_{p}$ variables can satisfy the former case. If $\Pi_{\text {out }}^{\prime}$ is a set of disjoint paths, then these variables are distinct, so the number of processor $p$ 's Stores between $s$ and $t$ must be at least $\max \left(0,\left|\Pi_{\text {out }}^{\prime}\right|-M_{p}\right\}$. The desired conclusion follows from the fact that there exists a $\Pi_{\text {in }}^{\prime} \subseteq \Pi_{\text {in }}(K)$ comprising $Q_{\text {in }}(K)$ disjoint paths and a $\Pi_{\text {out }}^{\prime} \subseteq \Pi_{\text {out }}(K)$ comprising $Q_{\text {out }}(K)$ disjoint paths.

Now it is time to prove our main lower bound result, Theorem 2.1. We use a simple counting argument which we call segmentation, in which we apply Lemma 2.1 to subchains of the chain $K$. A similar argument was introduced by Hong-Kung in their the proof of 21 , Lemma 3.1] and can be seen in many works building on Hong-Kung's. A distinction between our segmentation argument and many prior ones is that we allow the segment size to vary, rather than fixing it equal to the local memory size (however, we later make a similar choice to obtain the simpler result in Corollary 2.1). This generalized segmentation argument can improve the constant factors in the lower bound: for example, recent work of Langou-Lowery [25] exploited this flexibility to improve the constants in the communication lower bounds for matrix multiplication derived previously in [22, 5].

Proof of Theorem 2.1. Consider the hypothesized $p \in P$ and $K \subseteq F_{p}$. Let's introduce some additional notation. For any subchain $J$ of processor $p$ 's chain of instructions, let $W_{J}$ denote the number of (processor $p$ 's) Loads and Stores in $J$, and let $\bar{J}$ denote all of processor $p$ 's assigned instructions between the first and last instruction of $J$, inclusive; with this notation, let $L=\bar{K}$. We will prove a stronger result, that the right-hand side of (2.1) is a lower bound on $W_{L}$, which in turn is the desired lower bound on $W_{p}$, processor $p$ 's communication cost.

For each $l \in \mathbb{N}_{\geq 1}$, there exists a unique $k_{l} \in \mathbb{N}_{\geq 1}$ such that $L$ is the concatenation $L=L_{1} \cdots L_{k_{l}}$ of $k_{l}$ subchains $L_{j}$ such that

$$
l=W_{L_{1}}=W_{L_{2}}=\cdots=W_{L_{k_{l}-1}} \geq W_{L_{k_{l}}}>0
$$

Since $W_{L}=\sum_{j=1}^{k_{l}} W_{L_{j}}$, we have that $l\left(k_{l}-1\right) \leq W_{L} \leq l k_{l}$, attaining equality in the upper bound when $l$ divides $W_{L}$. Also note that $l\left(k_{l}-1\right)=0$ if and only if $k_{l}=1$ if and only if $l \geq W_{L}$; since the execution is finite, $W_{L}<\infty$, and thus $l\left(k_{l}-1\right)$ is positive for at most finitely many $l \in \mathbb{N}_{\geq 1}$. Therefore, considering all $l \in \mathbb{N}_{\geq 1}$,

$$
\begin{equation*}
W_{L} \geq \max _{l \in \mathbb{N} \geq 1} l\left(k_{l}-1\right) \tag{2.3}
\end{equation*}
$$

We will now derive a lower bound on $k_{l}$. Fix any $l \in \mathbb{N}_{\geq 1}$ and consider the concatenation $L=L_{1} \cdots L_{k_{l}}$ as defined above; this defines the concatenation $K=K_{1} \cdots K_{k_{l}}$, where each $K_{j}=K \cap L_{j}$. Since $0<N=|K|=\sum_{j=1}^{k_{l}}\left|K_{j}\right| \leq k_{l} \max _{j=1}^{k_{l}}\left|K_{j}\right|$, and since $k_{l} \in \mathbb{Z}$,

$$
\begin{equation*}
k_{l} \geq\left\lceil\frac{N}{\max _{j=1}^{k_{l}}\left|K_{j}\right|}\right\rceil \tag{2.4}
\end{equation*}
$$

We now weaken this inequality by bounding $\max _{j=1}^{k_{l}}\left|K_{j}\right|$ in terms of the function $q$.
Consider any $j \in\left[k_{l}\right]$ such that $\left|K_{j}\right|>0$ (there is at least one since $N>0$ ): applying Lemma 2.1, we have that

$$
W_{\overline{K_{j}}} \geq \max \left(0, Q_{\mathrm{in}}\left(K_{j}\right)-M_{p}\right)+\max \left(0, Q_{\mathrm{out}}\left(K_{j}\right)-M_{p}\right) \geq Q_{\mathrm{in}}\left(K_{j}\right)+Q_{\mathrm{out}}\left(K_{j}\right)-2 M_{p}
$$

since $\overline{K_{j}} \subseteq L_{j}, W_{\overline{K_{j}}} \leq W_{L_{j}} \leq l$, and it follows that $Q_{\text {in }}\left(K_{j}\right)+Q_{\text {out }}\left(K_{j}\right) \leq 2 M_{p}+l$. By hypothesis, for any nonempty $J \subseteq K,\lceil q(|J|)\rceil \leq Q_{\text {in }}(J)+Q_{\text {out }}(J)$ (the ceiling function is implied in the hypothesis); therefore, setting $J=K_{j},\left\lceil q\left(\left|K_{j}\right|\right)\right\rceil \leq 2 M_{p}+l$, and

$$
\begin{equation*}
\left.\max _{j=1}^{k_{l}}\left|K_{j}\right| \leq \max \{i \in[N]\} \mid\lceil q(i)\rceil \leq 2 M_{p}+l\right\} \tag{2.5}
\end{equation*}
$$

where the maximum on the right-hand side is well defined since $\left|K_{j}\right| \leq N$ and since $q$ is nondecreasing.

We obtain the lower bound (2.1) by substituting (2.5) into (2.4) into 2.3) and noting that $W_{p} \geq W_{L}$.

Now consider for the moment the two weaker hypotheses: $q(|J|) \leq Q_{\text {in }}(J)$ for all $\emptyset \neq$ $J \subseteq K$, or $q(|J|) \leq Q_{\text {out }}(J)$ for all $\emptyset \neq J \subseteq K$. In either case, we can repeat the preceding argument (with a few adjustments), redefining $W_{J}$ to count only Loads or Stores, and obtain the same lower bound as (2.1) except with $M_{p}+l$ appearing in place of $2 M_{p}+l$. These individual bounds on Loads and Stores may be useful for the analysis of machines with asymmetric Load/Store costs: for example, Stores are more expensive on some memory devices, and some algorithms can account for this [13]. However, we will not distinguish Load/Store costs further in the present work.

Proof of Corollary 2.1. Continue the notation introduced at the beginning of the proof of Theorem 2.1. We first simplify the maximum in the denominator of the fractional term in the right-hand side of (2.1). For any $l \in \mathbb{N}_{\geq 1}$ and any $i \in[N]$,

$$
\left\lceil i^{1 / \sigma}\right\rceil \leq 2 M_{p}+l \quad \Leftrightarrow \quad i^{1 / \sigma} \leq 2 M_{p}+l \quad \Leftrightarrow \quad i \leq\left(2 M_{p}+l\right)^{\sigma} \quad \Leftrightarrow \quad i \leq\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor
$$

where the three forward implications are immediate from properties of the floor function, of exponentiation, and the fact that $i \in \mathbb{Z}$; to see that the fourth inequality implies the first, substitute $i=\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor$ into the first inequality: $\left\lceil\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor^{1 / \sigma}\right\rceil \leq\left\lceil 2 M_{p}+l\right\rceil=2 M_{p}+l$. Therefore,

$$
\left.\max \{i \in[N]\} \mid\lceil q(i)\rceil \leq 2 M_{p}+l\right\}=\min \left(|K|,\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor\right)
$$

and substituting into the right-hand side of (2.1),

$$
\begin{aligned}
& \max _{l \in \mathbb{N} \geq 1} l\left(\left[\frac{N}{\max \left\{i \in[N] \mid\lceil q(i)\rceil \leq 2 M_{p}+l\right\}}\right]-1\right) \\
& \quad=\max _{l \in \mathbb{N} \geq 1} l\left(\left\lceil\frac{N}{\min \left(|K|,\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor\right)}\right]-1\right)=\max _{l \in \mathbb{N} \geq 1} l\left(\left\lceil\frac{N}{\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor}\right]-1\right)
\end{aligned}
$$

this is the first (tighter) lower bound in (2.2).
To simplify further, we will weaken this lower bound by fixing $l=M_{p} \geq 1$ :

$$
\max _{l \in \mathbb{N} \geq 1} l\left(\left\lceil\frac{N}{\left\lfloor\left(2 M_{p}+l\right)^{\sigma}\right\rfloor}\right\rceil-1\right) \geq M_{p}\left(\left\lceil\frac{N}{\left\lfloor\left(3 M_{p}\right)^{\sigma}\right\rfloor}\right\rceil-1\right)
$$

supposing that $N>\left\lfloor\left(3 M_{p}\right)^{\sigma}\right\rfloor$,

$$
M_{p}\left(\left\lceil\frac{N}{\left\lfloor\left(3 M_{p}\right)^{\sigma}\right\rfloor}\right\rceil-1\right) \geq M_{p} \cdot \frac{N}{2\left\lfloor\left(3 M_{p}\right)^{\sigma}\right\rfloor} \geq \frac{1}{2 \cdot 3^{\sigma}} \cdot \frac{N}{M_{p}^{\sigma-1}}
$$

Note that the second (looser) lower bound in (2.2) assumes the simpler, stronger condition $N \geq\left(4 M_{p}\right)^{\sigma}>\left\lfloor\left(3 M_{p}\right)^{\sigma}\right\rfloor$.

It is possible to strengthen the second conclusion of Corollary 2.1 by assuming stronger hypotheses regarding $N$ and $M_{p}$, or by taking more care approximating the maximum argument $l$, rather than simply setting $l=M_{p}$.

### 2.4 Related Work

As mentioned previously, our work was influenced in several ways by Hong-Kung's paper [21]. (Note that the following observations about our model and Hong-Kung's are meant to be informal.) Hong-Kung's machine model is like a special case of ours with only one processor: roughly speaking, Hong-Kung's 'red/blue pebble game' model can be extended to ours by introducing pebbles in different shades of red, modeling different processors' local memories. However, to model a parallel computation, we must address contention for shared resources, like data races when accessing shared memory, and these issues did not arise in Hong-Kung's (sequential) model. Even in the sequential (one processor) case, there are a few technical distinctions between the models that complicate comparisons.

For example, in our model, Operations' output values can overwrite input values (sometimes called 'sliding pebbles'), and we allow values to reside in local memory before and after execution. Both of these extensions to Hong-Kung's model were considered by Savage [31], who also developed a stronger lower bound argument that synthesized the theories of space and communication costs.

Another substantial distinction between our model and Hong-Kung's is that ours disallows recomputation - each evaluation is associated with exactly one Operation in any implementation. Generally speaking, implementations with recomputation can potentially reduce communication cost, and it is future work to address recomputation within our model. We note that several practical applications of this work, like matrix multiplication, have the special property that each evaluation has a single successor in the algorithm DAG, and since there is no reason to recompute a value more times than it is read by its successors (in particular, there is no reason to recompute a vertex with outdegree-0), recomputation is unhelpful for reducing communication in these algorithms. Disallowing recomputation has become a popular assumption in the related literature; Elango-Rastello-Pouchet-Ramanujam-Sadayappan [18] recently showed that Hong-Kung's lower bound framework can be tightened to exploit this assumption.

Another improvement to Hong-Kung's approach, proposed in earlier conference talks by us 16 and by Langou-Lowery [25], is changing the segment size, the parameter $l$ in Theorem 2.1. which we already noted can improve the lower bounds versus the typical choice of setting $l$ to equal the local memory size. In the case of matrix multiplication, Langou-Lowery showed that this can sharpen the lower bound by a factor of 2 .

## Chapter 3

## Modeling Loop Nests

In this chapter we use algorithms (defined in Section 2.2) to model loop nests that access arrays by linear functions of the loop iteration vector. A section-by-section outline is as follows.

- In Section 3.1, we give motivating examples of algorithms to develop intuition for how loop nests can be modeled by algorithms.
- In Section 3.2, we define Hölder-Brascamp-Lieb (HBL) interpretations of algorithms, which model the linear mapping between operations and variables in many loop nests; this is the main contribution of this chapter.
- Then in Section 3.3, we review related work, in particular the papers of Irony-ToledoTiskin [22] and of Ballard-Demmel-Holtz-Schwartz [5], which studied a special class of algorithms with HBL interpretations (matrix multiplication and related algorithms).


### 3.1 Motivating Examples

A practical motivation of this work is transforming computer programs, especially nestedloop programs. For example, multiplying $N$-by- $N$ matrices, $C=A \cdot B$, can be programmed with triply nested loops (assuming $C$ is zero-initialized),

$$
\begin{gathered}
\text { for } i=1: N, \quad \text { for } j=1: N, \quad \text { for } k=1: N, \\
C(i, j)=C(i, j)+A(i, k) * B(k, j)
\end{gathered}
$$

This program executes the inner-loop statement " $C(i, j)=C(i, j)+A(i, k) * B(k, j)$ " exactly $N^{3}$ times, with different indices $i, j, k$ - however, there is freedom to execute these statements in a different order besides the prescribed (lexicographical) order on $(i, j, k) \in[N]^{3}$ and we would like to reorder the $N^{3}$ statements to expose data reuse. Our communication lower bounds apply to all reorderings, so if our implementation attains the lower bound, then there is no hope in trying other reorderings to further reduce communication - instead, we must seek a new algorithm.

To be clear, the preceding pseudocode programming language is meant to be informal, and we will not define a conversion between pseudocode and algorithms (as defined in Section (2.2); moreover, terms like "program", "loop", "statement", etc., are undefined in our


$$
\left\{b_{i}=f_{i}\left(a_{i}\right) \mid i \in[N]\right\}
$$

Figure 3.1: Example 1 (see Section 3.1.1). The DAG depicts the case $N=3$.


Figure 3.2: Example 2 (see Section 3.1.2). The DAG depicts the case $N=3, m=3, n=2$.
model. Ultimately, to apply this work to program transformations, we will need to formalize the conversion between some given program representation and an algorithm in our model. Elango-Rastello-Pouchet-Ramanujam-Sadayappan recently demonstrated such a conversion [18], paving the way for future applications of communication lower bounds to program transformations. We discuss program transformations abstractly (in terms of algorithms) in Chapter 5

We begin with some motivating examples of algorithms with particular input- and outputpath properties that we will characterize in a more general fashion in Section 3.2. In this examples, we will denote variables by subscripted expressions $a_{i}, b_{i, j}$, etc.; variables denoted in this manner are assumed to be distinct unless specified otherwise. We will also assume a given parameter $N \in \mathbb{N}_{\geq 1}$.

### 3.1.1 Example 1

Our first example is $N$ independent unary operations - see Figure 3.1. The input- and output-paths in this algorithm's DAG are the singletons $\left\{a_{i}\right\}$ and $\left\{b_{i}\right\}$ for each $i \in[N]$, which are always disjoint. Thus, for any $K \subseteq F$, its input-path cutsize $Q_{\text {in }}(K)$ and outputpath cutsize $Q_{\text {out }}(K)$ (defined in Section 2.3) both equal $|K|$.

### 3.1.2 Example 2

For the second example, a variant on the first, we suppose each operation has input arity $m \in \mathbb{N}_{\geq 1}$ and output arity $n \in \mathbb{N}_{\geq 1}$ - see Figure 3.2. The input- and output-paths in


$$
\left\{a_{i}=f_{i}\left(a_{i-1}\right) \mid i \in[N]\right\}
$$

Figure 3.3: Example 3 (see Section 3.1.3). The DAG depicts the case $N=3$.


$$
\left\{a_{i, j}=f_{i, j}\left(a_{i-1, j}\right) \mid(i, j) \in[N]^{2}\right\}
$$

Figure 3.4: Example 4 (see Section 3.1.4). The DAG depicts the case $N=3$.
this algorithm's DAG are still (disjoint) singletons, except that there are now more: for any $K \subseteq F, Q_{\text {in }}(K)=m|K|$ and $Q_{\text {out }}(K)=n|K|$.

### 3.1.3 Example 3

The third example is $N$ completely dependent unary operations, whose induced (algorithm) DAG is a directed path - see Figure 3.3. Since there is only one input $a_{0}$ and one output $a_{N}$, there can be at most one disjoint input-path and output-path. Therefore, for any nonempty $K \subseteq F, Q_{\text {in }}(K)=Q_{\text {out }}(K)=1$.

### 3.1.4 Example 4

The fourth example synthesizes the first and third examples, with $N$ independent sets of $N$ completely dependent unary operations - see Figure 3.4. In this case, each of the $N$ paths cut (intersected) by $K$ contributes exactly one to both the input- and output-path cutsize. That is, for any nonempty $K \subseteq F, Q_{\text {in }}(K)=Q_{\text {out }}(K)=\left|\left\{j \mid f_{i, j} \in K\right\}\right|$. For example, this number is one when $K=\left\{f_{1,1}, f_{2,1}, f_{3,1}\right\}$, but three when $K=\left\{f_{1,1}, f_{1,2}, f_{1,3}\right\}$.

### 3.1.5 Example 5

The fifth example is a essentially a product of the third example with itself, defining a "diamond" DAG structure - see Figure 3.5. There is only one output variable: $O$ was unspecified, so by convention $O=\operatorname{im}(y) \backslash \operatorname{im}(x)=\left\{a_{N, N}\right\}$. Thus there can be at most one disjoint output-path. On the other hand, the $2 N$ input variables $I=\left\{a_{0, k}, a_{k, 0} \mid k \in[N]\right\}$


$$
\left\{a_{i, j}=f_{i, j}\left(a_{i-1, j}, a_{i, j-1}\right) \mid(i, j) \in[N]^{2}\right\}
$$

Figure 3.5: Example 5 (see Section 3.1.5). The DAG depicts the case $N=3$.


Variant 1


Variant 2


Variant 3

Figure 3.6: Example 5 variants, DAGs for the case $N=3$ (see Section 3.1.5).
are a great source of disjoint input-paths. Indeed, for any nonempty $K \subseteq F$,

$$
Q_{\text {in }}(K)=\left|\left\{i \mid f_{i, j} \in K\right\}\right|+\left|\left\{j \mid f_{i, j} \in K\right\}\right|, \quad Q_{\text {out }}(K)=1 .
$$

The diamond DAG is a good example to demonstrate the dependence of our approach on input- and and output-paths. Consider the three variants of this example in Figure 3.6 . Variant 1 flips the original (variant 0 ), with one input variable $I=\left\{a_{1,1}\right\}$ but $2 N-1$ output variables $O=\left\{a_{N, k}, a_{k, N} \mid k \in[N]\right\}$ : for any nonempty $K \subseteq F$,

$$
Q_{\text {in }}(K)=1, \quad Q_{\text {out }}(K)=\left|\left\{i \mid f_{i, j} \in K\right\}\right|+\left|\left\{j \mid f_{i, j} \in K\right\}\right|-\left|\left\{k \mid f_{k, k} \in K\right\}\right| .
$$

Variant 2 combines variants 0 and 1 , with $I=\left\{a_{0, k}, a_{k, 0} \mid k \in[N]\right\}$ and $O=\left\{a_{N, k}, a_{k, N} \mid\right.$ $k \in[N]\}$, so

$$
\begin{aligned}
Q_{\text {in }}(K) & =\left|\left\{i \mid f_{i, j} \in K\right\}\right|+\left|\left\{j \mid f_{i, j} \in K\right\}\right|, \\
Q_{\text {out }}(K) & =\left|\left\{i \mid f_{i, j} \in K\right\}\right|+\left|\left\{j \mid f_{i, j} \in K\right\}\right|-\left|\left\{k \mid f_{k, k} \in K\right\}\right| .
\end{aligned}
$$

Variant 3 has $I=\left\{a_{1,1}\right\}$ and $O=\left\{a_{N, N}\right\}$, so

$$
Q_{\text {in }}(K)=Q_{\text {out }}(K)=1
$$

We will see in Chapter 5 that our communication lower bounds are asymptotically attainable for variants 0,1 , and 2 , but not for variant 3 . For variant 3 , our communication lower bound is zero, but it is possible to show that a nonzero lower bound exists when $N$ is sufficiently large (see, e.g., [11]). Our approach inherits this weakness from Hong-Kung's approach: their lower bound [21, Theorem 3.1] is also trivial for this example (note that variant 0 is treated in [21, Section 7]). Savage [31] strengthened this aspect of Hong-Kung's approach by showing that Hong-Kung's hypothesis regarding dominator set size (analogous to $Q_{\text {in }}$ ) can be relaxed. While Savage's approach generally yields tighter lower bounds than HongKung's for examples like variant 3 with few inputs and outputs, due to technical distinctions between Hong-Kung's and Savage's models - Savage allows an instruction's output value to overwrite one of its input values ("sliding pebbles"; see also [19]) and allows inputs/outputs to reside in local memory before/after execution - Savage's lower bounds can also be smaller than Hong-Kung's. Hong-Kung's and Savage's approaches were compared and extended by Bilardi-Pietracaprina-D'Alberto [11]; it is future work to incorporate these refinements in our approach.

These diamond DAGs generalize from two dimensions to $r \in \mathbb{N}_{\geq 2}$ dimensions: e.g., the original (variant 0) becomes

$$
\left\{a_{i_{1}, \ldots, i_{r}}=f_{i_{1}, \ldots, i_{r}}\left(a_{i_{1}, \ldots, i_{j-1}, i_{j}-1, i_{j+1}, \ldots, i_{r}} \mid j \in[r]\right) \mid\left(i_{1}, \ldots, i_{r}\right) \in[N]^{r}\right\} .
$$

There is still only one output variable, $O=\left\{a_{N, \ldots, N}\right\}$, but there are now $r N^{r-1}$ input variables $I=\left\{a_{i_{1}, \ldots, i_{r}} \mid(\exists k \in[r]) i_{k}=0\right\}$. For any nonempty $K \subseteq F, Q_{\text {out }}(K)=1$ but

$$
Q_{\mathrm{in}}(K)=\sum_{j=1}^{r}\left|\left\{\left(i_{1}, \ldots, i_{j-1}, i_{j+1}, \ldots, i_{r}\right) \mid f_{i_{1}, \ldots, i_{r}} \in K\right\}\right| .
$$

Each summand (indexed by $j$ ) is the cardinality of a subset of $\mathbb{Z}^{r-1}$ obtained from a subset of $\mathbb{Z}^{r}$ by 'forgetting' the latter's elements' $j$-th components. These $r$ projections, functions from $\mathbb{Z}^{r}$ to $\mathbb{Z}^{r-1}$, are special examples of group homomorphisms from the Abelian group $\mathbb{Z}^{r}$ to the Abelian group $\mathbb{Z}^{r-1}$. We will revisit this example in Section 5.5.5.

### 3.1.6 Example 6

The next example, called a reduction, means any algorithm whose DAG has a single output - see Figure 3.7. For any nonempty $K \subseteq F, 1 \leq Q_{\text {in }}(K) \leq N$ and $Q_{\text {out }}(K)=1$. In Figure 3.7. we also show two special reductions which attain the lower and upper bounds, respectively, on $Q_{\text {in }}(K)$, for any nonempty $K \subseteq F$.

Similarly to reductions we have expansions - see Figure 3.8 For any nonempty $K \subseteq F$, $Q_{\text {in }}(K)=1$ and $1 \leq Q_{\text {out }}(K) \leq N$. In Figure 3.8, we show two special expansions which attain the lower and upper bounds, respectively, on $Q_{\text {out }}(K)$, for any nonempty $K \subseteq F$.

### 3.2 HBL Interpretations

In this section we define a group-theoretic interpretation for an algorithm called an $H B L$ (Hölder-Brascamp-Lieb) interpretation, which will describe a broad range of algorithms, and


$$
\begin{gathered}
V \supseteq\left\{s_{i}, t \mid i \in[N]\right\} \\
I=\left\{s_{i} \mid i \in[N]\right\}, \quad O=\{t\}
\end{gathered}
$$



Figure 3.7: Example 6, reduction (see Section 3.1.6). The top row gives two cartoon characterizations of the corresponding class of DAGs (when $N=3$ ) which are meant to be equivalent: the 'pencil' version is convenient in later examples (see Sections 5.5 .2 and 5.5.3) with multiple reductions intersecting at different angles. The bottom row gives two concrete special cases (when $N=3$ ) mentioned in Section 3.1.6.
let us systematically derive communication lower bounds and (sometimes) optimal implementations using techniques in Chapters 4 and 5 .

Computational data is commonly structured as arrays (scalars, vectors, matrices, etc.). Given some $r \in \mathbb{N}$ and $Z \subseteq \mathbb{Z}^{r}$, an array (of variables) is an injective function $A: Z \rightarrow V$. This is an informal analogy for array data structures commonly found in programming languages, as well as more general data structures with indirection like " $B$ (pointer $(i)$ )", provided that pointer $(\cdot)$ is injective; due to our abstract model, our arrays have no assumptions on rectangularity, type homogeneity, data-layout regularity, etc. When we consider multiple arrays, e.g., $A_{1}, A_{2}, A_{3}$, their elements may alias, e.g., $A_{1}(i)=A_{2}(i, j)=A_{3}(i, j, k)$ may all denote the same variable. We can disallow this possibility by assuming no aliasing, meaning that the images of all arrays are disjoint. However, as we will explain in Chapter 5, our communication lower bound analysis assumes maximal aliasing, meaning that the constants in our lower bounds could be sharpened in the absence of aliasing.

Consider any algorithm with parameters $V, F, x, y, I, O$ and its DAG ( $V, E$ ). Consider any

- $K$, nonempty subset of $F$,
- $m, r \in\{1,2, \ldots\}, r_{1}, \ldots, r_{m} \in\{0,1, \ldots\}$,
- $Z$, injection from $K$ to $\mathbb{Z}^{r}$,


$$
\begin{gathered}
V \supseteq\left\{s_{i}, t \mid i \in[N]\right\} \\
I=\{s\}, \quad O=\left\{t_{i} \mid i \in[N]\right\}
\end{gathered}
$$



Figure 3.8: Example 6, expansion (Section 3.1.6). The top row gives two cartoon characterizations of the corresponding class of DAGs (when $N=3$ ) which are meant to be equivalent: the 'pencil' version is convenient in later examples (see Sections 5.5.2 and 5.5.3 with multiple expansions intersecting at different angles. The bottom row gives two concrete special cases (when $N=3$ ) mentioned in Section 3.1.6.

- $\phi_{1}, \ldots, \phi_{m}$, group homomorphisms from $\mathbb{Z}^{r}$ to $\mathbb{Z}^{r_{1}}, \ldots, \mathbb{Z}^{r_{m}}$, and
- $A_{1}, \ldots, A_{m}$, injections from $\phi_{1}(\operatorname{im}(Z)), \ldots, \phi_{m}(\operatorname{im}(Z))$ to $I \cup O$.

For each $\emptyset \neq J \subseteq K$, letting $A(J)=\bigcup_{j=1}^{m} A_{j}\left(\phi_{j}(Z(J))\right)$, suppose there are disjoint paths in ( $V, E$ )

- from each $u \in A(J) \cap I$ to distinct $v \in \bigcup_{f \in J}\left\{x_{f, 1}, \ldots, x_{f, m_{f}}\right\}$ and
- to each $v \in A(J) \cap O$ from distinct $u \in \bigcup_{f \in J}\left\{y_{f, 1}, \ldots, y_{f, n_{f}}\right\}$.

We say that the given algorithm admits an HBL interpretation (with respect to these parameters).

For example, we see that the diamond DAG in Section 3.1.5, in particular, the variant in Figure 3.5, admits an HBL interpretation with

$$
\begin{gathered}
K=\left\{f_{i, j} \mid(i, j) \in[N]^{2}\right\}, \quad m=r=2, \quad r_{1}=r_{2}=1 \quad, Z\left(f_{i, j}\right)=(i, j), \\
\phi_{1}(i, j)=i, \quad \phi_{2}(i, j)=j, \quad A_{1}(i)=a_{i, 0}, \quad A_{2}(j)=a_{0, j} .
\end{gathered}
$$

For any $\emptyset \neq J \subseteq K$, we can identify the desired set of disjoint paths by tracing each of the $2 N$ inputs (positioned along the one of the coordinate axes in the plane) in a straight line until intersecting $J$ (or ignore those inputs if $J$ isn't intersected).

Having identified a set of disjoint input- and output-paths for each $\emptyset \neq J \subseteq K$, with each element of $A(J)$ included in a distinct path, we conclude that $Q_{\text {in }}(J)+Q_{\text {out }}(J) \geq|A(J)|$, the key property for proving the following result.

Theorem 3.1. If $\bigcap_{j=1}^{m} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$, then there exists $\sigma \in[1, m]$ such that

$$
\begin{equation*}
|J|^{1 / \sigma} \leq Q_{\text {in }}(J)+Q_{\text {out }}(J) \quad \text { for all } \emptyset \neq J \subseteq K \tag{3.1}
\end{equation*}
$$

Proof. See Section 5.1 .
Theorem 3.1 yields communication lower bounds via Corollary 2.1. That is, given an algorithm and HBL interpretation (continuing notation from earlier in Section 3.2), consider any implementation of that algorithm where some processor $p$, with $M_{p}$ local memory cells, is assigned the $N=|K|$ Operations corresponding to the evaluations $K \subseteq F$. By Corollary 2.1. if $N \geq\left(4 M_{p}\right)^{\sigma}$, then processor $p$ performs at least

$$
\frac{1}{2 \cdot 3^{\sigma}} \cdot \frac{N}{M_{p}^{\sigma-1}}=\Omega\left(\frac{N}{M_{p}^{\sigma-1}}\right)
$$

Loads and Stores between their first and last Operation corresponding to $K$.
As we will see in Section 5.1, $\sigma$ can be computed by solving the linear program,

$$
\begin{equation*}
\sigma=\min \left\{\sum_{j=1}^{m} s_{j} \mid s \in[0, \infty)^{m} \wedge\left(\forall H \leq \mathbb{Z}^{r}\right) \operatorname{rank}(H) \leq \sum_{j=1}^{m} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)\right\} \tag{3.2}
\end{equation*}
$$

### 3.3 Related Work

The class of algorithms studied in this chapter builds on earlier works by Irony-Toledo-Tiskin [22] and Ballard-Demmel-Holtz-Schwartz [5], which studied a special class of HBL interpretations, based on the Loomis-Whitney inequality - a special case of Hölder-Brascamp-Liebtype inequalities, which we define next in Chapter 4- in order to target matrix multiplication and related algorithms. While there are many similarities in the class of algorithms we consider, our models of computation have some key distinctions. In particular, our algorithm model is in some ways more constrained, avoiding the complication of 'R2/D2 operands' (mentioned in Section 2.3) at the cost of having to develop the input- and output-path-based analysis. We also address the "no interleaving" constraint in our earlier work [15] by allowing operations with multiple outputs (we are not aware of an earlier model that did this), thus enforcing atomicity of inner-loop statements.

## Chapter 4

## Hölder-Brascamp-Lieb-type Inequalities

Our communication lower bound analysis in Section 2.3 formalized the intuitive argument that a bounded number of operands implies a bounded number of operations can be performed. The key property of our target class of computations which enables this intuitive argument is that their operations and operands can be identified by tuples of integers, and the mapping from an operation to its operands can be modeled by linear functions between integer tuples. We exploit this algebraic relationship via the following result.

Theorem 4.1 ([15, Theorem 3.2]). Suppose we are given a positive integer $m$, nonnegative integers $r, r_{1}, \ldots, r_{m}$, and Abelian group homomorphisms $\phi_{1}, \ldots, \phi_{m}$ from $\mathbb{Z}^{r}$ to $\mathbb{Z}^{r_{1}}, \ldots, \mathbb{Z}^{r_{m}}$, respectively. Let $s_{1}, \ldots, s_{m}$ be nonnegative real numbers.

$$
\begin{array}{ll}
\text { If } & \operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) \\
\text { for all subgroups } H \text { of } \mathbb{Z}^{r},  \tag{4.2}\\
\text { then } & |E| \leq \prod_{j}\left|\phi_{j}(E)\right|^{s_{j}}
\end{array} \text { for all nonempty finite subsets } E \text { of } \mathbb{Z}^{r} .
$$

Proof. See Section 4.6
Each $x \in \mathbb{Z}^{r}$ identifies an operation whose operands are identified by $\phi_{1}(x), \ldots, \phi_{m}(x)$; thus, Theorem 4.1 gives an upper bound on the number $|E|$ of operations that can be performed when certain numbers $\left|\phi_{1}(E)\right|, \ldots,\left|\phi_{m}(E)\right|$ of operands are available.

We will actually prove a more general version of Theorem 4.1, stated as Theorem 4.2 in Section 4.2. While Theorem 4.1 suffices to analyze our target class of computations, we believe there are applications for the more general result in Theorem 4.2, in particular regarding groups with torsion.

This chapter is largely based on material that appeared in our earlier report 15, Sections 3 and 5]. A section-by-section outline of this chapter is as follows.

- In Section 4.1, we define Hölder-Brascamp-Lieb-type (HBL-type) inequalities, which generalize the inequalities (4.2) in the conclusion of Theorem 4.1 that we use to bound the number of doable operations given a bounded number of operands.
- In Section 4.2, we state our main results: Theorem 4.2 is a generalization of Theorem 4.1 , and Theorems 4.3 and 4.4 concern the decidability of the hypotheses (4.1) of Theorem 4.1.
- In Section 4.3, we review related work on HBL-type inequalities, in particular, the paper by Bennett-Carbery-Christ-Tao [9] that inspired this line of inquiry, as well as our previous report [15].
- In Section 4.4, we study the hypothesis of Theorem 4.2, a generalization of 4.1).
- In Section 4.5, we study the conclusion of Theorem 4.2, a generalization of 4.2).
- In Section 4.6, we synthesize the results developed in Sections 4.4 and 4.5 to conclude the proof of Theorem 4.2, and then finally of Theorem 4.1, as a special case.
- Lastly, in Section 4.7, we study the decidability of the hypothesis of Theorem 4.2, proving Theorems 4.3 and 4.4 .


### 4.1 Definitions and Notation

We define Hölder-Brascamp-Lieb-type inequalities using some measure-theoretic notation (definitions can be found in standard textbooks, e.g., $[30 \mid)$. Let $(X, \Sigma, \mu)$ be a measure space and let $f: X \rightarrow \mathbb{C} \cup\{\infty\}$ be a $\Sigma$-measurable function. For each $p \in(0, \infty]$, we define

$$
\|f\|_{L^{p}(X, \Sigma, \mu)}= \begin{cases}\left(\int_{(X, \Sigma)}|f|^{p} \mathrm{~d} \mu\right)^{1 / p} & p<\infty \\ \underset{(X, \Sigma, \mu)}{\operatorname{ess} \sup |f|} & p=\infty\end{cases}
$$

Our work considers the special case of discrete measure spaces, where $\Sigma=2^{X}$ is the discrete $\sigma$-algebra and $\mu=|\cdot|$ is counting measure; we specialize the preceding notation for this discrete setting, defining

$$
\|f\|_{\ell^{p}(X)}=\|f\|_{L^{p}\left(X, 2^{X},|\cdot|\right)}= \begin{cases}\left(\sum_{x \in X}|f(x)|^{p}\right)^{1 / p} & p<\infty \\ \sup _{x \in X}|f(x)| & p=\infty\end{cases}
$$

Additionally, it will be more convenient for us to work with the reciprocals $s=1 / p \in[0, \infty)$ of the exponents $p \in(0, \infty]$; in case $p=\infty$, we define $L^{1 / 0}=L^{\infty}$ and $\ell^{1 / 0}=\ell^{\infty}$.

A Hölder-Brascamp-Lieb-type inequality is defined by the following parameters:

- a positive integer $m \in\{1,2, \ldots\}$, and dedicating $j$ to range over $\{1, \ldots, m\}$,
- $m+1$ measure spaces $(G, \Sigma, \mu),\left(G_{j}, \Sigma_{j}, \mu_{j}\right)$ where $G, G_{j}$ are Abelian groups,
- $m$ Abelian group homomorphisms $\phi_{j}: G \rightarrow G_{j}$,
- $m \Sigma_{j}$-measurable functions $f_{j}: G_{j} \rightarrow[0, \infty]$,
- $m$ nonnegative real numbers $s_{j} \in[0, \infty)$, and
- an extended-real number $C \in \mathbb{R} \cup\{ \pm \infty\}$.

Together these parameters define an inequality,

$$
\begin{equation*}
\int_{(G, \Sigma)} \prod_{j} f_{j} \circ \phi_{j} \mathrm{~d} \mu \leq C \prod_{j}\left\|f_{j}\right\|_{L^{1 / s_{j}}\left(G_{j}, \Sigma_{j}, \mu_{j}\right)} \tag{4.3}
\end{equation*}
$$

where we adopt the convention that $0 \cdot \pm \infty=0$.
We will study families of Hölder-Brascamp-Lieb-type inequalities, where some parameters are fixed, while others are treated as variables. In particular, we will frequently fix the number $m$, the measure spaces $(G, \Sigma, \mu),\left(G_{j}, \Sigma_{j}, \mu_{j}\right)$, and the homomorphisms $\phi_{j}$, while varying the functions $f_{j}$, numbers $s_{j}$, and number $C$.

Since this work targets discrete measure spaces, our notation suppresses the discrete $\sigma$ algebras $\Sigma, \Sigma_{j}$ and counting measures $\mu, \mu_{j}$ to simplify in three ways: (1) we collect the fixed parameters as the triple $\mathcal{G}=\left(G,\left(G_{j}\right)_{j},\left(\phi_{j}\right)_{j}\right)$ called an HBL datum, where $m$ is implicit, (2) we use the shorthand notation $\|\cdot\|_{\ell^{p}(X)}=\|\cdot\|_{L^{p}\left(X, 2^{x},|\cdot|\right)}$ defined above, and (3) we avoid mentioning the (trivial) $\Sigma_{j}$-measurability of the functions $f_{j}$. We also collect the variable parameters $f_{j}, s_{j}$ as $m$-tuples called $m$-functions $f=\left(f_{j}\right)_{j} \in X_{j}\left(G_{j} \rightarrow[0, \infty]\right)$ and $m$-exponents $s=\left(s_{j}\right)_{j} \in[0, \infty)^{m}$.

### 4.2 Main Results

Consider an HBL datum $\mathcal{G}=\left(G,\left(G_{j}\right)_{j},\left(\phi_{j}\right)_{j}\right)$. Let $\mathcal{P}(\mathcal{G})$ denote the (convex) set of all $m$-exponents $s \in[0, \infty)^{m}$, such that

$$
\begin{equation*}
\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) \quad \text { for all finite rank subgroups } H \text { of } G ; \tag{4.4}
\end{equation*}
$$

additionally, for each $s \in[0, \infty)^{m}$, let $A(\mathcal{G}, s), B(\mathcal{G}, s), C(\mathcal{G}, s) \in \mathbb{R} \cup\{ \pm \infty\}$ denote the infima of $A, B, C \in \mathbb{R} \cup\{ \pm \infty\}$ such that

$$
\begin{align*}
&|H| \leq A \prod_{j}\left|\phi_{j}(H)\right|^{s_{j}} \text { for all finite subgroups } H \text { of } G,  \tag{4.5}\\
&|E| \leq B \prod_{j}\left|\phi_{j}(E)\right|^{s_{j}} \text { for all nonempty finite subsets } E \text { of } G,  \tag{4.6}\\
& \sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right) \leq C \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \text { for all m-functions }  \tag{4.7}\\
&
\end{align*}
$$

We will prove Theorem 4.1 as a special case of the following more general result, which relates the four objects $\overline{\mathcal{P}(\mathcal{G})}, A(\mathcal{G}, \cdot), B(\mathcal{G}, \cdot), C(\mathcal{G}, \cdot)$ defined by $\mathcal{G}$.

Theorem $4.2\left(\boxed{14}\right.$, Theorem 1.2]). Consider any $s \in[0, \infty)^{m}$ and let $r=\left(\min \left(1, s_{j}\right)\right)_{j}$. If $s \in \mathcal{P}(\mathcal{G})$, then $B(\mathcal{G}, s)=C(\mathcal{G}, s)=A(\mathcal{G}, r)$. Conversely, if $B(\mathcal{G}, s)<\infty$, then $s \in \mathcal{P}(\mathcal{G})$.

Proof. See Section 4.6.
A supplemental argument (see Lemmas 4.27and 4.39) shows that there is no loss of generality to apply Theorem 4.2 only to $m$-exponents $s \in[0,1]^{m}$, enabling a simpler formulation (see [14, Theorem 1.2]). While we ignore this simplification here, we will exploit it at many points in our analysis.

A practical question is whether the hypothesis $s \in \mathcal{P}(\mathcal{G})$ of Theorem 4.2 can be verified computationally for any $m$-exponent $s$. While the general case is open, we found an affirmative answer in a special case, when $G$ has finite rank: in this case, since ranks are natural numbers, and since $G$ 's rank is an upper bound on all its subgroups' ranks, only finitely many distinct inequalities can appear in (4.4). Our algorithm that proves the following result depends crucially on this fact.
Theorem 4.3 ( $[15$, Theorem 5.1]). If $\operatorname{rank}(G)<\infty$, then there exists an algorithm that computes a finite subset of the inequalities in (4.4) which define the same set $\mathcal{P}(\mathcal{G})$.

Proof. See Section 4.7, there are additional assumptions on how $\mathcal{G}$ is represented.
While this result suffices to verify whether a given $m$-exponent $s \in \mathcal{P}(\mathcal{G})$, it does not address the computability of the set of inequalities in (4.4). That is, one can in principle query the existence of a particular inequality in (4.4) by asking whether a subgroup $H$ of $G$ exists with certain values of $\operatorname{rank}(H), \operatorname{rank}\left(\phi_{j}(H)\right)$.
Theorem 4.4 ( $\sqrt{15}$, Theorem 5.9]). If $\operatorname{rank}(G)<\infty$, then there exists an algorithm to compute the set of all inequalities in (4.4) if and only if there exists an algorithm to determine whether an arbitrary system of polynomial equations over $\mathbb{Q}$ is solvable over $\mathbb{Q}$.

Proof. See Section 4.7.1. there are additional assumptions on how $\mathcal{G}$ is represented.
The question of whether the latter algorithm exists is known as Hilbert's Tenth Problem over $\mathbb{Q}$, and this question remains open (see, e.g., 28]).

### 4.3 Background and Related Work

Hölder-Brascamp-Lieb-type inequalities include those studied by Rogers-Hölder [29, 20] (see also [26]), Young [36] (see also [7, 12]), Loomis-Whitney [24], and Brascamp-Lieb [12]. Typically, studies of Hölder-Brascamp-Lieb-type inequalities have considered the continuum setting, where the measure spaces $(G, \Sigma, \mu),\left(G_{j}, \Sigma_{j}, \mu_{j}\right)$ are finite-dimensional $\mathbb{R}$-vector spaces with Lebesgue measure and the homomorphisms $\phi_{j}$ are $\mathbb{R}$-linear maps (see [6] for references). A pair of papers [10, 9] by Bennett-Carbery-Christ-Tao derived necessary and sufficient conditions for Hölder-Brascamp-Lieb-type inequalities to hold with $C<\infty$, both in the continuum setting and, in [9], in the case of discrete Abelian groups. In particular, [9, Theorem 2.4] is the starting point for Theorem 4.2 we now restate this result in the present notation, assuming a given HBL datum $\mathcal{G}$.
Proposition 4.1 ( 9 , Theorem 2.4]). Suppose $G, G_{j}$ are finitely generated and consider any $s \in[0, \infty)$. If $s \in \widehat{\mathcal{P}}(\mathcal{G})$, then $B(\mathcal{G}, s) \leq C(\mathcal{G}, s)<\infty$. Conversely, if $B(\mathcal{G}, s)<\infty$, then $s \in \mathcal{P}(\mathcal{G})$.

In their proof, Bennett-Carbery-Christ-Tao observed a link between the values $B(\mathcal{G}, s)$ and $C(\mathcal{G}, s)$ and the structure of the torsion subgroup $T(G)$ of $G$. We pursued this link in a paper [15], strengthening Proposition 4.1 with sharper bounds on these constants.

Proposition 4.2 ([15, Theorem 3.12]). Suppose $G, G_{j}$ are finitely generated and consider any $s \in[0, \infty)$. If $s \in \mathcal{P}(\mathcal{G})$, then $B(\mathcal{G}, s) \leq C(\mathcal{G}, s) \leq|T(G)|$. Conversely, if $B(\mathcal{G}, s)<\infty$, then $s \in \mathcal{P}(\mathcal{G})$.

The inequality in (4.6) for $E=T(G) \cap \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)$ shows that $B(\mathcal{G}, s) \geq|E| \geq 1$ for all mexponents $s$. Therefore, $B(\mathcal{G}, s)=C(\mathcal{G}, s)=|T(G)|$ whenever $T(G) \leq \cap_{j} \operatorname{ker}\left(\phi_{j}\right)$, e.g., when $G$ is torsion-free, or when all $G_{j}$ are torsion-free. In the general case, however, Proposition 4.2 leaves a gap of up to a factor of $|T(G)|$ between $B(\mathcal{G}, s), C(\mathcal{G}, s)$ and their upper bounds. This gap was closed by Christ [14], who precisely determined $B(\mathcal{G}, s), C(\mathcal{G}, s)$ in the more general setting where $G, G_{j}$ are not necessarily finitely generated. Christ's main result [14, Theorem 1.2], stated above as Theorem 4.2, thus completed the study of the optimal constants for discrete Abelian groups.

In [15], we also studied the computability of the hypothesis $s \in \mathcal{P}(\mathcal{G})$; we already summarized our findings in Theorems 4.3 and 4.4. We mention here that our algorithm for Theorem 4.3 is similar to an algorithm of Valdimarsson [35, Theorem 1.8], in the continuum setting, for computing the hypothesis of [9, Theorem 2.1].

### 4.4 Analysis of $\mathcal{P}$

Now we study the set $\mathcal{P}(\mathcal{G})$, defined by an HBL datum $\mathcal{G}$ via (4.4), in more depth. We let $\mathcal{H}^{*}(G), \mathcal{H}_{\mathbb{N}}^{*}(G), \mathcal{H}_{\mathrm{fg}}^{*}(G), \mathcal{H}_{0}^{*}(G), \mathcal{H}_{\mathrm{f}}^{*}(G)$ denote the sets of all, all finite-rank, all finitely generated, all rank-zero, and all finite subgroups of $G$, respectively (definitions can be found in standard texts, e.g., 23 ); these sets are partially ordered by containment,

$$
\mathcal{H}^{*}(G) \supseteq \mathcal{H}_{\mathbb{N}}^{*}(G) \supseteq \mathcal{H}_{\mathrm{fg}}^{*}(G) \cup \mathcal{H}_{0}^{*}(G) \supseteq \mathcal{H}_{\mathrm{fg}}^{*}(G) \cap \mathcal{H}_{0}^{*}(G)=\mathcal{H}_{\mathrm{f}}^{*}(G)
$$

Most results in this section are stated assuming a general HBL datum $\mathcal{G}$, although most will only be applied in the case where $G$ is finitely generated, i.e., where

$$
\mathcal{H}^{*}(G)=\mathcal{H}_{\mathbb{N}}^{*}(G)=\mathcal{H}_{\mathrm{fg}}^{*}(G) \supseteq \mathcal{H}_{0}^{*}(G)=\mathcal{H}_{\mathrm{f}}^{*}(G) .
$$

Given $\mathcal{G}$ and any $\mathcal{H} \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$, we define

$$
\begin{equation*}
\mathcal{P}(\mathcal{G}, \mathcal{H})=\left\{s \in[0, \infty)^{m} \mid(\forall H \in \mathcal{H}) \operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)\right\} . \tag{4.8}
\end{equation*}
$$

That is, $\mathcal{P}(\mathcal{G}, \mathcal{H})$ is defined by a nonempty, countable set of $\mathbb{Z}$-linear inequalities constraining $s \in \mathbb{R}^{m}$ : there are $m$ inequalities of the form $s_{j} \geq 0$ and countably many of the form $\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$, even if $\mathcal{H}$ is uncountable, since finite ranks are integers. Since $\mathcal{P}(\mathcal{G}, \mathcal{H})$ equals the intersection of closed half-spaces, it is closed and convex.

### 4.4.1 Group-Theoretic Tools

For the following three elementary results, let $G$ be any Abelian group with finite rank, not necessarily associated with an HBL datum.

Lemma 4.1. For any $H \leq G$, there exists $K \leq G$ such that $H \leq K$, $\operatorname{rank}(H)=\operatorname{rank}(K)$, and $G / K$ is torsion-free.

Proof. Define $K=\{x \mid(\exists a \in \mathbb{Z}) a x \in H\}$; it follows that $H \leq K \leq G$. If $x+K \in G / K$ such that $a(x+K)=0+K$ for some $a \in \mathbb{Z}$, then $y=a x \in K$; however, by definition of $K$, there exists $b \in \mathbb{Z}$ such that $b y \in H$, i.e., bax $\in H$, thus $x \in K$, so $x+K=0+K$, i.e., $G / K$ is torsion-free. Since each element of $K / H$ is a torsion element, $K / H$ is a torsion group, so $\operatorname{rank}(K / H)=0$, i.e., $\operatorname{rank}(H)=\operatorname{rank}(K)$.

Lemma 4.2. For any $H \leq G$, there exists $K \leq G$ such that $\operatorname{rank}(G)=\operatorname{rank}(H)+\operatorname{rank}(K)$ and $H \cap K=\{0\}$.

Proof. By Lemma 4.1, we can find $H \leq F \leq G$ such that $G / F$ is torsion-free and $\operatorname{rank}(H)=$ $\operatorname{rank}(F)$. Since $G / F$ has finite rank, let $U$ be the subgroup of $G / F$ generated by any $\operatorname{rank}(G / F)$ independent elements of $G / F: U$ is finitely generated with $\operatorname{rank}(U)=\operatorname{rank}(G / F)$, and we can write $U=L / F$ for some $F \leq L \leq G$ such that $\operatorname{rank}(L)=\operatorname{rank}(G)$. Since $L / F$ is a torsion-free finitely generated Abelian group, it is projective (moreover, free) as a $\mathbb{Z}$ module, meaning that for any surjective Abelian group homomorphism $\psi$ with domain $V$ and codomain $L / F$, there exists $W \leq V$ such that $W$ and $\operatorname{ker}(\psi)$ are complementary subgroups of $V$. Letting $V=L$, we have that $\operatorname{ker}(\psi)=F$, and so taking $K=W \leq L$, we have that $L=F+K$, so $\operatorname{rank}(G)=\operatorname{rank}(L)=\operatorname{rank}(F)+\operatorname{rank}(K)=\operatorname{rank}(H)+\operatorname{rank}(K)$; moreover, $H \cap K \leq F \cap K=\{0\}$, i.e., $H \cap K=\{0\}$.

Lemma 4.3. If $\phi$ is an Abelian group homomorphism with domain $G$ and if there exist $H, K \leq G$ such that $H \leq K \leq G$ and $\operatorname{rank}(H)=\operatorname{rank}(K)$, then $\operatorname{rank}(\phi(H))=\operatorname{rank}(\phi(K))$.

Proof. $0=\operatorname{rank}(K / H) \geq \operatorname{rank}(\phi(K) / \phi(H)) \geq 0$, so $0=\operatorname{rank}(\phi(K))-\operatorname{rank}(\phi(H))$.

### 4.4.2 Embedding $\mathbb{Z}$ in $\mathbb{Q}$

We digress briefly to mention an arguably simpler approach to studying $\mathcal{P}(\mathcal{G}, \cdot)$ which we will exploit later to prove Theorem 4.4.

Consider any HBL datum $\mathcal{G}=\left(G,\left(G_{j}\right)_{j},\left(\phi_{j}\right)_{j}\right)$, but treat $G, G_{j}$ as $\mathbb{Z}$-modules (defining multiplication in terms of addition) and $\phi_{j}$ as $\mathbb{Z}$-module homomorphisms. Let $f$ be a ring homomorphism from $\mathbb{Z}$ to an integral domain $\mathbb{D}$. Define $G^{\mathbb{D}}, G_{j}^{\mathbb{D}}, \phi_{j}^{\mathbb{D}}$ via extension of scalars: that is, regarding $\mathbb{D}$ as a $\mathbb{Z}$-module via $f, G^{\mathbb{D}}=G \otimes_{\mathbb{Z}} \mathbb{D}$ and $G_{j}^{\mathbb{D}}=G_{j} \otimes_{\mathbb{Z}} \mathbb{D}$ are $\mathbb{D}$-modules and $\phi_{j}^{\mathbb{D}}=\phi_{j} \otimes_{\mathbb{Z}} \mathrm{id}_{\mathbb{D}}$ are $\mathbb{D}$-module homomorphisms ( $\otimes_{\mathbb{Z}}$ denotes the tensor product of $\mathbb{Z}$-modules or $\mathbb{Z}$-module homomorphisms). We also introduce the notation $U \leq_{\mathbb{D}} V$, asserting that $U$ is a sub(-DD-)module of a $\mathbb{D}$-module $V$, and $\operatorname{rank}_{\mathbb{D}}(V)$, denoting the rank (over $\mathbb{D}$ ) of a $\mathbb{D}$-module $V$, but often omit the subscript- $\mathbb{D}$ s in the case $\mathbb{D}=\mathbb{Z}$ since ambiguity in this case causes no harm.

Lemma 4.4. Consider any $h, h_{j} \in \mathbb{N}$. If there exists $H \leq G$ such that $\operatorname{rank}(H)=h$ and each $\operatorname{rank}\left(\phi_{j}(H)\right)=h_{j}$, then there exists $V \leq_{\mathbb{D}} G^{\mathbb{D}}$ such that $\operatorname{rank}_{\mathbb{D}}(V)=h$ and each $\operatorname{rank}_{\mathbb{D}}\left(\phi_{j}^{\mathbb{D}}(V)\right)=h_{j}$. The converse also holds if $\mathbb{D}=\mathbb{Q}$.

Proof. Letting $V=H^{\mathbb{D}}=H \otimes_{\mathbb{Z}} \mathbb{D} \leq_{\mathbb{D}} G^{\mathbb{D}}$, the rank properties follow as consequences of the $\mathbb{Z}$-linearity of $\otimes_{\mathbb{Z}}$, i.e., if elements $e_{i} \in G$ are $\mathbb{Z}$-linearly independent then elements $e_{i} \otimes 1 \in G^{\mathbb{D}}$ are $\mathbb{D}$-linearly independent. To show the converse, note that $\mathbb{Q}$ is the fraction field of $\mathbb{Z}$. Therefore, for any $V \leq_{\mathbb{Q}} G^{\mathbb{Q}}$, any $\mathbb{Q}$-basis of $V$ (at least one exists) can be converted to a (same cardinality) $\mathbb{Z}$-basis of some $V_{\mathbb{Z}} \leq G$ such that $V=\left(V_{\mathbb{Z}}\right)^{\mathbb{Q}}$.

Therefore, it is possible to analyze $\mathcal{P}(\mathcal{G}, \cdot)$ via $\mathcal{P}_{\mathbb{Q}}\left(\mathcal{G}^{\mathbb{Q}}, \cdot\right)$, where $\mathcal{G}^{\mathbb{Q}}=\left(G^{\mathbb{Q}},\left(G_{j}^{\mathbb{Q}}\right)_{j},\left(\phi_{j}^{\mathbb{Q}}\right)_{j}\right)$ and $\mathcal{P}_{\mathbb{Q}}$ is obtained from $\mathcal{P}$ by replacing Abelian group ranks with $\mathbb{Q}$-module ranks in (4.8). The primary advantage of this approach is that $\mathbb{Q}$-modules are free, so some arguments simplify: in particular, submodules always have complements, avoiding the need for Lemmas 4.1 and 4.2, and proper submodules always have strictly smaller rank, avoiding the need for Lemma 4.3 .

### 4.4.3 Properties of $\mathcal{P}(\mathcal{G}, \mathcal{H})$

We now derive some basic results regarding $\mathcal{P}(\mathcal{G}, \mathcal{H})$ for any fixed $\mathcal{G}$ and $\mathcal{H} \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$.
Lemma 4.5. If $s \in \mathcal{P}(\mathcal{G}, \mathcal{H})$, then $t \in \mathcal{P}(\mathcal{G}, \mathcal{H})$ for all $s \leq t \in[0, \infty)^{m}$.
Proof. If the inequality induced by any $H \in \mathcal{H}$ is satisfied by $s$ then it is by $t$ as well.
Lemma 4.6. If $\mathcal{H}^{\prime} \subseteq \mathcal{H}$ then $\mathcal{P}\left(\mathcal{G}, \mathcal{H}^{\prime}\right) \supseteq \mathcal{P}(\mathcal{G}, \mathcal{H})$.
Proof. $\mathcal{P}\left(\mathcal{G}, \mathcal{H}^{\prime}\right)$ is defined by a subset of the inequalities that define $\mathcal{P}(\mathcal{G}, \mathcal{H})$.
There is no loss of generality to suppose each $\phi_{j}$ is surjective, i.e., $G_{j}=\phi_{j}(G)$.
Lemma 4.7. $\mathcal{P}(\mathcal{G}, \mathcal{H})=\mathcal{P}\left(\left(G,\left(\phi_{j}(G)\right)_{j},\left(\chi_{j}\right)_{j}\right), \mathcal{H}\right)$, where $\chi_{j}: G \rightarrow \phi_{j}(G): x \mapsto \phi_{j}(x)$.
Proof. For each $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ and each $j, \operatorname{rank}\left(\chi_{j}(H)\right)=\operatorname{rank}\left(\psi_{j}(H)\right)$.
Additionally, there is no loss of generality to ignore the rank-0 subgroups of $G$.
Lemma 4.8. $\mathcal{P}(\mathcal{G}, \mathcal{H})=\mathcal{P}\left(\mathcal{G}, \mathcal{H} \backslash \mathcal{H}_{0}^{*}(G)\right)$.
Proof. Each $H \in \mathcal{H}_{0}^{*}(G)$ induces the inequality $0 \leq \sum_{j} s_{j} \cdot 0$, satisfied by all $s \in \mathbb{R}^{m}$.
Note that if $\mathcal{H} \subseteq \mathcal{H}_{0}^{*}(G)$, e.g., $\mathcal{H}=\emptyset$ or $G \in \mathcal{H}_{0}^{*}(G)$, then $\mathcal{P}(\mathcal{G}, \mathcal{H})=[0, \infty)^{m}$.
Lemma 4.9. If $r=\sup _{H \in \mathcal{H}} \operatorname{rank}(H)<\infty$, then $\mathcal{P}\left(\mathcal{G}, \mathcal{H}^{\prime}\right)=\mathcal{P}(\mathcal{G}, \mathcal{H})$ for some finite $\mathcal{H}^{\prime} \subseteq$ $\mathcal{H}$.

Proof. The ranks of $H$ and its homomorphic images $\phi_{j}(H)$ are integers between 0 and $r \in \mathbb{N}$; therefore, the set $\mathcal{H}$, possibly infinite, generates at most $(r+1)^{m+1}<\infty$ distinct inequalities. The conclusion follows by choosing any subset $\mathcal{H}^{\prime}$ of $\mathcal{H}$ that induces these inequalities.

Lemma 4.10. If $\operatorname{rank}\left(H \cap \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)\right)=0$ for all $H \in \mathcal{H}$, then $[1, \infty)^{m} \subseteq \mathcal{P}(\mathcal{G}, \mathcal{H})$.

Proof. If $\operatorname{rank}(G)=0$, then $\mathcal{P}(\mathcal{G}, \mathcal{H})=[0, \infty)^{m}$ by Lemma 4.8, so the conclusion is immediate. If $\operatorname{rank}(G)>0$, then consider any $H \in \mathcal{H}$ and the corresponding homomorphism $\Phi: H \rightarrow \bigoplus_{j} \phi_{j}(H): x \mapsto\left(\phi_{j}(x)\right)_{j}$; let $K=\operatorname{ker}(\Phi)=H \cap \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right) \in \mathcal{H}_{0}^{*}(G)$, so by the first isomorphism theorem $(\Phi(H) \cong H / K)$, the fact that $\operatorname{rank}(\Phi(H)) \leq \operatorname{rank}\left(\oplus_{j} \phi_{j}(H)\right)=$ $\sum_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$, and the hypothesis,

$$
0=\operatorname{rank}(K)=\operatorname{rank}(H)-\operatorname{rank}(\Phi(H)) \geq \operatorname{rank}(H)-\sum_{j} \operatorname{rank}\left(\phi_{j}(H)\right)
$$

so for all $H \in \mathcal{H}, \operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$ holds with $s_{j}=1$ for each $j$. Thus by Lemma 4.5, $t \in \mathcal{P}(\mathcal{G}, \mathcal{H})$ for all $t \geq(1, \ldots, 1)$.

Lemma 4.11. If $\operatorname{rank}\left(H \cap \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)\right)>0$ for some $H \in \mathcal{H}$, then $\mathcal{P}(\mathcal{G}, \mathcal{H})=\emptyset$.
Proof. Observe that $\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=0<\operatorname{rank}(H)$ for all $s \in \mathbb{R}^{m}$.
Lemma 4.12. If $H \in \mathcal{H} \backslash \mathcal{H}_{0}^{*}(G)$, then $\sum_{j} s_{j} \geq 1$ for all $s \in \mathcal{P}(\mathcal{G}, \mathcal{H})$.
Proof. Rearranging $H$ 's induced inequality, $1 \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) / \operatorname{rank}(H) \leq \sum_{j} s_{j}$.
Lemma 4.13. Suppose $m \geq 2$ and consider any $s \in[0, \infty)^{m}$ with some component $s_{k}=0$. Then, $s \in \mathcal{P}(\mathcal{G}, \mathcal{H})$ if and only if $\left(s_{j}\right)_{j \neq k} \in \mathcal{P}\left(\left(G,\left(G_{j}\right)_{j \neq k}\left(\phi_{j}\right)_{j \neq k}\right), \mathcal{H}\right)$.

Proof. For any $H \in \mathcal{H}_{\mathbb{N}}^{*}(G), \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$.
Lemma 4.14. Suppose that there exist $k, l \in\{1, \ldots, m\}$ such that $k<l$ and $\operatorname{rank}\left(\operatorname{ker}\left(\phi_{k}\right) \cap\right.$ $\left.\operatorname{ker}\left(\phi_{l}\right)\right)=\operatorname{rank}\left(\operatorname{ker}\left(\phi_{k}\right)\right)=\operatorname{rank}\left(\operatorname{ker}\left(\phi_{l}\right)\right)<\infty$. Let $\mathcal{G}^{\prime}=\left(G,\left(G_{j}\right)_{j \neq l},\left(\phi_{j}\right)_{j \neq l}\right)$. If $s \in$ $\mathcal{P}(\mathcal{G}, \mathcal{H})$, then $t \in \mathcal{P}\left(\mathcal{G}^{\prime}, \mathcal{H}\right)$ where

$$
t=\left(s_{1}, \ldots, s_{k-1}, s_{k}+s_{l}, s_{k+1}, \ldots, s_{l-1}, s_{l+1}, \ldots, s_{m}\right)
$$

conversely, if $t \in \mathcal{P}\left(\mathcal{G}^{\prime}\right)$, then for all $\epsilon \in\left[0, t_{k}\right], s(\epsilon) \in \mathcal{P}(\mathcal{G})$ where

$$
s(\epsilon)=\left(t_{1}, \ldots, t_{k-1}, \epsilon, t_{k+1}, \ldots, t_{l-1}, t_{k}-\epsilon, t_{l}, \ldots, t_{m-1}\right)
$$

Proof. The inequality induced by any $H \leq G$ can be rewritten,

$$
\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=\left(s_{k}+s_{l}\right) \operatorname{rank}\left(\phi_{k}(H)\right)+\sum_{j \neq k, l} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)
$$

Lemma 4.15. Suppose that there exists $k \in\{1, \ldots, m\}$ such that $\operatorname{rank}\left(\phi_{k}(G)\right)=0$. Let $\mathcal{G}^{\prime}=\left(G,\left(G_{j}\right)_{j \neq k},\left(\phi_{j}\right)_{j \neq k}\right)$. If $s \in \mathcal{P}(\mathcal{G}, \mathcal{H})$, then $t=\left(s_{j}\right)_{j \neq k} \in \mathcal{P}\left(\mathcal{G}^{\prime}, \mathcal{H}\right)$; conversely, if $t \in \mathcal{P}\left(\mathcal{G}^{\prime}, \mathcal{H}\right)$, then for all $\epsilon \in[0, \infty), s(\epsilon) \in \mathcal{P}(\mathcal{G}, \mathcal{H})$ where

$$
s(\epsilon)=\left(t_{1}, \ldots, t_{k-1}, \epsilon, t_{k}, \ldots, t_{m-1}\right) .
$$

Proof. The inequality induced by any $H \leq G$ can be rewritten,

$$
\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) .
$$

Lemma 4.16. If there exists a $k \in\{1, \ldots, m\}$ such that $\operatorname{rank}\left(\phi_{k}(G)\right)=\operatorname{rank}(G)<\infty$ and if there exists an $H \in \mathcal{H}$ such that $\operatorname{rank}(H)>0$, then $e_{k} \in \mathcal{P}(\mathcal{G}, \mathcal{H})$.

Proof. For any $K \leq G, \operatorname{rank}\left(\phi_{k}(K)\right)=\operatorname{rank}(K)$, so the the inequality induced by $K$ can be rewritten,

$$
\operatorname{rank}(K) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right)=s_{k} \operatorname{rank}(K)+\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right)
$$

If $s_{k} \geq 1$, the preceding inequality holds for any $\left(s_{j}\right)_{j \neq k} \in[0, \infty)^{m-1}$.

### 4.4.4 Extreme Points of $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$

We continue studying arbitrary $\mathcal{G}$ and $\mathcal{H} \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$, which remain fixed across the following results. As mentioned earlier, $\mathcal{P}(\mathcal{G}, \mathcal{H})$ is a closed convex subset of $\mathbb{R}^{m}$. The following results consider intersecting $\mathcal{P}(\mathcal{G}, \mathcal{H})$ with the unit $m$-cube $[0,1]^{m}$. We will motivate the choice $[0,1]^{m}$ later in Lemma 4.27 for now, we only exploit the fact that this intersection results in a bounded set. Therefore, as a bounded and closed convex subset of $\mathbb{R}^{m}, \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$ equals the convex hull of its extreme points (in particular, such points exist).

Lemma 4.17. Suppose $m \geq 2$ and consider any extreme point sof $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$. Either some component $s_{k} \in\{0,1\}$ or there exists $H \in \mathcal{H}$ such that $0<\operatorname{rank}(H)<\operatorname{rank}(G)$ and $\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$.

Proof. If $\operatorname{rank}(G)=0$ then by Lemma 4.8, $\mathcal{P}(\mathcal{G}, \mathcal{H})=[0, \infty)^{m}$, so the first case applies. Otherwise, suppose $\operatorname{rank}(G)>0$ and $s \in(0,1)^{m}$. As an extreme point of $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap$ $[0,1]^{m}, s$ must lie at the intersection of $m \mathbb{Z}$-linearly independent hyperplanes that support $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, and since $s \in(0,1)^{m}$, these hyperplanes support $\mathcal{P}(\mathcal{G}, \mathcal{H})$ but not $[0,1]^{m}$. In other words, we can find $H_{1}, H_{2} \in \mathcal{H}$ such that $0<\operatorname{rank}\left(H_{i}\right)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)$ for both $i \in\{1,2\}$. Moreover, the $\mathbb{Z}$-linear independence of the corresponding hyperplanes means that for some $j, \operatorname{rank}\left(\phi_{j}\left(H_{1}\right)\right) \neq \operatorname{rank}\left(\phi_{j}\left(H_{2}\right)\right)$. Therefore, the ranks of $H_{1}$ and $H_{2}$ cannot both equal $\operatorname{rank}(G)$ : if $\operatorname{rank}(G)<\infty$, this follows from Lemma 4.3, and if $\operatorname{rank}(G)=\infty$, this follows since $H_{1}, H_{2} \in \mathcal{H}_{\mathbb{N}}^{*}(G)$. Therefore there exists $H_{i} \in \mathcal{H}$ such that $0<\operatorname{rank}\left(H_{i}\right)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)<\operatorname{rank}(G)$.

Lemma 4.18. If $\sum_{j} \operatorname{ker}\left(\phi_{j}\right) \in \mathcal{H}_{0}^{*}(G)$ and $s$ is an extreme point of $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, then $s \in\{0,1\}^{m}$.

Proof. By the first hypothesis, $\operatorname{rank}(H)=\operatorname{rank}\left(\phi_{j}(H)\right)$ for all $H \in \mathcal{H}$, and every positive rank subgroup induces the same inequality $1 \leq \sum_{j} s_{j}$. This inequality corresponds to a hyperplane in $\mathbb{R}^{m}$ with normal $(1, \ldots, 1)$, which intersects $[0,1]^{m}$ at $e_{1}, \ldots, e_{m}$ (the standard basis vectors in $\mathbb{R}^{m}$ ), thus all extreme points in $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$ are in $\{0,1\}^{m}$.

Lemma 4.19. If $s$ is an extreme point of $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$ with some component $s_{k} \in(0,1)$, then $\operatorname{rank}\left(\phi_{k}(H)\right)>0$ for all $H \in \mathcal{H}$ with $\operatorname{rank}(H)=\operatorname{rank}(G)$.

Proof. Suppose towards obtaining a contradiction that the hypotheses hold but rank $\left(\phi_{k}(H)\right)=$ 0 for some $H \in \mathcal{H}$ with $\operatorname{rank}(H)=\operatorname{rank}(G)$. By Lemma 4.3 and the definition of rank,
$\operatorname{rank}\left(\phi_{k}(K)\right)=0$ for all $K \in \mathcal{H} ;$ therefore, $\sum_{j} t_{j} \operatorname{rank}\left(\phi_{j}(K)\right)=t_{k} \cdot 0+\sum_{j \neq k} t_{j} \operatorname{rank}\left(\phi_{j}(K)\right)$ for all $K \in \mathcal{H}$ and all $t \in \mathcal{P}(\mathcal{G}, \mathcal{H})$. Thus, both

$$
t_{-}=\left(s_{1}, \ldots, s_{k-1}, 0, s_{k+1}, \ldots, s_{m}\right), \quad t_{+}=\left(s_{1}, \ldots, s_{k-1}, 1, s_{k+1}, \ldots, s_{m}\right)
$$

are in $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$ and $s \neq t_{ \pm}$by hypothesis. Since $s$ lies on an open line segment between $t_{ \pm} \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, it cannot be an extreme point, a contradiction.

Lemma 4.20. If $s$ is an extreme point of $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, $\sup _{H \in \mathcal{H}} \operatorname{rank}(H)<\infty$, and $\operatorname{rank}(H)<\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$ for all $H \in \mathcal{H} \backslash \mathcal{H}_{0}^{*}(G)$, then $s \in\{0,1\}^{m}$.

Proof. Suppose towards obtaining a contradiction that the hypotheses hold but some component $s_{k} \in(0,1)$. First, observe that

$$
t_{+}=\left(s_{1}, \ldots, s_{k-1}, 1, s_{k+1}, \ldots, s_{m}\right) \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}
$$

via Lemma 4.5, and that $s \neq t_{+}$. Second, under the present hypotheses that $s_{k}>0$ and the inequality generated by each $H \in \mathcal{H}$ with $\operatorname{rank}(H)>0$ is strict, it follows that for each such $H$ we can find a positive $\epsilon$ such that $\operatorname{rank}(H) \leq \sum_{j}\left(s_{j}-\epsilon\right) \operatorname{rank}\left(\phi_{j}(H)\right)$. From Lemma 4.9, it suffices to restrict our attention to the inequalities induced by some finite $\mathcal{H}^{\prime} \subseteq \mathcal{H}$, in which case only finitely many $\epsilon$ arise and their minimum $\epsilon^{*}$ is well-defined, and, in particular, positive. Thus we can define

$$
t_{-}=\left(s_{1}, \ldots, s_{k-1}, s_{k}-\epsilon^{*}, s_{k+1}, \ldots, s_{m}\right) \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}
$$

and $s \neq t_{-}$. Since $s$ lies on an open line segment between $t_{ \pm} \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, it cannot be an extreme point, a contradiction.

Lemma 4.21. If $s$ is an extreme point of $\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, $\sup _{H \in \mathcal{H}} \operatorname{rank}(H)<\infty$, and $\operatorname{rank}(H)<\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$ for all $H \in \mathcal{H} \backslash \mathcal{H}_{0}^{*}(G)$ such that $\operatorname{rank}(H)<\operatorname{rank}(G)$, then at most one component $s_{k} \notin\{0,1\}$.

Proof. Suppose towards obtaining a contradiction that the hypotheses hold but $s_{k}, s_{l} \in(0,1)$ for two distinct indices $k, l$. By Lemma 4.19, $\operatorname{rank}\left(\phi_{k}(G)\right), \operatorname{rank}\left(\phi_{l}(G)\right)>0$; in particular, $\operatorname{rank}(G)>0$. Define $t(\epsilon) \in \mathbb{R}^{m}$ componentwise by

$$
t_{j}(\epsilon)= \begin{cases}s_{k}+\epsilon \operatorname{rank}\left(\phi_{l}(G)\right) & j=k \\ s_{l}-\epsilon \operatorname{rank}\left(\phi_{k}(G)\right) & j=l \\ s_{j} & j \notin\{k, l\}\end{cases}
$$

We now determine constraints on $\epsilon$ so that $t(\epsilon) \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$. For any $H \in \mathcal{H}$, we will require that

$$
\begin{aligned}
\operatorname{rank}(H) & \leq \sum_{j} t_{j}(\epsilon) \operatorname{rank}\left(\phi_{j}(H)\right) \\
& =\epsilon\left(\operatorname{rank}\left(\phi_{l}(G)\right) \operatorname{rank}\left(\phi_{k}(H)\right)-\operatorname{rank}\left(\phi_{k}(G)\right) \operatorname{rank}\left(\phi_{l}(H)\right)\right)+\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) ;
\end{aligned}
$$

defining

$$
\begin{aligned}
d(H) & =\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)-\operatorname{rank}(H) \in[0, \infty), \\
r(H) & =\operatorname{rank}\left(\phi_{l}(G)\right) \operatorname{rank}\left(\phi_{k}(H)\right)-\operatorname{rank}\left(\phi_{k}(G)\right) \operatorname{rank}\left(\phi_{l}(H)\right) \in\left\{-\operatorname{rank}(G)^{2}, \ldots, \operatorname{rank}(G)^{2}\right\},
\end{aligned}
$$

the constraint on $\epsilon$ that each $H \in \mathcal{H}$ induces is $-\epsilon r(H) \leq d(H)$, or more strictly, $|\epsilon||r(H)| \leq$ $d(H)$. Actually, there is no constraint on $\epsilon$ in the case $r(H)=0$, e.g., when $\operatorname{rank}(H)=$ $\operatorname{rank}(G)$ or $\operatorname{rank}(H)=0$. Thus it suffices to consider only the subset $\mathcal{H}^{\prime} \subseteq \mathcal{H}$ comprising the $H \in \mathcal{H}$ such that $0<\operatorname{rank}(H)<\operatorname{rank}(G)$. By the same argument used to prove Lemma 4.9, $d(H)$ takes on only finitely many distinct values as $H$ varies over $\mathcal{H}^{\prime}$, thus we can define $d^{*}=\min _{H \in \mathcal{H}^{\prime}} d(H)$. It remains to enforce that $t \in[0, \infty)^{m}$, and, in particular, that $t \in[0,1]^{m}: 0 \leq s_{k}+\epsilon \operatorname{rank}\left(\phi_{l}(G)\right) \leq 1$ and $0 \leq s_{l}-\epsilon \operatorname{rank}\left(\phi_{k}(G)\right) \leq 1$. Combining these constraints on $\epsilon$, replacing $d(H)$ by its minimum $d^{*}$, replacing $|r(H)|, \operatorname{rank}\left(\phi_{k}(G)\right)$, and $\operatorname{rank}\left(\phi_{l}(G)\right)$ by a common upper bound $\operatorname{rank}(G)^{2}>0$, and symmetrizing (tightening at most one of) the upper and lower bounds,

$$
|\epsilon| \leq \frac{\min \left\{d^{*}, s_{k}, 1-s_{k}, s_{l}, 1-s_{l}\right\}}{\operatorname{rank}(G)^{2}} ;
$$

by hypothesis, $d^{*}>0$ and $s_{k}, s_{l} \in(0,1)$, so the right-hand side is positive. Thus, there exists $\epsilon_{ \pm} \in \mathbb{R}$ such that $\epsilon_{-}<0<\epsilon_{+}, t\left(\epsilon_{ \pm}\right) \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, and $s \neq t\left(\epsilon_{ \pm}\right)$. Since $s$ lies on an open line segment between $t\left(\epsilon_{ \pm}\right) \in \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$, it cannot be an extreme point, a contradiction.

### 4.4.5 Properties of $\mathcal{P}(\mathcal{G})=\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathbb{N}}^{*}(G)\right)$

For any $\operatorname{HBL}$ datum $\mathcal{G}$, the set $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathbb{N}}^{*}(G)\right)$ is of principal interest, so we introduce the notation $\mathcal{P}(\mathcal{G})=\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathbb{N}}^{*}(G)\right)$.

Before studying $\mathcal{P}(\mathcal{G})$, we mention that we could also define $\mathcal{P}(\mathcal{G})$ in terms of $\mathcal{H}_{\mathrm{fg}}^{*}(G)$.
Lemma 4.22. $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathbb{N}}^{*}(G)\right)=\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathrm{fg}}^{*}(G)\right)$.
Proof. We have right-to-left inclusion from the fact that $\mathcal{H}_{\mathrm{fg}}^{*}(G) \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$. To observe left-to-right inclusion, for each $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, pick any $\operatorname{rank}(G)$ independent elements of $G$ to generate $K \in \mathcal{H}_{\mathrm{fg}}^{*}(G)$ such that $\operatorname{rank}(H)=\operatorname{rank}(K)$; by Lemma 4.3, for each $j$, $\operatorname{rank}\left(\phi_{j}(H)\right)=\operatorname{rank}\left(\phi_{j}(K)\right)$.

Lemma 4.23. If $G \in \mathcal{H}_{\mathbb{N}}^{*}(G), H \in \mathcal{H}_{\mathrm{fg}}^{*}(G)$, and $\operatorname{rank}(H)=\operatorname{rank}(G)$, then $\mathcal{P}(\mathcal{G})=$ $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathrm{fg}}^{*}(H)\right)$.

Proof. We have right-to-left inclusion from the fact that $\mathcal{H}_{\mathrm{fg}}^{*}(H) \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$. (Note that all subgroups of $H$ are finitely generated.) To observe left-to-right inclusion, consider each $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ : letting $F=K \cap H \in \mathcal{H}_{\mathrm{fg}}^{*}(H)$, it follows from the argument in Lemma 4.3 that $\operatorname{rank}(K)=\operatorname{rank}(F)$ and for each $j, \operatorname{rank}\left(\phi_{j}(K)\right)=\operatorname{rank}\left(\phi_{j}(F)\right)$.

A key property of $\mathcal{P}(\mathcal{G})$ is that it factors in the following sense. Given the HBL datum $\mathcal{G}=\left(G,\left(G_{j}\right),\left(\phi_{j}\right)\right)$ and any $H \leq G$, we define two associated HBL data called the factor data: the restricted datum $\left.\mathcal{G}\right|_{H}$ and the quotient datum $\left.\mathcal{G}\right|^{H}$ :

$$
\begin{array}{ll}
\left.\mathcal{G}\right|_{H}=\left(H,\left(\phi_{j}(H)\right)_{j},\left(\chi_{j}\right)_{j}\right), & \chi_{j}: H \rightarrow \phi_{j}(H): x \mapsto \phi_{j}(x), \\
\left.\mathcal{G}\right|^{H}=\left(G / H,\left(G_{j} / \phi_{j}(H)\right)_{j},\left(\psi_{j}\right)_{j}\right), & \psi_{j}: G / H \rightarrow G_{j} / \phi_{j}(H): x+H \mapsto \phi_{j}(x)+\phi_{j}(H) .
\end{array}
$$

In our analysis, the symbols $\chi_{j}$ and $\psi_{j}$ will always denote the homomorphisms associated with the factor data $\left.\mathcal{G}\right|_{H}$ and $\left.\mathcal{G}\right|^{H}$, i.e., $\chi_{j}$ and $\psi_{j}$ are implied by the notations $\left.\mathcal{G}\right|_{H}$ and $\left.\mathcal{G}\right|^{H}$. Also note an asymmetry in these definitions: $\chi_{j}$ is always surjective, whereas $\psi_{j}$ is surjective only if $\phi_{j}$ is; this property is exploited later in Lemma 4.43 however, the homomorphisms' codomains play no role in the analysis of $\mathcal{P}$.

Lemma 4.24. If $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, then $\mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right) \cap \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right) \subseteq \mathcal{P}(\mathcal{G}) \subseteq \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right)$.

Proof. Consider any $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ : it follows that $K \cap H \in \mathcal{H}_{\mathbb{N}}^{*}(H)$ and $(K+H) / H \in$ $\mathcal{H}_{\mathbb{N}}^{*}(G / H)$, even if $\operatorname{rank}(G)=\infty$. Now consider any $s \in \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right) \cap \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right)$ : it follows that $\operatorname{rank}(K \cap H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\chi_{j}(K \cap H)\right)$ and $\operatorname{rank}((K+H) / H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\psi_{j}((K+H) / H)\right)$. Combining these results,

$$
\begin{aligned}
\operatorname{rank}(K) & =\operatorname{rank}(K \cap H)+\operatorname{rank}((K+H) / H) \\
& \leq \sum_{j} s_{j} \operatorname{rank}\left(\chi_{j}(K \cap H)\right)+\sum_{j} s_{j} \operatorname{rank}\left(\psi_{j}((K+H) / H)\right) \\
& =\sum_{j} s_{j}\left(\operatorname{rank}\left(\chi_{j}(K \cap H)\right)+\operatorname{rank}\left(\psi_{j}((K+H) / H)\right)\right) \\
& =\sum_{j} s_{j}\left(\operatorname{rank}\left(\phi_{j}(K \cap H)\right)+\operatorname{rank}\left(\left(\phi_{j}(K)+\phi_{j}(H)\right) / \phi_{j}(H)\right)\right) \\
& =\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right) .
\end{aligned}
$$

Repeating for all $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, we conclude that $s \in \mathcal{P}(\mathcal{G})$; the first inclusion follows by repeating this argument for every $s \in \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right) \cap \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right)$.

To show the second inclusion, consider any $s \in \mathcal{P}(\mathcal{G})$. By definition, $\mathcal{H}_{\mathbb{N}}^{*}(H) \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$. Since $\phi_{j}(K)=\chi_{j}(K)$ for any $K \in \mathcal{H}_{\mathbb{N}}^{*}(H), s \in \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right)$. Considering every $s \in \mathcal{P}(\mathcal{G})$ we conclude $\mathcal{P}(\mathcal{G}) \subseteq \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right)$.

Lemma 4.25. If $s \in \mathcal{P}(\mathcal{G}), H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, and $\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$, then $s \in$ $\mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right) \cap \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right)$.

Proof. By Lemma $4.24, s \in \mathcal{P}(\mathcal{G}) \subseteq \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right)$, so it remains to show that $s \in \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right)$. Each element of $\mathcal{H}_{\mathbb{N}}^{*}(G / H)$ can be written as $(K+H) / H$ for some $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, even if $\operatorname{rank}(G)=$ $\infty$. Additionally, $\operatorname{rank}((K+H) / H)=\operatorname{rank}(K+H)-\operatorname{rank}(H)$ and, for each $j, \operatorname{rank}\left(\psi_{j}((K+\right.$ $H) / H))=\operatorname{rank}\left(\left(\phi_{j}(K)+\phi_{j}(H)\right) / \phi_{j}(H)\right)=\operatorname{rank}\left(\phi_{j}(K)+\phi_{j}(H)\right)-\operatorname{rank}\left(\phi_{j}(H)\right)$. Combining
these identities and the hypotheses that $s \in \mathcal{P}(\mathcal{G})$ and $\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$,

$$
\begin{aligned}
\operatorname{rank}((K+H) / H) & =\operatorname{rank}(K+H)-\operatorname{rank}(H) \\
& \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K+H)\right)-\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) \\
& =\sum_{j} s_{j}\left(\operatorname{rank}\left(\phi_{j}(K+H)\right)-\operatorname{rank}\left(\phi_{j}(H)\right)\right) \\
& =\sum_{j} s_{j} \operatorname{rank}(\psi((K+H) / H)) .
\end{aligned}
$$

Repeating for all $(K+H) / H \in \mathcal{H}_{\mathbb{N}}^{*}(G / H)$, we conclude that $s \in \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right)$.
Lemma 4.26. Consider any $s \in[0, \infty)^{m}$ with some component $s_{k} \geq 1$ and let $K=\operatorname{ker}\left(\phi_{k}\right)$. Then $s \in \mathcal{P}(\mathcal{G})$ if and only if $\left(s_{j}\right)_{j \neq k} \in \mathcal{P}\left(\left(G,\left(G_{j}\right)_{j \neq k},\left(\left.\phi_{j}\right|_{K}\right)_{j \neq k}\right)\right.$.

Proof. Sufficiency follows from the fact that, for any $H \in \mathcal{H}_{\mathbb{N}}^{*}(K), \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=$ $\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$. To show necessity, consider any $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ : let $F_{1}=H \cap K$ and apply Lemma 4.2 to define $F_{2} \leq H$ such that $\operatorname{rank}(H)=\operatorname{rank}\left(F_{1}\right)+\operatorname{rank}\left(F_{2}\right)$. It follows that both $F_{1}, F_{2} \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, so

$$
\begin{aligned}
& \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) \\
& \quad \geq \operatorname{rank}\left(\phi_{k}(H)\right)+\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) \geq \operatorname{rank}\left(\phi_{k}\left(F_{2}\right)\right)+\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}\left(F_{1}\right)\right) \\
& \quad \geq \operatorname{rank}\left(\phi_{k}\left(F_{2}\right)\right)+\operatorname{rank}\left(F_{1}\right)=\operatorname{rank}\left(F_{2}\right)+\operatorname{rank}\left(F_{1}\right)=\operatorname{rank}(H)
\end{aligned}
$$

Lemma 4.27. If $s \in \mathcal{P}(\mathcal{G})$, then $t \in \mathcal{P}(\mathcal{G})$ for all $\left(\min \left\{1, s_{j}\right\}\right)_{j} \leq t \in[0, \infty)^{m}$.
Proof. First suppose some $s_{k}>1$ and define $t=\left(s_{1}, \ldots, s_{k-1}, 1, s_{k+1}, \ldots, s_{m}\right)$; we now show that $t \in \mathcal{P}(\mathcal{G})$. Consider any $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ and let $F=H \cap \operatorname{ker}\left(\phi_{k}\right)$; since $F \in \mathcal{H}_{\mathbb{N}}^{*}(G)$,

$$
\operatorname{rank}(F) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(F)\right)=s_{k} \cdot 0+\sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(F)\right)=\sum_{j \neq k} t_{j} \operatorname{rank}\left(\phi_{j}(F)\right) .
$$

By Lemma 4.2, there exists $K \leq H$ such that $\operatorname{rank}(H)=\operatorname{rank}(F)+\operatorname{rank}(K)$ and $\operatorname{rank}(F \cap$ $K)=0$. Since $K \cap F=K \cap\left(H \cap \operatorname{ker}\left(\phi_{k}\right)\right)=K \cap \operatorname{ker}\left(\phi_{k}\right)$ and $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$,

$$
\operatorname{rank}(K)=\operatorname{rank}(K)-\operatorname{rank}\left(K \cap \operatorname{ker}\left(\phi_{k}\right)\right)=\operatorname{rank}\left(\phi_{k}(K)\right) ;
$$

moreover, since $t_{k}=1$,

$$
\operatorname{rank}(H)=\operatorname{rank}(K)+\operatorname{rank}(F) \leq t_{k} \operatorname{rank}\left(\phi_{k}(K)\right)+\sum_{j \neq k} t_{j} \operatorname{rank}\left(\phi_{j}(F)\right) \leq \sum_{j} t_{j} \operatorname{rank}\left(\phi_{k}(H)\right)
$$

Thus we have identified $t \in \mathcal{P}(\mathcal{G})$ which, relative to $s$, has one fewer component greater than one. We can repeat this argument for each such component to show that $\left(\min \left\{1, s_{j}\right\}\right)_{j} \in$ $\mathcal{P}(\mathcal{G})$. The conclusion for all $t \geq\left(\min \left\{1, s_{j}\right\}\right)_{j}$ follows from Lemma 4.5.

That is, $\mathcal{P}(\mathcal{G})$ is completely specified by $\mathcal{P}(\mathcal{G}) \cap[0,1]^{m}$. As mentioned in Section 4.4.4, we prefer to analyze $\mathcal{P}(\mathcal{G}) \cap[0,1]^{m}$ because it equals the convex hull of its extreme points.

Lemma 4.28. If $s \in[0,1]^{m}$ has exactly one component $s_{k} \in(0,1)$ and there exists $H \in$ $\mathcal{H}_{\mathbb{N}}^{*}(G)$ such that $\operatorname{rank}(G)=\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$ and $\operatorname{rank}\left(\phi_{k}(H)\right)>0$, then there exists $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ such that $\operatorname{rank}(K)>\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right)$.

Proof. Note that these hypotheses imply $G \in \mathcal{H}_{\mathbb{N}}^{*}(G)$. Letting $I=\left\{j \mid s_{j}=1\right\}$, we rewrite a hypothesis regarding $H$,

$$
\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=s_{k} \operatorname{rank}\left(\phi_{k}(H)\right)+\sum_{i \in I} s_{i} \operatorname{rank}\left(\phi_{i}(H)\right)
$$

If $I=\emptyset$ then since $\operatorname{rank}\left(\phi_{k}(H)\right)>0, s_{k}=\operatorname{rank}(H) / \operatorname{rank}\left(\phi_{k}(H)\right)$. If $\operatorname{rank}\left(\phi_{k}(H)\right)=$ $\operatorname{rank}(H)$ then $s_{k}=1$, a contradiction to the hypothesis that $s_{k} \in(0,1)$. So it must be that $\operatorname{rank}\left(\phi_{k}(H)\right)<\operatorname{rank}(H), \operatorname{sorank}\left(\operatorname{ker}\left(\phi_{k}\right)\right)>0$; taking $K=\operatorname{ker}\left(\phi_{k}\right) \in \mathcal{H}_{\mathbb{N}}^{*}(G)$,

$$
\operatorname{rank}(K)>0=s_{k} \operatorname{rank}\left(\phi_{k}(K)\right)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{k}(K)\right) .
$$

Otherwise, supposing $I \neq \emptyset$, by the hypotheses $s_{k}$ and $\operatorname{rank}\left(\phi_{k}(H)\right)>0$,

$$
\sum_{i \in I} \operatorname{rank}\left(\phi_{i}(H)\right)=\sum_{i \in I} s_{i} \operatorname{rank}\left(\phi_{i}(H)\right)=\operatorname{rank}(H)-s_{k} \operatorname{rank}\left(\phi_{k}(H)\right)<\operatorname{rank}(H) .
$$

We now show that the existence of $H$ implies $\operatorname{rank}\left(\bigcap_{i \in I} \operatorname{ker}\left(\phi_{i}\right)\right)>0$. Consider the homomorphism $\Phi: G \rightarrow \bigoplus_{i \in I} \phi_{i}(G): x \mapsto\left(\phi_{i}(x)\right)_{i \in I}$, and let $K=\operatorname{ker}(\Phi) \in \mathcal{H}_{\mathbb{N}}^{*}(G)$. Since $\bigcap_{i \in I} \operatorname{ker}\left(\phi_{i}\right)=K, \operatorname{rank}(H)=\operatorname{rank}(G)<\infty, \operatorname{rank}(\Phi(G)) \leq \operatorname{rank}\left(\bigoplus_{i \in I} \phi_{i}(G)\right)=\sum_{i \in I} \operatorname{rank}\left(\phi_{i}(G)\right)$, and Lemma 4.3.

$$
\begin{aligned}
\operatorname{rank}(K) & =\operatorname{rank}(G)-\operatorname{rank}(\Phi(G)) \geq \operatorname{rank}(G)-\sum_{i \in I} \operatorname{rank}\left(\phi_{i}(G)\right) \\
& =\operatorname{rank}(H)-\sum_{i \in I} \operatorname{rank}\left(\phi_{i}(H)\right)=s_{k} \operatorname{rank}\left(\phi_{k}(H)\right)>0
\end{aligned}
$$

Because $\phi_{i}(K)=\{0\}$ for all $i \in I$,

$$
\begin{aligned}
\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right) & =s_{k} \operatorname{rank}\left(\phi_{k}(K)\right)+\sum_{i \in I} \operatorname{rank}\left(\phi_{i}(K)\right) \\
& =s_{k} \operatorname{rank}\left(\phi_{k}(K)\right) \leq s_{k} \operatorname{rank}(K)
\end{aligned}
$$

Since $s_{k}<1$ and $\operatorname{rank}(K)>0$, we have that $\operatorname{rank}(K)>\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right)$.
Lemma 4.29. If $\operatorname{rank}(G)<\infty$ and $s$ is an extreme point of $\mathcal{P}(\mathcal{G}) \cap[0,1]^{m}$, then $s \in\{0,1\}^{m}$ or there exists $H \in \mathcal{H}_{\mathbb{N}}^{*}(G) \backslash \mathcal{H}_{0}^{*}(G)$ such that $\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)<\operatorname{rank}(G)$.

Proof. Suppose toward a contradiction that neither case applies, i.e., that some component $s_{k} \in(0,1)$ and there does not exist $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ such that $0<\operatorname{rank}(H)<\operatorname{rank}(G)$ and $\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$. Under these hypotheses, by Lemma 4.21, there is exactly one such index $k$. Additionally, under these hypotheses, by the contrapositive of Lemma 4.20
(noting that $\operatorname{rank}(G)<\infty$ implies $\left.\sup _{H \in \mathcal{H}_{\mathbb{N}}^{*}(G)} \operatorname{rank}(H)<\infty\right)$, there exists $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ with $\operatorname{rank}(H)=\operatorname{rank}(G)$ and $\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$. By Lemma 4.19, $\operatorname{rank}\left(\phi_{k}(H)\right)>0$, and finally by Lemma 4.28, we obtain $K \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ such that $\operatorname{rank}(K)>\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right)$, which contradicts the hypothesis that $s \in \mathcal{P}(\mathcal{G})$.

The product case refers to a special class of HBL data where $\mathcal{P}(\mathcal{G})$ is defined by a simpler set of inequalities. Here we consider the product case in the context of finitely generated torsion-free Abelian groups; a more general setting was considered in [9, Section 7]. The following result is a special case of the more general result [9, Proposition 7.1].

Lemma 4.30. Consider any HBL datum $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)$ for some $r \in \mathbb{N}$; suppose there exists a $\mathbb{Z}$-basis $B=\left\{e_{1}, \ldots, e_{r}\right\}$ of $\mathbb{Z}^{r}$ such that each $\operatorname{ker}\left(\phi_{j}\right)=\left\langle K_{j}\right\rangle$ for some $K_{j} \subseteq B$. Let $\mathcal{H}=\left\{H_{1}=\left\langle e_{1}\right\rangle, \ldots, H_{r}=\left\langle e_{r}\right\rangle\right\}$. Then for any $s \in[0, \infty)^{m}, s \in \mathcal{P}(\mathcal{G})$ if and only if $s \in \mathcal{P}(\mathcal{G}, \mathcal{H})$.

Proof. Necessity follows from Lemma 4.6, i.e., the inequalities defining $\mathcal{P}(\mathcal{G}, \mathcal{H})$ also appear in the definition of $\mathcal{P}(\mathcal{G})$.

To show sufficiency, rewrite the inequality for each $H_{i} \in \mathcal{H}$,

$$
1=\operatorname{rank}\left(H_{i}\right) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)=\sum_{j} s_{j} \delta_{i, j}
$$

where $\delta_{i, j}=1_{K_{j}}\left(e_{i}\right)$, i.e., $\delta_{i, j}=1$ if $e_{i} \in K_{j}$, and $\delta_{i, j}=0$ otherwise. Now consider any $H \leq \mathbb{Z}^{r}$, and let $h=\operatorname{rank}(H)$. Let $M_{H}$ be an $r$-by- $h \mathbb{Z}$-basis matrix of $H$ (coordinates with respect to $B$ ). Since the rank of $M_{H}$ is $h$, at least one $h$-by- $h$ minor of $M_{H}$ is nonzero; pick any such minor and let $R \subseteq\{1, \ldots, r\}$ comprise the $h$ corresponding row indices of $M_{H}$. It follows that $\operatorname{rank}\left(\phi_{j}(H)\right) \geq \sum_{i \in R} \delta_{i, j}$. Therefore,

$$
\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right) \geq \sum_{j} s_{j} \sum_{i \in R} \delta_{i, j}=\sum_{i \in R} \sum_{j} s_{j} \delta_{i, j} \geq \sum_{i \in R} \operatorname{rank}\left(H_{i}\right)=\sum_{i \in R} 1=h=\operatorname{rank}(H) .
$$

### 4.5 Analysis of $A, B$, and $C$

It is useful to define $A(\mathcal{G}, s, H), B(\mathcal{G}, s, E)$, and $C(\mathcal{G}, s, f)$ to denote the values of $A, B$, and $C$ in the individual inequalities in (4.5), 4.6), and (4.7) such that equality holds. Thus, $A(\mathcal{G}, s), B(\mathcal{G}, s)$, and $C(\mathcal{G}, s)$ are the suprema of $A(\mathcal{G}, s, H), B(\mathcal{G}, s, E)$, and $C(\mathcal{G}, s, f)$ over the corresponding sets of subgroups $H$, subsets $E$, and $m$-functions $f$.

In the case of $A$, recalling from Section 4.4 that $\mathcal{H}_{\mathrm{f}}^{*}(G)$ comprises the finite subgroups of $G$, for any $s \in[0, \infty)^{m}$ and any $H \in \mathcal{H}_{\mathrm{f}}^{*}(G)$, we define

$$
\begin{equation*}
A(\mathcal{G}, s, H)=B(\mathcal{G}, s, H)=\frac{|H|}{\prod_{j}\left|\phi_{j}(H)\right|^{s_{j}}} \in(0, \infty) \tag{4.9}
\end{equation*}
$$

Similarly in the case of $B$, letting $\mathcal{E}^{*}(G)$ denote the set of all nonempty finite subsets of $G$, for any $s \in[0, \infty)^{m}$ and any $E \in \mathcal{E}^{*}(G)$, we define

$$
\begin{equation*}
B(\mathcal{G}, s, E)=C\left(\mathcal{G}, s,\left(1_{\phi_{j}(E)}\right)_{j}\right)=\frac{|E|}{\prod_{j}\left|\phi_{j}(E)\right|^{s_{j}}} \in(0, \infty) \tag{4.10}
\end{equation*}
$$

In the cases of both (4.5) and (4.6), there are well-defined minimum values $A$ and $B$ for each individual inequality (for each $H \in \mathcal{H}_{\mathrm{f}}^{*}(G)$ and $E \in \mathcal{E}^{*}(G)$, respectively). However, in the case of (4.7), there are two degenerate classes of $m$-functions where a minimum value for $C$ is not well defined. In the first degenerate class of $m$-functions $f, \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}=0$, i.e., one or more functions $f_{j}=0$; in this case, equality holds in 4.7) with any $C$. In the second degenerate class of $m$-functions $f, \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}\left(G_{j}\right)}}=\infty$; in this case, equality holds in 4.7) with any $C>0$, recalling the convention that $0 \cdot \infty=0$. There is no loss of generality to ignore these classes of $m$-functions since they do not influence the value of $C(\mathcal{G}, s)$ : to see this, consider the $m$-function with $f_{j}=1_{\left\{\phi_{j}(0)\right\}}=1_{\{0\}}$ : 4.7) holds with equality with $C=1$, so $C(\mathcal{G}, s) \geq 1$, and we can take $C$ in both degenerate cases to be less than 1 . Since $\Pi_{j}$ is a finite product, $\Pi_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \in(0, \infty)$ is equivalent to $\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \in(0, \infty)$ for all $j$. With these observations, we collect the nondegenerate $m$-functions,

$$
\mathcal{F}^{*}(\mathcal{G}, s)=\left\{\left(f_{j}\right)_{j} \in \underset{j}{X}\left(G_{j} \rightarrow[0, \infty]\right) \mid\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \in(0, \infty)\right\} .
$$

With this notation, for any $s \in[0, \infty)^{m}$ and any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, we can now define

$$
\begin{equation*}
C(\mathcal{G}, s, f)=\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \in[0, \infty] \tag{4.11}
\end{equation*}
$$

note that we still permit $m$-functions $f$ where $\operatorname{supp}\left(\prod_{j} f_{j} \circ \phi_{j}\right)=\bigcap_{j} \operatorname{supp}\left(f_{j} \circ \phi_{j}\right)=\emptyset$ : these are the only $f \in \mathcal{F}^{*}(\mathcal{G}, s)$ where $C(\mathcal{G}, s, f)=0$. In conclusion, we have

$$
A(\mathcal{G}, s)=\sup _{H \in \mathcal{H}_{f}^{*}(G)} A(\mathcal{G}, s, H), \quad B(\mathcal{G}, s)=\sup _{E \in \mathcal{E}^{*}(G)} B(\mathcal{G}, s, E), \quad C(\mathcal{G}, s)=\sup _{f \in \mathcal{F}^{*}(\mathcal{G}, s)} C(\mathcal{G}, s, f) .
$$

### 4.5.1 Basic Properties of $A(\mathcal{G}, s), B(\mathcal{G}, s)$, and $C(\mathcal{G}, s)$

Lemma 4.31. For any $s \in[0, \infty)^{m}, 1 \leq A(\mathcal{G}, s) \leq B(\mathcal{G}, s) \leq C(\mathcal{G}, s) \leq|G|$.
Proof. Considering the trivial subgroup $\{0\} \in \mathcal{H}_{\mathrm{f}}^{*}(G)$, we have that $A(\mathcal{G}, s) \geq A(\mathcal{G}, s,\{0\})=$ 1. Then we have that $A(\mathcal{G}, s) \leq B(\mathcal{G}, s) \leq C(\mathcal{G}, s)$, since their defining sets are nested. Lastly, note that for any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$,

$$
\begin{aligned}
\sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right) & \leq|G| \sup \left\{\prod_{j} f_{j}\left(\phi_{j}(x)\right) \mid x \in G\right\} \leq|G| \prod_{j} \sup \left\{f_{j}\left(\phi_{j}(x)\right) \mid x \in G\right\} \\
& =|G| \prod_{j} \sup \left\{f_{j}\left(y_{j}\right) \mid y_{j} \in \phi_{j}(G)\right\} \leq|G| \prod_{j} \sup \left\{f_{j}\left(y_{j}\right) \mid y_{j} \in G_{j}\right\} \\
& =|G| \prod_{j}\left\|f_{j}\right\|_{\ell^{\infty}\left(G_{j}\right)} \leq|G| \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}\left(G_{j}\right)}},
\end{aligned}
$$

so $C(\mathcal{G}, s) \leq|G|$.
A function $F:[0, \infty)^{m} \rightarrow[0, \infty]$ is antitone if, for each $s, t \in[0, \infty)^{m}, s \leq t$ (componentwise) implies $F(s) \geq F(t)$.

Lemma 4.32. $A(\mathcal{G}, s), B(\mathcal{G}, s)$ and $C(\mathcal{G}, s)$ are antitone functions of $s \in[0, \infty)^{m}$.

Proof. For each $H \in \mathcal{H}_{\mathrm{f}}^{*}(G)$, since $1 \leq\left|\phi_{j}(H)\right| \leq|H|<\infty$ for each $j, A(\mathcal{G}, s, H)$ is a nonincreasing function of each component $s_{j} \in[0, \infty)$, and thus $A(\mathcal{G}, s)$ is an antitone function of $s \in[0, \infty)^{m}$. Similarly, for each $E \in \mathcal{E}^{*}(G)$, since $1 \leq\left|\phi_{j}(E)\right| \leq|E|<\infty$ for each $j, B(\mathcal{G}, s, E)$ is a nonincreasing function of each component $s_{j} \in[0, \infty)$, and thus $B(\mathcal{G}, s)$ is an antitone function of $s \in[0, \infty)^{m}$. More generally, for any discrete measure space $\left(X, 2^{X},|\cdot|\right)$, any $g: X \rightarrow \mathbb{C} \cup\{\infty\}$, and any $p, q \in[0, \infty), p \leq q$ implies $\|g\|_{\ell^{1 / p}(X)} \leq\|g\|_{\ell^{1 / q}(X)}$. Thus, for any $m$-exponents $s, t$ and any $m$-function $f, s \leq t$ implies that $\prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \leq$ $\prod_{j}\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}$ and, if $f \in \mathcal{F}^{*}(\mathcal{G}, t)$, then $f \in \mathcal{F}^{*}(\mathcal{G}, s)$. We conclude that $C(\mathcal{G}, s) \geq C(\mathcal{G}, t)$ by taking suprema over $\mathcal{F}^{*}(\mathcal{G}, s)$ and $\mathcal{F}^{*}(\mathcal{G}, t)$.

Lemma 4.33. If $s \in[0, \infty)^{m}$ then $A(\mathcal{G}, s) \leq|T(G)|$.
Proof. If $H \in \mathcal{H}_{\mathrm{f}}^{*}(G)$ then $H \leq T(G)$.
Lemma 4.34. If $K \leq \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right),|K|<\infty$, and $s \in[0, \infty)^{m}$, then $A(\mathcal{G}, s) \geq|K|$.
Proof. $A(\mathcal{G}, s) \geq A(\mathcal{G}, s, K)=|K| \Pi_{j}\left|\phi_{j}(K)\right|^{-s_{j}}=|K| \cdot 1$.
Lemma 4.35. If $K \leq \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right),|K|=\infty$, and $s \in[0, \infty)^{m}$, then $B(\mathcal{G}, s)=\infty$.
Proof. $K$ has a sequence $E_{1}, E_{2}, \ldots$ of nonempty subsets of increasing cardinality; each $E_{i} \in \mathcal{E}^{*}(G)$, so $B\left(\mathcal{G}, s, E_{i}\right)=\left|E_{i}\right| \prod_{j}\left|\phi_{j}\left(E_{i}\right)\right|^{-s_{j}}=\left|E_{i}\right|$, so $B(\mathcal{G}, s)=\infty$.

Lemma 4.36. If $I=\left\{j \mid s_{j} \geq 1\right\}$, $K=\bigcap_{i \in I} \operatorname{ker}\left(\phi_{i}\right) \in \mathcal{H}_{\mathrm{f}}^{*}(G)$, and $s \in[0, \infty)^{m}$, then $C(\mathcal{G}, s) \leq|K|$.

Proof. If $I=\emptyset$, then $K=G$, and we have that $C(\mathcal{G}, s) \leq|G|$ by Lemma 4.31, regardless of $|K|$. Otherwise, for any $m$-tuple $f=\left(f_{j}\right)_{j}$ of functions $f_{j}: G_{j} \rightarrow \mathbb{C} \cup\{\infty\}$, because $G \rightarrow X_{j} G_{j}: x \mapsto\left(\phi_{i}(x)\right)_{i \in I}$ is $|K|$-to- 1,

$$
\prod_{i \in I}\left\|f_{i}\right\|_{\ell^{1 / s_{i}}\left(G_{i}\right)} \geq \prod_{i \in I}\left\|f_{i}\right\|_{\ell^{1}\left(G_{i}\right)}=\sum_{\left(y_{i}\right)_{i} \in X_{i \in I}} \prod_{G_{i}}\left|f_{i}\left(y_{i}\right)\right| \geq \frac{1}{|K|} \sum_{x \in G} \prod_{i \in I}\left|f_{i}\left(\phi_{i}(x)\right)\right| .
$$

Therefore,

$$
\begin{aligned}
\sum_{x \in G} \prod_{j}\left|f_{j}\left(\phi_{j}(x)\right)\right| & \leq\left(\prod_{j \notin I}\left\|f_{j}\right\|_{\ell \infty\left(G_{j}\right)}\right) \sum_{x \in G} \prod_{i \in I}\left|f_{i}\left(\phi_{i}(x)\right)\right| \\
& \leq\left(\prod_{j \notin I}\left\|f_{j}\right\|_{\ell \infty\left(G_{j}\right)}\right)\left(|K| \prod_{i \in I}\left\|f_{i}\right\|_{\ell^{1 / s_{i}}\left(G_{i}\right)}\right) \leq|K| \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)},
\end{aligned}
$$

so, specializing for $m$-functions $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, we see that $C(\mathcal{G}, s) \leq|K|$.
Lemma 4.37. If $s \in[0, \infty)^{m}$, $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, and $\emptyset \neq I \subseteq\{1, \ldots, m\}$, then $C(\mathcal{G}, s, f) \leq$ $C\left(\left(\phi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I},\left(f_{i}\right)_{i \in I}\right)$, with equality when $s_{j}=0$ or $\phi_{j}(G)=\{0\}$ for each $j \notin I$.

Proof. The inequality follows from the definition of $C(\mathcal{G}, s, f)$,

$$
\begin{aligned}
C(\mathcal{G}, s, f) & =\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sum_{x \in G}\left(\prod_{i \in I} \frac{f_{i}\left(\phi_{i}(x)\right)}{\left\|f_{i}\right\|_{\ell^{1 / s_{i}}\left(G_{i}\right)}}\right)\left(\prod_{j \notin I} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}\right) \\
& \leq \sum_{x \in G}\left(\prod_{i \in I} \frac{f_{i}\left(\phi_{i}(x)\right)}{\left\|f_{i}\right\|_{\ell^{1 / s_{i}}\left(G_{i}\right)}}\right) \cdot 1=C\left(\left(\phi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I},\left(f_{i}\right)_{i \in I}\right),
\end{aligned}
$$

and the fact that the projection $\mathcal{F}^{*}(\mathcal{G}, s) \rightarrow \mathcal{F}^{*}\left(\left(\phi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I}\right): f \mapsto\left(f_{i}\right)_{i \in I}$ is surjective.
To show equality when $s_{j}=0$ or $\phi_{j}(G)=\{0\}$, for all $j \notin I$, define the $m$-function $g$ with $g_{i}=f_{i}$ for each $i \in I$ and $g_{j}(y)=\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \cdot 1_{\phi_{j}(G)}(y)$ for each $j \notin I$. For each $j \notin I$, we observe in both cases $s_{j}=0$ and $\phi_{j}(G)=\{0\}$ that $\left\|g_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}=\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$, and, for all $x \in G, g_{j}\left(\phi_{j}(x)\right)=\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$. Thus we confirm that $g \in \mathcal{F}^{*}(\mathcal{G}, s)$, and, for all $x \in G$, that

$$
\prod_{j \notin I} \frac{g_{j}\left(\phi_{j}(x)\right)}{\left\|g_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\prod_{j \notin I} \frac{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=1 .
$$

Therefore,

$$
C(\mathcal{G}, s, f) \leq C(\mathcal{G}, s, g)=C\left(\left(\phi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I},\left(g_{i}\right)_{i \in I}\right)=C\left(\left(\phi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I},\left(f_{i}\right)_{i \in I}\right)
$$

since for any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$ we can find such a $g \in \mathcal{F}^{*}(\mathcal{G}, s)$, the second part of the conclusion (equality) follows by taking suprema over $f \in \mathcal{F}^{*}(\mathcal{G}, s)$.

Lemma 4.38. For all $s \in[0, \infty)^{m}, C(\mathcal{G}, s)$ equals the supremum of $C\left(\mathcal{G}, s, f^{E}\right)$ over all $E \in \mathcal{E}^{*}(G)$ and $f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)$ where each $\operatorname{supp}\left(f_{j}^{E}\right)=\phi_{j}(E)$.

Proof. Consider any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, and for each $E \in \mathcal{E}^{*}(G)$, query whether $\sum_{x \in E} \Pi_{j} f_{j}\left(\phi_{j}(x)\right)>$ 0 . If so, it follows that $\left\|\left.f_{j}\right|_{\phi_{j}(E)}\right\|_{\ell^{1 / s_{j}}\left(\phi_{j}(E)\right)}>0$; thus, we can define another $m$-function $f^{E}$ with components

$$
f_{j}^{E}(y)=\frac{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}{\left\|\left.f_{j}\right|_{\phi_{j}(E)}\right\|_{\ell^{1 / s_{j}}\left(\phi_{j}(E)\right)}} \cdot f_{j}(y) \cdot 1_{\phi_{j}(E)}(y)
$$

it follows that $\left\|f_{j}^{E}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}=\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$, and so $f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)$ as well. Otherwise, if $\sum_{x \in E} \prod_{j} f_{j}\left(\phi_{j}(x)\right)=0$, set each $f_{j}^{E}=0$, so $f^{E} \notin \mathcal{F}^{*}(\mathcal{G}, s)$. By the definitions of $C(\mathcal{G}, s, \cdot)$ and the Lebesgue integral for counting measure,

$$
\begin{aligned}
C(\mathcal{G}, s, f) & =\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sup _{E \in \mathcal{E}^{*}(G)} \sum_{x \in E} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sup _{\substack{E \in \mathcal{E}^{*}(G) \\
f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)}} \sum_{x \in E} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \\
& =\sup _{\substack{E \in \mathcal{E}^{*}(G) \\
f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)}} \sum_{x \in G} \prod \frac{f_{j}^{E}\left(\phi_{j}(x)\right)}{\left\|f_{j}^{E}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sup _{\substack{E \in \mathcal{E}^{*}(G) \\
f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)}} C\left(\mathcal{G}, s, f^{E}\right) .
\end{aligned}
$$

The conclusion follows by taking a supremum over $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, redefining its associated $m$-functions $f^{E}$ appropriately.

### 4.5.2 Properties of $B(\mathcal{G}, s)$ and $C(\mathcal{G}, s)$ when $s \in[0,1]^{m}$

Lemma 4.39. If $s \in[0, \infty)^{m}$ and $r=\left(\min \left(1, s_{j}\right)\right)_{j}$, then $B(\mathcal{G}, s)=B(\mathcal{G}, r)$ and $C(\mathcal{G}, s)=$ $C(\mathcal{G}, r)$.

Proof. We prove the second conclusion first, and then specialize that proof to obtain the first conclusion.

Suppose $s_{k}>1$ for some $k \in\{1, \ldots, m\}$ and define $t \in[0, \infty)^{m}$ such that $t_{k}=1$ and $t_{j}=s_{j}$ for $j \neq k$. Consider any $E \in \mathcal{E}^{*}(G)$ and any $m$-function $f$ where for each $j$, $0 \neq f_{j}<\infty$ and $\operatorname{supp}\left(f_{j}\right)=\phi_{j}(E)$. Therefore, $f \in \mathcal{F}^{*}(\mathcal{G}, r)$ for all $r \in[0, \infty)^{m}$, in particular, for $r=s$ and $r=t$. For each $y \in \phi_{k}(E)$, let $E_{y}=E \cap \phi_{k}^{-1}(y)$, and define another $m$-function $f^{y}$ with components

$$
f_{j}^{y}(z)= \begin{cases}f_{j}(z) \neq 0 & z \in \phi_{j}\left(E_{y}\right) \\ 0 & z \notin \phi_{j}\left(E_{y}\right)\end{cases}
$$

It follows that $f_{k}=\sum_{y \in \phi_{k}(E)} f_{k}^{y}, \bigcap_{j} \operatorname{supp}\left(f_{j}^{y} \circ \phi_{j}\right)=E_{y}$, and $\operatorname{supp}\left(f_{j}^{y}\right)=\phi_{j}\left(E_{y}\right)$; in particular, $\operatorname{supp}\left(f_{k}^{y}\right)=\{y\}$. Therefore, for each $j, 0<\left\|f_{j}^{y}\right\|_{\ell^{1 / s_{j}\left(G_{j}\right)}} \leq\left\|f_{j}\right\|_{\ell^{1 / s_{j}\left(G_{j}\right)}}<\infty$, and

$$
\begin{aligned}
C\left(\mathcal{G}, s, f^{y}\right) & =\sum_{x \in G} \prod_{j} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}^{y}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sum_{x \in E_{y}} \prod_{j} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}^{y}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \\
& =\left(\sum_{x \in E_{y}} \prod_{j \neq k} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}^{y}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}\right) \cdot \frac{f_{k}^{y}(y)}{\left\|f_{k}^{y}\right\|_{\ell^{1 / s_{k}\left(G_{k}\right)}}}=\left(\sum_{x \in E_{y}} \prod_{j \neq k} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}^{y}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}\right) \cdot \frac{f_{k}^{y}(y)}{f_{k}^{y}(y)} \\
& =\sum_{x \in E_{y}} \prod_{j \neq k} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}^{y}\right\|_{\ell^{1 / t_{j}\left(G_{j}\right)}}} \geq \sum_{x \in E_{y}} \prod_{j \neq k} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}} .
\end{aligned}
$$

We will substitute this inequality in the definition of $C(\mathcal{G}, t, f)$,

$$
\begin{aligned}
C(\mathcal{G}, t, f) & =\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}}=\sum_{x \in E} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}}=\sum_{y \in \phi_{k}(E)} \sum_{x \in E_{y}} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}} \\
& =\sum_{y \in \phi_{k}(E)} \sum_{x \in E_{y}} \prod_{j} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}}=\sum_{y \in \phi_{k}(E)}\left(\sum_{x \in E_{y}} \prod_{j \neq k} \frac{f_{j}^{y}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / t_{j}\left(G_{j}\right)}}}\right) \cdot \frac{f_{k}^{y}(y)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \\
& \leq \sum_{y \in \phi_{k}(E)} C\left(\mathcal{G}, s, f^{y}\right) \cdot \frac{f_{k}^{y}(y)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \leq\left(\max _{y \in \phi_{k}(E)} C\left(\mathcal{G}, s, f^{y}\right)\right) \cdot \sum_{y \in \phi_{k}(E)} \frac{f_{k}^{y}(y)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \\
& =\left(\max _{y \in \phi_{k}(E)} C\left(\mathcal{G}, s, f^{y}\right)\right) \cdot \frac{\left\|f_{k}^{y}\right\|_{\ell^{1}\left(G_{k}\right)}}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \leq\left(\max _{y \in \phi_{k}(E)} C\left(\mathcal{G}, s, f^{y}\right)\right) \cdot 1 \leq C(\mathcal{G}, s) .
\end{aligned}
$$

Taking a supremum over all $E \in \mathcal{E}^{*}(G)$ and all $f \in \mathcal{F}^{*}(\mathcal{G}, t)$ such that each $\operatorname{supp}\left(f_{j}\right)=\phi_{j}(E)$ and then applying Lemma 4.38, we have that $C(\mathcal{G}, t) \leq C(\mathcal{G}, s)$, thus by Lemma 4.32, $C(\mathcal{G}, t)=C(\mathcal{G}, s)$. This argument can be repeated for each index $k$ such that $s_{k}>1$ to obtain $t=\left(\min \left(1, s_{j}\right)\right)_{j}$.

We obtain the first conclusion $B(\mathcal{G}, s)=B\left(\mathcal{G},\left(\min \left(1, s_{j}\right)\right)_{j}\right)$ as a special case of the preceding argument. In more detail, for each $E \in \mathcal{E}^{*}(G)$, we consider just the $m$-function
$f$ with components $f_{j}=1_{\phi_{j}(E)}$, so $C(\mathcal{G}, t, f)=B(\mathcal{G}, t, E)$; it follows that its associated $m$-functions $f^{y}$ have components $f_{j}^{y}=1_{\phi_{j}\left(E_{y}\right)}$, so $C\left(\mathcal{G}, s, f^{y}\right)=B\left(\mathcal{G}, s, E_{y}\right)$, therefore $\max _{y \in \phi_{k}(E)} C\left(\mathcal{G}, s, f^{y}\right) \leq B(\mathcal{G}, s)$. We then again take a supremum over all $E \in \mathcal{E}^{*}(G)$, obtaining $B(\mathcal{G}, t, E) \leq B(\mathcal{G}, s)$ for each, so $B(\mathcal{G}, t) \leq B(\mathcal{G}, s)$, and thus by Lemma 4.32, $B(\mathcal{G}, t)=B(\mathcal{G}, s)$.

In light of this result, it suffices to study $B(\mathcal{G}, s)$ and $C(\mathcal{G}, s)$ for just $s \in[0,1]^{m}$, rather than all $s \in[0, \infty)^{m}$ : in particular, their values on the $m$ faces of $[0,1]^{m}$ that intersect $(1, \ldots, 1)$ determine their values for all $s \in[0, \infty)^{m} \backslash[0,1]^{m}$.

Lemma 4.40. If $S$ is a nonempty closed subset of $[0,1]^{m}, c \in[1, \infty)$, and $C(\mathcal{G}, s) \leq c$ for all $s \in S$, then $C(\mathcal{G}, s) \leq c$ for all $s$ in the (closed) convex hull of $S$.

Proof. We will use some additional measure-theoretic notation in the following proof, in order to apply a multilinear variant of the Riesz-Thorin interpolation theorem (see, e.g., 8 , Theorem 2.7]).

Given a measure space $(X, \Sigma, \mu)$, a function $f: X \rightarrow \mathbb{C}$ is $\Sigma$-simple if $f=\sum_{i=1}^{n} c_{i} 1_{E_{i}}$ for some $n \in\{1,2, \ldots\}, c_{1}, \ldots, c_{n} \in \mathbb{C}$, and $E_{1}, \ldots, E_{n} \in \Sigma$.

Let $\left(X_{j}\right)_{j}$ be an $m$-tuple of sets; a function $T: X_{j}\left(X_{j} \rightarrow \mathbb{C}\right) \rightarrow \mathbb{C}$ is a multilinear form if, for any $f \in X_{j}\left(X_{j} \rightarrow \mathbb{C}\right)$ with component $f_{k}=c g+d h$ for $c, d \in \mathbb{C}$ and $g, h: X_{k} \rightarrow \mathbb{C}$,

$$
T(f)=c T\left(f_{1}, \ldots, f_{k-1}, g, f_{k+1}, \ldots, f_{m}\right)+d T\left(f_{1}, \ldots, f_{k-1}, h, f_{k+1}, \ldots, f_{m}\right)
$$

( $T$ does not denote a torsion subgroup, as it does elsewhere.)
Now let $\left(\left(X_{j}, \Sigma_{j}, \mu_{j}\right)\right)_{j}$ be an $m$-tuple of measure spaces and let $S_{j}$ denote the set of $\Sigma_{j^{-}}$ simple functions $f_{j}: X_{j} \rightarrow \mathbb{C}$. With this notation, a variant of the Riesz-Thorin interpolation theorem states the following: given a multilinear form $T$ and $r, t \in[0,1]^{m}$, if there exist $\rho, \tau \in[0, \infty)$ such that

$$
|T(f)| \leq \rho \prod_{j}\|f\|_{L^{1 / r_{j}\left(X_{j}, \Sigma_{j}, \mu_{j}\right)}} \quad \text { and } \quad|T(f)| \leq \tau \prod_{j}\left\|f_{j}\right\|_{L^{1 / t_{j}\left(X_{j}, \Sigma_{j}, \mu_{j}\right)}} \quad \text { for all } f \in \underset{j}{ } S_{j}
$$

then, for any $\theta \in(0,1)$, letting $\sigma=\rho^{\theta} \tau^{1-\theta}$ and $s=\theta r+(1-\theta) t$,

$$
|T(g)| \leq \sigma \prod_{j}\left\|g_{j}\right\|_{L^{1 / s_{j}}\left(X_{j}, \Sigma_{j}, \mu_{j}\right)} \quad \text { for all } f \in \underset{j}{X} S_{j}
$$

To apply this result in the present setting of the discrete Abelian groups $G, G_{j}$, we will consider the multilinear form $T(f)=\sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right)$, where $f \in X_{j}\left(G_{j} \rightarrow \mathbb{C}\right)$. Here we abuse our notation to suppose, for all $s \in[0, \infty)^{m}, \mathcal{F}^{*}(\mathcal{G}, s) \subseteq X_{j}\left(G_{j} \rightarrow \mathbb{C}\right)$ : this embedding is possible because for each $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, each of its components $f_{j}<\infty$, so we can restrict the codomain each $f_{j}$ from $[0, \infty]$ to $[0, \infty)$ and then embed $[0, \infty)$ in $\mathbb{C}$.

We first show that $|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$ for all $f \in X_{j} S_{j}$ and all $s \in S$. Consider any $s \in S$ and any $f \in X_{j} S_{j}$. Suppose first that $\left(\left|f_{j}\right|\right)_{j} \notin \mathcal{F}^{*}(\mathcal{G}, s)$ : this leads to the following two cases. In the first case, some $f_{j}=0$, in which case $|T(f)|=0=c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$. In the second case, some $\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}=\infty$; assuming the first case does not apply, we have
that $|T(f)| \leq \infty=c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$. Otherwise, suppose that $\left(\left|f_{j}\right|\right)_{j} \in \mathcal{F}^{*}(\mathcal{G}, s) \cap X_{j} S_{j}$ and apply the hypothesis $C\left(\mathcal{G}, s,\left(\left|f_{j}\right|\right)_{j}\right) \leq C(\mathcal{G}, s) \leq c$ :

$$
\begin{aligned}
|T(f)| & =\left|\sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right)\right| \leq \sum_{x \in G} \prod_{j}\left|f_{j}\left(\phi_{j}(x)\right)\right|=C\left(\mathcal{G}, s,\left(\left|f_{j}\right|\right)_{j}\right) \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \\
& \leq C(\mathcal{G}, s) \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)} .
\end{aligned}
$$

Thus, we confirm that $|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$ for all $f \in X_{j} S_{j}$ and all $s \in S$.
Now let $\mathcal{P}$ denote the convex hull of $S$, the smallest convex set in $\mathbb{R}^{m}$ containing $S$ : since $S \subseteq[0,1]^{m}$ is closed and bounded by hypothesis, it folows that $\mathcal{P} \subseteq[0,1]^{m}$ is also closed and bounded, and moreover each of $\mathcal{P}$ 's extreme points is an element of $S$. By Carathéodory's theorem, any point in $\mathcal{P}$ is a convex combination of at most $m+1$ extreme points. Since the version of the Riesz-Thorin theorem above is stated in terms of pairwise convex combinations rather than general convex combinations, we will apply it iteratively, in a sequence of at most $m+1$ rounds, to obtain the conclusion for all $s \in \mathcal{P}$.

The zeroth round was already established: $|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$ for all $f \in X_{j} S_{j}$ at each point $s$ in $\mathcal{P}_{0}=S$. For each round $k \in\{1,2, \ldots\}$, assume the induction hypothesis, that the set $\mathcal{P}_{k-1}$ obtained in the $(k-1)$-th round includes all $k$-element convex combinations of $S$, and that $|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$ for all $f \in X_{j} S_{j}$ and all $s \in \mathcal{P}_{k-1}$. For the $k$-th round, we consider all pairwise convex combinations of $r \in \mathcal{P}_{k-1}$ and $t \in \mathcal{P}_{0}$. By the induction hypothesis (and base case), for all $f \in X_{j} S_{j},|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{1 / r_{j}}$ for all $r \in \mathcal{P}_{k-1}$ and $|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / t_{j}}\left(G_{j}\right)}$ for all $t \in \mathcal{P}_{0}$. So, applying the Riesz-Thorin interpolation theorem with $T, r, t$, and $\rho=\tau=c$, we have that $|T(f)| \leq \sigma \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$ with $\sigma=c$ for all $f \in X_{j} S_{j}$ at all points $s$ on the line segment between $r$ and $t$ (i.e., $s=\theta r+(1-\theta) t$ for $\theta \in[0,1]$ ). This can be repeated for all $s$ equal to a pairwise convex combination of $r \in \mathcal{P}_{k-1}$ and $t \in \mathcal{P}_{0}$; the set $\mathcal{P}_{k}$ of all such $s$ therefore includes all $(k+1)$-element convex combinations of $S$. After $m+1$ rounds, we have established $|T(f)| \leq c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}\left(G_{j}\right)}}$ for all $f \in X_{j} S_{j}$ and all $s \in \mathcal{P}_{m}=\mathcal{P}$.

Now, for any $f \in X_{j} S_{j}$,

$$
|T(f)|=\left|\sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right)\right| \leq \sum_{x \in G} \prod_{j}\left|f_{j}\left(\phi_{j}(x)\right)\right|=T\left(\left(\left|f_{j}\right|\right)_{j}\right),
$$

which holds with equality when each $f_{j} \geq 0$. Therefore, for any $s \in \mathcal{P}$ and any $f \in$ $\mathcal{F}^{*}(\mathcal{G}, s) \cap X_{j} S_{j}$,

$$
\begin{aligned}
C(\mathcal{G}, s, f) & =\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\frac{\sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right)}{\prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \\
& =\frac{T(f)}{\prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \leq \frac{c \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}{\prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=c .
\end{aligned}
$$

Since the simple functions in the discrete setting include all functions with finite supports, applying Lemma 4.38 for each $s \in \mathcal{P}$, we have that $C(\mathcal{G}, s, f) \leq c$ for all $f \in \mathcal{F}^{*}(\mathcal{G}, s)$; finally, taking a supremum over $\mathcal{F}^{*}(\mathcal{G}, s)$, we conclude that $C(\mathcal{G}, s) \leq c$ for all $s \in \mathcal{P}$.

### 4.5.3 Factorization of HBL Data

We recall the notation introduced in Section 4.4.5regarding factor data: given an HBL datum $\mathcal{G}=\left(G,\left(G_{j}\right),\left(\phi_{j}\right)\right)$ and any $H \leq G$, its restricted datum $\left.\mathcal{G}\right|_{H}=\left(H,\left(\phi_{j}(H)\right)_{j},\left(\chi_{j}\right)_{j}\right)$ has homomorphisms $\chi_{j}: H \rightarrow \phi_{j}(H): x \mapsto \phi_{j}(x)$ and its quotient datum $\left.\mathcal{G}\right|^{H}=\left(G / H,\left(G_{j} / \phi_{j}(H)\right)_{j},\left(\psi_{j}\right)_{j}\right)$ has homomorphisms $\psi_{j}: G / H \rightarrow G_{j} / \phi_{j}(H): x+H \mapsto \phi_{j}(x)+\phi_{j}(H)$.

Lemma 4.41. If $s \in[0, \infty)^{m}$ and $H \leq G$, then $A\left(\left.\mathcal{G}\right|_{H}, s\right) \leq A(\mathcal{G}, s)$, with equality when $\mathcal{H}_{\mathrm{f}}^{*}(H)=\mathcal{H}_{\mathrm{f}}^{*}(G)$.

Proof. Generally $\mathcal{H}_{\mathrm{f}}^{*}(H) \subseteq \mathcal{H}_{\mathrm{f}}^{*}(G)$, and for all $K \in \mathcal{H}_{\mathrm{f}}^{*}(H),\left|\chi_{j}(K)\right|=\left|\phi_{j}(K)\right|$.
Lemma 4.42. If $H \in \mathcal{H}_{\mathrm{f}}^{*}(G), s \in[0, \infty)^{m}$, and $A\left(\left.\mathcal{G}\right|_{H}, s, H\right)=A\left(\left.\mathcal{G}\right|_{H}, s\right)$, then $A\left(\left.\mathcal{G}\right|_{H}, s\right) A\left(\left.\mathcal{G}\right|^{H}, s\right) \leq$ $A(\mathcal{G}, s)$.

Proof. Every element of $\mathcal{H}_{\mathrm{f}}^{*}(G / H)$ can be written as $K / H$ for some $H \leq K \in \mathcal{H}_{\mathrm{f}}^{*}(G)$ (i.e., some $K \in \mathcal{H}_{\mathrm{f}}^{*}(G)$ such that $\left.H \leq K\right)$; likewise, for each $j$, each element of $\mathcal{H}_{\mathrm{f}}^{*}\left(\phi_{j}(G) / \phi_{j}(H)\right)$ can be written as $\phi_{j}(K) / \phi_{j}(H)=\psi_{j}(K / H)$ for some $H \leq K \in \mathcal{H}_{\mathrm{f}}^{*}(G)$. By Lagrange's theorem, $|K / H|=|K| /|H|$, and, for each $j,\left|\phi_{j}(K) / \phi_{j}(H)\right|=\left|\phi_{j}(K)\right| /\left|\phi_{j}(H)\right|$. So, for any $H \leq K \in \mathcal{H}_{\mathrm{f}}^{*}(G)$,

$$
\begin{aligned}
A(\mathcal{G}, s) & \geq A(\mathcal{G}, s, K)=\frac{|K|}{\prod_{j}\left|\phi_{j}(K)\right|^{s_{j}}}=\frac{|H| \cdot|K / H|}{\prod_{j}\left(\left|\phi_{j}(H)\right| \cdot\left|\phi_{j}(K) / \phi_{j}(H)\right|\right)^{s_{j}}} \\
& =\left(\frac{|H|}{\prod_{j}\left|\chi_{j}(H)\right|^{s_{j}}}\right)\left(\frac{|K / H|}{\prod_{j}\left|\psi_{j}(K / H)\right|^{s_{j}}}\right)=A\left(\left.\mathcal{G}\right|_{H}, s, H\right) A\left(\left.\mathcal{G}\right|^{H}, s, K / H\right) \\
& =A\left(\left.\mathcal{G}\right|_{H}, s\right) A\left(\left.\mathcal{G}\right|^{H}, s, K / H\right)
\end{aligned}
$$

Repeating for each $H \leq K \in \mathcal{H}_{\mathrm{f}}^{*}(G)$, the value $A\left(\left.\mathcal{G}\right|^{H}, s, K / H\right)$ for each $K / H \in \mathcal{H}_{\mathrm{f}}^{*}(G)$ appears on the right-hand side at least once, so taking a supremum over $H \leq K \in \mathcal{H}_{\mathrm{f}}^{*}(G)$, the preceding inequality becomes $A(\mathcal{G}, s) \geq A\left(\left.\mathcal{G}\right|_{H}, s\right) A\left(\left.\mathcal{G}\right|^{H}, s\right)$.
Lemma 4.43. If $H \leq G$ and $s \in[0, \infty)^{m}$, then $C(\mathcal{G}, s) \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s\right)$.
Proof. First we suppose, without loss of generality, that $s \in(0, \infty)^{m}$, by application of Lemma 4.37 .

Consider any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$. First consider the quotient datum $\left.\mathcal{G}\right|^{H}$. Define the $m$-function $h$ componentwise by

$$
h_{j}: G_{j} / \phi_{j}(H) \rightarrow[0, \infty): y+\phi_{j}(H) \mapsto\left(\sum_{z \in \phi_{j}(H)} f_{j}(z+y)^{1 / s_{j}}\right)^{s_{j}}
$$

we now check that each component $h_{j}$ is a function of the coset $y+\phi_{j}(H)$, meaning that it is independent of the choice of the representative $y \in G_{j}$ : considering any other possible representative $w \in G_{j}$, we have that $v=w-y \in \phi_{j}(H)$, so

$$
\begin{aligned}
h_{j}\left(w+\phi_{j}(H)\right)^{1 / s_{j}} & =\sum_{u \in \phi_{j}(H)} f_{j}(u+w)^{1 / s_{j}}=\sum_{u \in \phi_{j}(H)} f_{j}(u+v+y)^{1 / s_{j}} \\
& =\sum_{z-v \in \phi_{j}(H)} f_{j}(z+y)^{1 / s_{j}}=\sum_{z \in \phi_{j}(H)} f_{j}(z+y)^{1 / s_{j}}=h_{j}\left(y+\phi_{j}(H)\right)^{1 / s_{j}}
\end{aligned}
$$

since $z-v \in \phi_{j}(H)$ if and only if $z \in \phi_{j}(H)$. Moreover,

$$
\begin{aligned}
\left\|h_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j} / \phi_{j}(H)\right)}^{1 / s_{j}} & =\sum_{y+\phi_{j}(H) \in G_{j} / \phi_{j}(H)}\left(\left(\sum_{z \in \phi_{j}(H)} f_{j}(z+y)^{1 / s_{j}}\right)^{s_{j}}\right)^{1 / s_{j}} \\
& \left.=\sum_{y+\phi_{j}(H) \in G_{j} / \phi_{j}(H)} \sum_{z \in \phi_{j}(H)} f_{j}(z+y)^{1 / s_{j}}=\sum_{y \in G_{j}} f_{j}(y)^{1 / s_{j}}=\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}^{1 / s_{j}}\right)
\end{aligned}
$$

because, letting $Y$ denote the set of coset representatives $y$ appearing in the summation subscript $y+\phi_{j}(H) \in G_{j} / \phi_{j}(H)$, the map $Y \times \phi_{j}(H) \rightarrow G_{j}:(y, z) \mapsto y+z$ is a bijection. It follows that $h \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|^{H}, s\right)$, thus by definition,

$$
C\left(\left.\mathcal{G}\right|^{H}, s, h\right)=\sum_{x+H \in G / H} \prod_{j} \frac{h_{j}\left(\psi_{j}(x+H)\right)}{\left\|h_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j} / \phi_{j}(H)\right)}}
$$

Next we consider the restricted datum $\left.\mathcal{G}\right|_{H}$. For each $w \in G$, define the $m$-function $g^{(w)}$ with $g_{j}^{(w)}: \phi_{j}(H) \rightarrow[0, \infty]: y \mapsto f_{j}\left(y+\phi_{j}(w)\right)$. For each $j,\left\|g_{j}^{(w)}\right\|_{\ell^{1 / s_{j}}\left(\phi_{j}(H)\right)} \leq\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$, so there are two mutually exclusive cases. Either $g_{j}^{(w)}=0$ for some $j$, or $g^{(w)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{H}, s\right)$. In the former case, $\sum_{x \in H} \prod_{j} g^{(w)}\left(\phi_{j}(x)\right)=0$; in the latter case,

$$
\begin{aligned}
C\left(\left.\mathcal{G}\right|_{H}, s, g^{(w)}\right) & =\sum_{x \in H} \prod_{j} \frac{g_{j}^{(w)}\left(\chi_{j}(x)\right)}{\left\|g_{j}^{(w)}\right\|_{\ell^{1 / s_{j}}\left(\phi_{j}(H)\right.}}=\sum_{x \in H} \prod_{j} \frac{\left.f_{j}\left(\phi_{j}(x)+\phi_{j}(w)\right)\right)}{\left(\sum_{y \in \phi_{j}(H)} f_{j}\left(y+\phi_{j}(w)\right)^{1 / s_{j}}\right)^{s_{j}}} \\
& =\sum_{x \in H} \prod_{j} \frac{f_{j}\left(\phi_{j}(x+w)\right)}{h_{j}\left(\psi_{j}(w+H)\right)} .
\end{aligned}
$$

Now allowing $v$ to vary over a set of representatives of the cosets $v+H \in G / H$,

$$
\begin{aligned}
\sum_{x \in G} \prod_{j} f_{j}\left(\phi_{j}(x)\right) & =\sum_{\substack{v+H \in G / H}} \sum_{x \in H} \prod_{j} f_{j}\left(\phi_{j}(x+v)\right) \\
& =\sum_{\substack{v+H \in G / H \\
g^{(v)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{H}, s\right)}} C\left(\left.\mathcal{G}\right|_{H}, s, g^{(v)}\right) \cdot \prod_{j} h_{j}\left(\psi_{j}(v+H)\right) \\
& \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) \sum_{\substack{v+H \in G / H \\
g^{(v)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{H}, s\right)}} \prod_{j} h_{j}\left(\psi_{j}(v+H)\right) \\
& \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) \sum_{x+H \in G / H} \prod_{j} h_{j}\left(\psi_{j}(x+H)\right) \\
& =C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s, h\right) \prod_{j}\left\|h_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j} / \phi_{j}(H)\right)} \\
& \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s\right) \prod_{j}\left\|h_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j} / \phi_{j}(H)\right)} \\
& =C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s\right) \prod_{j}\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}
\end{aligned}
$$

so $C(\mathcal{G}, s, f) \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s\right)$. The conclusion is obtained by repeating this argument for each $f \in \mathcal{F}^{*}(\mathcal{G}, s)$ and observing the same conclusion each time.

### 4.5.4 When $G$ Is Finite

The following three results, in tandem with Lemma 4.42 above, complete our summary of the machinery introduced in [14] to precisely determine $B(\mathcal{G}, s), C(\mathcal{G}, s)$ when $G$ is a finite Abelian group (i.e., a torsion group). In Section 4.6, we synthesize these results with the results from [15] for torsion-free Abelian groups to address the general case. Here we follow [14] and restrict $s \in[0,1]^{m}$; later, in the proof of Theorem 4.2, we apply Lemmas 4.27 and Lemma 4.39 to address the case of $s \in[0, \infty)^{m}$.

Lemma 4.44. If $|G|<\infty$ and $s \in[0,1]^{m}$ with all $s_{j}=0$ except one $s_{k} \geq 0$, then $A(\mathcal{G}, s)=$ $A(\mathcal{G}, s, G)=|G|^{1-s_{k}}\left|\operatorname{ker}\left(\phi_{k}\right)\right|^{s_{k}}=C(\mathcal{G}, s)$.

Proof. Let $K=\operatorname{ker}\left(\phi_{k}\right)$ and consider any $H \leq G\left(\right.$ noting that $\left.H, K, G \in \mathcal{H}_{\mathrm{f}}^{*}(G)\right)$ :

$$
\begin{aligned}
A(\mathcal{G}, s, H) & =\frac{|H|}{\prod_{j}\left|\phi_{j}(H)\right|^{s_{j}}}=\frac{|H|}{\left|\phi_{k}(H)\right|^{s_{k}}}=\frac{|H|}{\left(\frac{|H|}{|H \cap K|}\right)^{s_{k}}} \\
& =|H|^{1-s_{k}}|H \cap K|^{s_{k}} \leq|G|^{1-s_{k}}|H \cap K|^{s_{k}} \leq|G|^{1-s_{k}}|K|^{s_{k}}=A(\mathcal{G}, s, G)
\end{aligned}
$$

so $A(\mathcal{G}, s)=A(\mathcal{G}, s, G)=|G|^{1-s_{k}}|K|^{s_{k}}$, the two equalities in the conclusion.
Now consider any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$ :

$$
C(\mathcal{G}, s, f)=\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sum_{x \in G}\left(\prod_{j \neq k} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{\infty}\left(G_{j}\right)}}\right) \frac{f_{k}\left(\phi_{j}(x)\right)}{\left\|f_{k}\right\|_{\ell^{1 / s_{k}\left(G_{k}\right)}}} \leq \sum_{x \in G} \frac{f_{k}\left(\phi_{j}(x)\right)}{\left\|f_{k}\right\|_{\ell^{1 / s_{k}\left(G_{k}\right)}}}
$$

and, letting $V$ denote a set of representatives of the cosets $v+K \in G / K$,

$$
\begin{aligned}
\sum_{x \in G} f_{k}\left(\phi_{k}(x)\right) & =\sum_{v \in V} \sum_{u \in K} f_{k}\left(\phi_{k}(u+v)\right)=\sum_{v \in V}|K| f_{k}\left(\phi_{k}(v)\right)=|K|\left\|\left.f_{k}\right|_{\phi_{k}(V)}\right\|_{\ell^{1}\left(\phi_{k}(V)\right)} \\
& \leq|K|\left|\phi_{k}(V)\right|^{1-s_{k}}\left\|\left.f_{k}\right|_{\phi_{k}(V)}\right\|_{\ell^{1 / s_{k}}\left(\phi_{k}(V)\right)} \leq|K|\left|\phi_{k}(V)\right|^{1-s_{k}}\left\|f_{k}\right\|_{\ell^{1 / s_{k}}\left(G_{k}\right)} \\
& =|K|\left|\phi_{k}(G)\right|^{1-s_{k}}\left\|f_{k}\right\|_{\ell^{1 / s_{k}}\left(G_{k}\right)}=|K|\left(\frac{|G|}{|K|}\right)^{1-s_{k}}\left\|f_{k}\right\|_{\ell^{1 / s_{k}\left(G_{k}\right)}} \\
& =|K|^{s_{k}}|G|^{1-s_{k}}\left\|f_{k}\right\|_{\ell^{1 / s_{k}\left(G_{k}\right)}}=A(\mathcal{G}, s)\left\|f_{k}\right\|_{\ell^{1 / s_{k}\left(G_{k}\right)}},
\end{aligned}
$$

so $C(\mathcal{G}, s, f) \leq A(\mathcal{G}, s)$ for all $f \in \mathcal{F}^{*}(\mathcal{G}, s)$, thus $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$. By Lemma 4.31 $A(\mathcal{G}, s) \leq C(\mathcal{G}, s)$, so we conclude $A(\mathcal{G}, s)=C(\mathcal{G}, s)$.

Lemma 4.45. If $|G|<\infty$ and $s \in[0,1]^{m}$ with all $s_{j} \in\{0,1\}$ except one $s_{k} \in[0,1]$, then $A(\mathcal{G}, s)=C(\mathcal{G}, s)$.

Proof. Let $I=\left\{j \neq k \mid s_{j}=1\right\}$. Let $H=\bigcap_{i \in I} \operatorname{ker}\left(\phi_{i}\right)$ and consider the factor data $\left.\mathcal{G}\right|_{H}$ and $\left.\mathcal{G}\right|^{H}$.

First we study the restricted datum $\left.\mathcal{G}\right|_{H}$ and show that $A\left(\left.\mathcal{G}\right|_{H}, s\right)=C\left(\left.\mathcal{G}\right|_{H}, s\right)$. By Lemma 4.37. $C\left(\left.\mathcal{G}\right|^{H}, s\right) \leq C\left(\left(\psi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}\right)$. Now apply Lemma 4.44 to $\left(\chi_{j}\right)_{j \notin I}$ and $\left(s_{j}\right)_{j \notin I}$ :

$$
A\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}, H\right)=A\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}\right)=C\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}\right)=|H|^{1-s_{k}}\left|\operatorname{ker}\left(\phi_{k}\right)\right|^{s_{k}} .
$$

Note that $\prod_{i \in I}\left|\phi_{i}(H)\right|^{s_{i}}=1$, so $A\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}, H\right)=A\left(\left.\mathcal{G}\right|_{H}, s, H\right)$ and $A\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}\right)=$ $A\left(\left.\mathcal{G}\right|_{H}, s\right)$. Therefore, applying Lemma 4.37,

$$
A\left(\left.\mathcal{G}\right|_{H}, s\right)=A\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}\right)=C\left(\left(\chi_{j}\right)_{j \notin I},\left(s_{j}\right)_{j \notin I}\right) \geq C\left(\left.\mathcal{G}\right|_{H}, s\right)
$$

by Lemma 4.31, $A\left(\left.\mathcal{G}\right|_{H}, s\right) \leq C\left(\left.\mathcal{G}\right|_{H}, s\right)$, so we must have that $A\left(\left.\mathcal{G}\right|_{H}, s\right)=C\left(\left.\mathcal{G}\right|_{H}, s\right)$.
Next we study the quotient datum $\left.\mathcal{G}\right|_{H}$ and show that $A\left(\left.\mathcal{G}\right|^{H}, s\right)=C\left(\left.\mathcal{G}\right|^{H}, s\right)$. Since $H=\bigcap_{i \in I} \operatorname{ker}\left(\phi_{i}\right)$, it follows that $\bigcap_{i \in I} \operatorname{ker}\left(\psi_{i}\right)=\{0\}$. By Lemma 4.36, $C\left(\left(\psi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I}\right) \leq 1$. By Lemma 4.37, $C\left(\left.\mathcal{G}\right|^{H}, s\right) \leq C\left(\left(\psi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I}\right)$. By Lemma 4.31, $1 \leq A\left(\left.\mathcal{G}\right|^{H}, s\right) \leq C\left(\left.\mathcal{G}\right|^{H}, s\right)$, so $A\left(\left.\mathcal{G}\right|^{H}, s\right)=C\left(\left.\mathcal{G}\right|^{H}, s\right)$.

Since we concluded above that $A\left(\left.\mathcal{G}\right|_{H}, s, H\right)=A\left(\left.\mathcal{G}\right|_{H}, s\right)$, by Lemma 4.42, $A(\mathcal{G}, s) \geq$ $A\left(\left.\mathcal{G}\right|_{H}, s\right) A\left(\left.\mathcal{G}\right|^{H}, s\right)$. By Lemma 4.43 .

$$
C(\mathcal{G}, s) \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s\right)=A\left(\left.\mathcal{G}\right|_{H}, s\right) A\left(\left.\mathcal{G}\right|^{H}, s\right) \leq A(\mathcal{G}, s)
$$

By Lemma 4.31, $A(\mathcal{G}, s) \leq C(\mathcal{G}, s)$, so it must be that $A(\mathcal{G}, s)=C(\mathcal{G}, s)$.
Lemma 4.46. If $|G|<\infty$ and $s \in[0,1]^{m}$, then $A(\mathcal{G}, s)=C(\mathcal{G}, s)$.
Proof. By Lemma 4.31, we have that $A(\mathcal{G}, s) \leq C(\mathcal{G}, s)$, so it remains to show the reverse inequality $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$. We do so by induction on $|G|, m \in\{1,2, \ldots\}$. In the base cases with $|G|=1$, by Lemma 4.31, we have that $C(\mathcal{G}, s) \leq|G|=1 \leq A(\mathcal{G}, s)$. In the base cases with $m=1$, by Lemma 4.44, we have that $A(\mathcal{G}, s)=C(\mathcal{G}, s)$.

Now we suppose $|G|>1$ and $m>1$, and assume the induction hypotheses for all HBL data $\mathcal{G}^{\prime}$ with $\left|G^{\prime}\right|<|G|$ and $m^{\prime}<m$, for all $s \in[0,1]^{m^{\prime}}$. We will invoke the induction hypotheses in three scenarios; after concluding the proof in these cases, we will show how the general case reduces to one of these three.

In the first scenario, some component $s_{k}=0$ : by the induction hypotheses and Lemma 4.37, $A(\mathcal{G}, s)=A\left(\left(\phi_{j}\right)_{j \neq k},\left(s_{j}\right)_{j \neq k}\right)=C\left(\left(\phi_{j}\right)_{j \neq k},\left(s_{j}\right)_{j \neq k}\right)=C(\mathcal{G}, s)$.

In the second scenario, some component $s_{k}=1$; let $K=\operatorname{ker}\left(\phi_{k}\right)$, let $\left.\mathcal{G}\right|_{K}=\left(K,\left(\phi_{j}(K)\right)_{j},\left(\chi_{j}\right)_{j}\right)$ be the restricted datum, let $\left.\mathcal{G}^{\prime}\right|_{K}=\left(K,\left(\phi_{j}(K)\right)_{j \neq k},\left(\chi_{j}\right)_{j \neq k}\right)$ be $\left.\mathcal{G}\right|_{K}$ with its $k$-th component omitted, and let $V$ be a set of representatives of the cosets $v+K \in G / K$. For each $w \in G$, define $g^{(w)}$ with components $g_{j}^{(w)}: \phi_{j}(K) \rightarrow[0, \infty]: y \mapsto f_{j}\left(y+\phi_{j}(w)\right)$. For each $j$, $\left\|g_{j}^{(w)}\right\|_{\ell^{1 / s_{j}}\left(\phi_{j}(K)\right)} \leq\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}$, so there are two mutually exclusive cases. Either $g_{j}^{(w)}=0$ for some $j$, or $g^{(w)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{K}, s\right)$. In the former case, $\sum_{x \in K} \prod_{j} g_{j}^{(w)}\left(\phi_{j}(x)\right)=0$. In the latter case, $\left(g_{j}^{(w)}\right)_{j \neq k} \in \mathcal{F}^{*}\left(\left.\mathcal{G}^{\prime}\right|_{K},\left(s_{j}\right)_{j \neq k}\right)$. For any $f \in \mathcal{F}^{*}(\mathcal{G}, s)$,

$$
\begin{aligned}
C(\mathcal{G}, s, f) & =\sum_{x \in G} \prod_{j} \frac{f_{j}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sum_{v \in V} \sum_{x \in K} \prod_{j} \frac{f_{j}\left(\phi_{j}(v+x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}}=\sum_{\substack{v \in V \\
g^{(v)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{K}, s\right)}} \sum_{x \in K} \prod_{j} \frac{g_{j}^{(v)}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \\
& =\sum_{\substack{v \in V \\
g^{(v)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{K}, s\right)}} \frac{f_{k}\left(\phi_{k}(v)\right)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)} \sum_{x \in K} \prod_{j \neq k} \frac{g_{j}^{(v)}\left(\phi_{j}(x)\right)}{\left\|f_{j}\right\|_{\ell^{1 / s_{j}}\left(G_{j}\right)}} \leq \sum_{\substack{v \in V \\
g^{(v)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{K}, s\right)}} \frac{f_{k}\left(\phi_{k}(v)\right)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \sum_{x \in K} \prod_{j \neq k} \frac{g_{j}^{(v)}\left(\phi_{j}(x)\right)}{\left\|g_{j}^{(v)}\right\|_{\ell^{1 / s_{j}}\left(\phi_{j}(K)\right)}}} \begin{aligned}
& \leq \sum_{\substack{v \in V \\
g^{(v)} \in \mathcal{F}^{*}(\mathcal{G} \mid K, s)}} \frac{f_{k}\left(\phi_{k}(v)\right)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} C\left(\left.\mathcal{G}^{\prime}\right|_{K},\left(s_{j}\right)_{j \neq k}\right) .
\end{aligned}
\end{aligned}
$$

Now by the induction hypothesis, we have that $C\left(\left.\mathcal{G}^{\prime}\right|_{K},\left(s_{j}\right)_{j \neq k}\right)=A\left(\left.\mathcal{G}^{\prime}\right|_{K},\left(s_{j}\right)_{j \neq k}\right)$. Moreover, since $\phi_{k}(K)=\{0\}$, applying Lemma 4.37 with $m$-functions $f \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{K}, s\right)$ such that $f_{j}=1_{\phi_{j}(H)}$ for some $H \in \mathcal{H}_{\mathrm{f}}^{*}(K)$, thus taking a supremum of the conclusion over all such $m$-functions, we have that $A\left(\left.\mathcal{G}^{\prime}\right|_{K},\left(s_{j}\right)_{j \neq k}\right)=A\left(\left.\mathcal{G}\right|_{K}, s\right)$. Lastly, by Lemma 4.41, we have that $A\left(\left.\mathcal{G}\right|_{K}, s\right) \leq A(\mathcal{G}, s)$. Therefore,

$$
C(\mathcal{G}, s, f) \leq A(\mathcal{G}, s) \sum_{\substack{v \in V \\ g^{(v)} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{K}, s\right)}} \frac{f_{k}\left(\phi_{k}(v)\right)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \leq A(\mathcal{G}, s) \sum_{v \in V} \frac{f_{k}\left(\phi_{k}(v)\right)}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}} \leq A(\mathcal{G}, s) \frac{\left\|f_{k}\right\|_{1}}{\left\|f_{k}\right\|_{\ell^{1}\left(G_{k}\right)}}=A(\mathcal{G}, s)
$$

since $\left.\phi_{k}\right|_{V}$ is an injection. The conclusion $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$ follows by taking a supremum over all $f \in \mathcal{F}^{*}(\mathcal{G}, s)$.

In the third scenario, $A(\mathcal{G}, s, K) \geq 1$ for some $\{0\}<K<G$. This implies the existence of some $\{0\}<H<G$ such that $A(\mathcal{G}, s, H)=A(\mathcal{G}, s)$, as we now show. If $A(\mathcal{G}, s, K)=$ $A(\mathcal{G}, s)$, then we pick $H=K$. Otherwise, $1 \leq A(\mathcal{G}, s, K)<A(\mathcal{G}, s)$. Since $\{0\}<K<$ $G$ and $|G|<\infty, \mathcal{H}=\operatorname{argmax}_{0<H<G} A(\mathcal{G}, s, H)$ is nonempty and finite. Each $H \in \mathcal{H}$ maximizes $A\left(\left.\mathcal{G}\right|_{H}, s, F\right)=A(\mathcal{G}, s, F)$ over all $\{0\} \leq F \leq H$, including $F=\{0\}$, since $A(\mathcal{G}, s,\{0\})=1 \leq A(\mathcal{G}, s, K)$. Thus, for any $H \in \mathcal{H}, A(\mathcal{G}, s, H)=A(\mathcal{G}, s)$. Considering $H$ obtained in either case, since $A(\mathcal{G}, s, H)=A\left(\left.\mathcal{G}\right|_{H}, s, H\right) \leq A\left(\left.\mathcal{G}\right|_{H}, s\right) \leq A(\mathcal{G}, s)$, we have that $A\left(\left.\mathcal{G}\right|_{H}, s, H\right)=A\left(\left.\mathcal{G}\right|_{H}, s\right)$, so by Lemma 4.43, the induction hypotheses, and Lemma 4.42,

$$
C(\mathcal{G}, s) \leq C\left(\left.\mathcal{G}\right|_{H}, s\right) C\left(\left.\mathcal{G}\right|^{H}, s\right)=A\left(\left.\mathcal{G}\right|_{H}, s\right) A\left(\left.\mathcal{G}\right|^{H}, s\right) \leq A(\mathcal{G}, s)
$$

We now show that it suffices to consider just the preceding three scenarios. For each $a \in[1, \infty)$, define $P(a)=\left\{t \in[0,1]^{m} \mid A(\mathcal{G}, t) \leq a\right\}$. Since $1 \leq A(\mathcal{G}, s) \leq|G|$ by Lemma 4.31 and $|G|<\infty$ by hypothesis, $s \in P(A(\mathcal{G}, s))$, thus $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$ is implied by the more general statement, $C(\mathcal{G}, t) \leq A(\mathcal{G}, s)$ for all $t \in P(A(\mathcal{G}, s))$. More generally still, we will show that for all $a \in[1, \infty)$ and all $t \in P(a), C(\mathcal{G}, t) \leq a$. For each $a \in[1, \infty)$, note that $P(a)$ is a closed and bounded convex set, so by Carathéodory's theorem, it equals its extreme points' convex hull; then, by Lemma 4.40, if $C(\mathcal{G}, t) \leq a$ for all extreme points $t$ of $P(a)$ then it holds for all $t \in P(a)$.

Consider any $a \in[1, \infty)$; if $P(a) \neq \emptyset$, examine any of its extreme points $t$. Since $t \in P(a)$, $A(\mathcal{G}, t) \leq a$. If any of the preceding three scenarios apply, then we conclude $C(\mathcal{G}, t) \leq$ $A(\mathcal{G}, t) \leq a$. Thus, it remains to treat the scenario where $t \in(0,1)^{m}$ and $A(\mathcal{G}, t, H)<$ $1 \leq A(\mathcal{G}, t)$ for all $\{0\}<H<G$. In particular, $A(\mathcal{G}, t)=\max (A(\mathcal{G}, t,\{0\}), A(\mathcal{G}, t, G))=$ $\max (1, A(\mathcal{G}, t, G))$.

We first show that $A(\mathcal{G}, t)=1>A(\mathcal{G}, t, G)$. Suppose towards a contradiction that $A(\mathcal{G}, t)=A(\mathcal{G}, t, G)$. By the definition of $P(a)$, we see that each $s \in P(a)$ satisfies the following finite system of linear inequalities (taking logarithms of $A(\mathcal{G}, s) \leq a$ ):

$$
0 \leq s_{j} \leq 1 \quad \text { for all } j, \quad \ln |H|-\sum_{j} s_{j} \ln \left|\phi_{j}(H)\right| \leq \ln a \quad \text { for all }\{0\} \neq H \leq G
$$

Under the present hypotheses regarding $t$, since $A(\mathcal{G}, s)$ is a continuous function of $s \in$ $[0, \infty)^{m}$, for all $s \in P(a)$ in a neighborhood of $t$, all of these inequalities hold strictly except possibly the one induced by $H=G$ (just in case $A(\mathcal{G}, t)=a$ ). Therefore, $t$ is contained in
at most one of the hyperplanes that bound $P(a)$ : this contradicts extremality of $t$, because $P(a)$ is defined by at least $m \geq 2$ and at most finitely many linear inequalities.

Therefore, $A(\mathcal{G}, t, G)<A(\mathcal{G}, t)=\max (1, A(\mathcal{G}, t, G))$. That is, $A(\mathcal{G}, t, G)<1=A(\mathcal{G}, t)$. Note that for all $H \leq G, A(\mathcal{G}, \alpha t, H)$ is a continuous nonincreasing function of $\alpha \in[0,1]$, so $A(\mathcal{G}, \alpha t)$ is too. For all $\{0\}<H \leq G, A(\mathcal{G}, 0 \cdot t, H)>1>A(\mathcal{G}, 1 \cdot t, H)$. Thus, by continuity and (anti-) monotonicity, for all $\{0\}<H \leq G$, there exists $\alpha_{H} \in(0,1)$ such that $A\left(\mathcal{G}, \alpha_{H} t, H\right)=1$. Let $\mathcal{K}$ be the set of all subgroups $\{0\}<K \leq G$ which maximize $\alpha_{K}$ : since $G$ has finitely many subgroups, $\mathcal{K} \neq \emptyset$ (i.e., this maximum is well defined), and it follows that $A\left(\mathcal{G}, \alpha_{K} t\right)=1=A(\mathcal{G}, t)$. Moreover, we can find at least one $K \in \mathcal{K}$ such that $K<G$ : to show this, suppose towards a contradiction that none exist. Then, abbreviating $r=\alpha_{G} t$, we must have that $A(\mathcal{G}, r, G)=A(\mathcal{G}, r)=1 \leq a$, i.e., $r \in P(a)$. By Lemma 4.32, $P(a)$ also contains every $r \leq q \in[0,1]^{m}$. By construction, $r \leq t$, and moreover some component $r_{k}<t_{k}$ since $\alpha_{G} \in(0,1)$. Fixing $q_{j}=t_{j}$ for all $j \neq k$ and $q_{k}=\theta r_{k}+(1-\theta) 1$ for some $\theta \in[0,1]$, we observe that $q \in P(a)$ for any $\theta \in[0,1]$. Since $t \in(0,1)^{m}$ is an extreme point of $P(a), r_{k}<t_{k}<1$, so there exists an open line segment in $P(a)$ containing $t$, defined by $q$ with parameter $\theta \in(0,1)$, contradicting extremality of $t$. In light of this contradiction, we have shown that there exists some $\{0\}<K<G$ such that $A\left(\mathcal{G}, \alpha_{K} t, K\right)=A\left(\mathcal{G}, \alpha_{K} t\right)$. Thus, the third scenario applies (replacing $s$ by $\alpha_{K} t$ ), and we conclude

$$
C(\mathcal{G}, t) \leq C\left(\mathcal{G}, \alpha_{K} t\right) \leq A\left(\mathcal{G}, \alpha_{K} t\right)=A(\mathcal{G}, t)=1 \leq a
$$

Repeating this argument and obtaining $C(\mathcal{G}, t) \leq a$ for each extreme point $t$ of $P(a)$ for each $a \in[1, \infty)$, we confirm that $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$. As mentioned above, the conclusion $A(\mathcal{G}, s)=C(\mathcal{G}, s)$ follows from Lemma 4.31 .

### 4.6 Relating $\mathcal{P}$ with $A, B$, and $C$

Now we connect the preceding results in order to complete the proof of Theorem 4.2, Theorem 4.1 follows as a special case. We suppose again that $\mathcal{G}$ is any HBL datum.

Lemma 4.47. If $s \in \mathcal{P}(\mathcal{G}) \cap[0,1]^{m}$ then $A(\mathcal{G}, s)=C(\mathcal{G}, s)$.
Proof. We first show that $\operatorname{rank}\left(\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)\right)=0$. If $\operatorname{rank}\left(\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)\right)>0$, then by Lemma 4.11, $\mathcal{P}(\mathcal{G})=\emptyset$, a contradiction to the hypothesis $s \in \mathcal{P}(\mathcal{G})$.

We first consider the special case where $G$ is finitely generated and torsion-free ( $T(G)=$ $\{0\}$ ), and treat the general case afterward. Applying Lemmas 4.31 and 4.33 when $G$ is torsion-free, we see that $1 \leq A(\mathcal{G}, s) \leq|T(G)|=1$, so it suffices to show that $C(\mathcal{G}, s)=1$. We know that $C(\mathcal{G}, s) \geq 1$ by Lemma 4.31, so it remains to show that $C(\mathcal{G}, s) \leq 1$, which we will do by induction on $\operatorname{rank}(G)$.

In the base case $\operatorname{rank}(G)=0$, since $G$ is finitely generated, $G=\{0\}$, i.e., $|G|=1$; therefore, $1 \leq C(\mathcal{G}, s) \leq|G|$ by Lemma 4.31 .

Now suppose that $\operatorname{rank}(G)>0$ and assume the induction hypothesis, that the conclusion holds for all HBL data $\mathcal{G}^{\prime}$ with $\operatorname{rank}\left(G^{\prime}\right)<\operatorname{rank}(G)$. By Lemma 4.40, it suffices to show that $C(\mathcal{G}, s) \leq 1$ at the extreme points of $\mathcal{P}(\mathcal{G}) \cap[0,1]^{m}$. By Lemma 4.29, if $s$ is an extreme point of $\mathcal{P}(\mathcal{G}) \cap[0,1]^{m}$, there are two cases, not necessarily exclusive.

In the first case, $s \in\{0,1\}^{m}$. Let $I=\left\{j \mid s_{j}=1\right\}$; note that $I \neq \emptyset$, since otherwise $1 \leq \operatorname{rank}(G) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(G)\right)=0$, contradicting the hypothesis that $s \in \mathcal{P}(\mathcal{G})$. Letting $K=\bigcap_{i \in I} \operatorname{ker}\left(\phi_{i}\right) \in \mathcal{H}_{\mathrm{fg}}^{*}(G)=\mathcal{H}_{\mathbb{N}}^{*}(G)$, since $s \in \mathcal{P}(\mathcal{G})$,

$$
\operatorname{rank}(K) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(K)\right)=\sum_{i \in I} \operatorname{rank}\left(\phi_{i}(K)\right)=0,
$$

so $\operatorname{rank}(K)=0$, i.e., $K=\{0\}$. By Lemmas 4.37 and 4.36 ,

$$
C(\mathcal{G}, s)=C\left(\left(\phi_{i}\right)_{i \in I},\left(s_{i}\right)_{i \in I}\right) \leq|K|=1 .
$$

In the second case, there exists $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)=\mathcal{H}_{\mathrm{fg}}^{*}(G)$ such that $0<\operatorname{rank}(H)=$ $\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)<\operatorname{rank}(G)$. By Lemma 4.25. $s \in \mathcal{P}\left(\left.\mathcal{G}\right|_{H}\right) \cap \mathcal{P}\left(\left.\mathcal{G}\right|^{H}\right)$. Apply Lemma 4.1 to obtain a torsion-free $F \in \mathcal{H}_{\mathrm{fg}}^{*}(G)$ such that $G / F$ is finitely generated and torsionfree, $\operatorname{rank}(F)=\operatorname{rank}(H)$, and $\operatorname{rank}\left(\phi_{j}(F)\right)=\operatorname{rank}\left(\phi_{j}(H)\right)$ for each $j$. Since $s \in[0,1]^{m}$, $0<\operatorname{rank}(F), \operatorname{rank}(G / F)<\operatorname{rank}(G)$, and both $F$ and $G / F$ are finitely generated and torsionfree, applying Lemma 4.43 and the induction hypotheses for $\left.\mathcal{G}\right|_{F}$ and $\left.\mathcal{G}\right|^{F}$,

$$
C(\mathcal{G}, s) \leq C\left(\left.\mathcal{G}\right|_{F}, s\right) C\left(\left.\mathcal{G}\right|^{F}, s\right) \leq 1 \cdot 1=1
$$

This concludes the proof in the case where $G$ is finitely generated and torsion free.
Now consider the more general case where $G$ is still finitely generated but not necessarily torsion free; our task is to show that $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$. Since $G$ is finitely generated, $H \in \mathcal{H}_{\mathrm{f}}^{*}(G)$ if and only if $H \leq T(G)$. Therefore, $A\left(\left.\mathcal{G}\right|_{T(G)}, s\right)=A(\mathcal{G}, s)$ by Lemma 4.41, and then applying Lemma 4.46, $A\left(\left.\mathcal{G}\right|_{T(G)}, s\right)=C\left(\left.\mathcal{G}\right|_{T(G)}, s\right)$. Since $G / T(G)$ is finitely generated and torsion-free, by the argument in the previous case, $A\left(\left.\mathcal{G}\right|^{T(G)}, s\right)=C\left(\left.\mathcal{G}\right|^{T(G)}, s\right)=1$, and so applying Lemma 4.43,

$$
C(\mathcal{G}, s) \leq C\left(\left.\mathcal{G}\right|_{T(G)}, s\right) C\left(\left.\mathcal{G}\right|^{T(G)}, s\right)=A\left(\left.\mathcal{G}\right|_{T(G)}, s\right) \cdot 1=A(\mathcal{G}, s)
$$

This concludes the proof in the case where $G$ is finitely generated.
Now we treat the most general case, where $G$ is not necessarily finitely generated or torsion-free; again, our task is to show that $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$. By Lemma 4.38, it suffices to show that $C\left(\mathcal{G}, s, f^{E}\right) \leq A(\mathcal{G}, s)$ for all $E \in \mathcal{E}^{*}(G)$ and all $f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)$ where each $\operatorname{supp}\left(f_{j}^{E}\right)=\phi_{j}(E)$. Pick any $E \in \mathcal{E}^{*}(G)$, let $G_{E} \leq G$ be the subgroup of $G$ generated by $E$, and examine the restricted datum $\left.\mathcal{G}\right|_{G_{E}}$. By Lemma 4.24, $s \in \mathcal{P}(\mathcal{G})$ implies $s \in \mathcal{P}\left(\left.\mathcal{G}\right|_{G_{E}}\right)$, so by applying this result in the finitely generated case (already established), we have that $C\left(\left.\mathcal{G}\right|_{G_{E}}, s\right) \leq A\left(\left.\mathcal{G}\right|_{G_{E}}, s\right)$, and moreover $A\left(\left.\mathcal{G}\right|_{G_{E}}, s\right) \leq A(\mathcal{G}, s)$ by Lemma 4.41. For any $f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)$ where each $\operatorname{supp}\left(f_{j}^{E}\right)=\phi_{j}(E)$, it follows that $\left(\left.f_{j}^{E}\right|_{\phi_{j}\left(G_{E}\right)}\right)_{j} \in \mathcal{F}^{*}\left(\left.\mathcal{G}\right|_{G_{E}}, s\right)$, therefore,

$$
C\left(\mathcal{G}, s, f^{E}\right)=C\left(\left.\mathcal{G}\right|_{G_{E}}, s,\left(\left.f_{j}^{E}\right|_{\phi_{j}\left(G_{E}\right)}\right)_{j}\right) \leq C\left(\left.\mathcal{G}\right|_{G_{E}}, s\right) \leq A(\mathcal{G}, s) ;
$$

thus we have $C(\mathcal{G}, s) \leq A(\mathcal{G}, s)$ by taking a supremum over all $E \in \mathcal{E}^{*}(G)$ and all associated $m$-functions $f^{E} \in \mathcal{F}^{*}(\mathcal{G}, s)$.

Lemma 4.48. For each $s \in[0, \infty)^{m}, B(\mathcal{G}, s)<\infty$ implies $s \in \mathcal{P}(\mathcal{G})$.

Proof. The key step of this proof is the following simple observation regarding free $\mathbb{Z}$ modules. Let $H, K$ be free $\mathbb{Z}$-modules with finite ranks $h, k$ and let $\phi: H \rightarrow K$ be a surjective $\mathbb{Z}$-module homomorphism (so $k \leq h$ ). Pick any $\mathbb{Z}$-bases of $H$ and $K$ and let $F \in \mathbb{Z}^{k \times h}$ be the (rank- $k$ ) matrix of $\phi$ with respect to these $\mathbb{Z}$-bases. Consider any element of $H$, represented as a column vector $x \in \mathbb{Z}^{h}$ in $H$ 's chosen $\mathbb{Z}$-basis, and the column vector $y \in \mathbb{Z}^{k}$ (representing an element of $K$ ) obtained by matrix-vector multiplication $y=F \cdot x$, defined componentwise for each $i \in\{1, \ldots, k\}$ by $y_{i}=\sum_{j=1}^{h} F_{i, j} x_{j}$. We have that

$$
\max _{i=1}^{k}\left|y_{i}\right|=\max _{i=1}^{k}\left|\sum_{j=1}^{h} F_{i, j} x_{j}\right| \leq \max _{i=1}^{k}\left(\sum_{j=1}^{h}\left|F_{i, j} x_{j}\right|\right) \leq\left(\max _{i=1}^{k} \sum_{j=1}^{h}\left|F_{i, j}\right|\right) \max _{j=1}^{h}\left|x_{j}\right|,
$$

Let $F^{*}=\max _{i=1}^{k} \sum_{j=1}^{h}\left|F_{i, j}\right|$; observe that $F^{*}$ depends only on $\phi$ and the chosen $\mathbb{Z}$-bases of $H$ and $K$. For each $N \in\{1,2, \ldots\}$, let $E_{N}=\{1, \ldots, N\}^{h}$ denote a set of $N^{h}$ column vectors representing elements of $H$ in its chosen $\mathbb{Z}$-basis. For each $x \in E_{N}$, we have that $\max _{j=1}^{h}\left|x_{j}\right| \leq N$. It follows that $\max _{i=1}^{k}\left|y_{i}\right| \leq F^{*} N$. Since this holds for all $y \in \phi\left(E_{N}\right)$, we have that $\phi\left(E_{N}\right) \subseteq\left\{-F^{*} N, \ldots, F^{*} N\right\}^{k}$, i.e., $\left|\phi\left(E_{N}\right)\right| \leq\left(2 F^{*}+1\right)^{k} N^{k}$.

Now we apply this observation for each homomorphism $\phi_{j}$ in the present setting. Consider any $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, and first consider the case where $H$ is finitely generated and torsion-free. As a torsion-free finitely generated Abelian group, by defining multiplication in terms of addition, $H$ becomes a free $\mathbb{Z}$-module. Letting $r=\operatorname{rank}(H)<\infty$, fix a $\mathbb{Z}$-basis $e_{1}, \ldots, e_{r}$ of $H$ and for each $N \in\{1,2, \ldots\}$, define $E_{N}=\left\{\sum_{i=1}^{r} n e_{i} \mid n \in\{1, \ldots, N\}\right\}$, so $\left|E_{N}\right|=N^{r}$.

Now consider each homomorphism $\phi_{j}$ in turn. Letting $H_{j}=\phi_{j}(H)$, we see that like $H$, $H_{j}$ is finitely generated and torsion-free and so can be treated as a free $\mathbb{Z}$-module. Defining $\chi_{j}: H \rightarrow H_{j}: x \mapsto \phi_{j}(x)$ to be a surjective $\mathbb{Z}$-module homomorphism, we can apply the preceding argument: therefore, there exists a constant $A_{j} \in[0, \infty)$ independent of $N$ such that $\left|\chi_{j}\left(E_{N}\right)\right| \leq A_{j} N^{r_{j}}$, where $r_{j}=\operatorname{rank}\left(H_{j}\right)$. Since the ranks of $H, H_{j}$ are same when interpreting $H, H_{j}$ as Abelian groups or as $\mathbb{Z}$-modules, and since $\chi_{j}\left(E_{N}\right)=\phi_{j}\left(E_{N}\right)$, we have that $\left|\phi_{j}\left(E_{N}\right)\right| \leq A_{j} N^{\operatorname{rank}\left(\phi_{j}(H)\right)}$. The hypothesis $B(\mathcal{G}, s)<\infty$ implies that $B\left(\mathcal{G}, s, E_{N}\right)<\infty$ for all $N \in\{1,2, \ldots\}$. In particular, there exists a constant $B \in(0, \infty)$ such that $\left|E_{N}\right| \leq$ $B \prod_{j}\left|\phi_{j}\left(E_{N}\right)\right|^{s_{j}}$, so

$$
N^{\operatorname{rank}(H)}=\left|E_{N}\right| \leq B \prod_{j}\left|\phi_{j}\left(E_{N}\right)\right|^{s_{j}} \leq B \prod_{j} A_{j} N^{s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)}=\left(B \prod_{j} A_{j}\right) N^{\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)} .
$$

Supposing $N \geq 2$ and taking base- $N$ logarithms, we have that $\operatorname{rank}(H) \leq \log _{N}\left(B \prod_{j} A_{j}\right)+$ $\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$. Since $B, A_{j}$ are independent of $N$, taking the limit as $N \rightarrow \infty$ we have that $\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$.

Now consider the case where $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$ is neither finitely generated nor torsion-free: pick any $\operatorname{rank}(H)$ independent elements of $H$ and examine the subgroup $K$ of $H$ they generate: $K$ is finitely generated, torsion free, and $\operatorname{rank}(H)=\operatorname{rank}(K)$, so by Lemma 4.3, $\operatorname{rank}\left(\phi_{j}(H)\right)=\operatorname{rank}\left(\phi_{j}(K)\right)$ for each $j$. We can thus apply the argument above to $K$ and obtain the same conclusion for $H$.

Repeating this argument for all $H \in \mathcal{H}_{\mathbb{N}}^{*}(G)$, we confirm that $s \in \mathcal{P}(\mathcal{G})$.
At this point, we are ready to assemble the preceding lemmas to prove Theorem 4.2, all that remains is to handle the cases of $s \in \mathcal{P}(\mathcal{G})$ such that $s \notin[0,1]^{m}$.

Proof of Theorem 4.2. If $s \in[0,1]^{m}$, then we apply Lemma 4.47 to show that $s \in \mathcal{P}(\mathcal{G}) \cap$ $[0,1]^{m}$ implies $A(\mathcal{G}, s)=C(\mathcal{G}, s)$; therefore by Lemma 4.31, $B(\mathcal{G}, s)=C(\mathcal{G}, s)=A(\mathcal{G}, s)$. Otherwise, letting $r=\left(\min \left(1, s_{j}\right)\right)_{j}$, we apply Lemma 4.39 to show that $B(\mathcal{G}, s)=B(\mathcal{G}, r)$ and $C(\mathcal{G}, s)=C(\mathcal{G}, r)$, and apply Lemma 4.5 to show that $r \in \mathcal{P}(\mathcal{G})$; since $r \in[0,1]^{m}$ as well, we conclude $B(\mathcal{G}, s)=C(\mathcal{G}, s)=A(\mathcal{G}, r)$ by the preceding case. This gives the first conclusion of Theorem 4.2. To show the second conclusion, for any $s \in[0, \infty)^{m}$, apply Lemma 4.48 to show that $B(\mathcal{G}, s)<\infty$ implies $s \in \mathcal{P}(\mathcal{G})$.

Finally, Theorem 4.1 follows as a special case of the first conclusion of Theorem 4.2.
Proof of Theorem 4.1. In this setting, we have the HBL datum $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)$; since $\mathbb{Z}^{r}$ is torsion-free, $A(\mathcal{G}, s)=1$ for all $s \in[0, \infty)^{m}$. Therefore, by Theorem 4.2, if $s \in \mathcal{P}(\mathcal{G})$, then $B(\mathcal{G}, s)=1$.

### 4.7 Computing $\mathcal{P}$

Now we address the decidability of the hypothesis of Theorem 4.2, that is, given an HBL datum $\mathcal{G}$ and $s \in[0, \infty)^{m}$, to decide whether $s \in \mathcal{P}(\mathcal{G})=\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathbb{N}}^{*}(G)\right)$. We will only address the case when $r=\operatorname{rank}(G)<\infty$, in order to exploit Lemma 4.9, we comment briefly on the general case after concluding the proof of Theorem 4.3. Then, in Section 4.7.1, we prove Theorem 4.4

By Lemma 4.23, since $r<\infty$, it suffices to compute $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{\mathrm{fg}}^{*}(H)\right)$ for any $H \in \mathcal{H}_{\mathrm{fg}}^{*}(G)$ with $\operatorname{rank}(H)=r$; moreover, there is no loss of generality to suppose $H$ is torsion-free. Therefore, there is no loss of generality hereafter to suppose that $G$ is finitely generated and torsion-free. In particular, there is no loss of generality to suppose $G=\mathbb{Z}^{r}$; therefore, $G$ 's elements ( $r$-tuples over $\mathbb{Z}$ ) and its action (componentwise addition) are computable.

When convenient, we will treat $G$ and its subgroups as $\mathbb{Z}$-modules and use moduletheoretic terminology; multiplication, defined in terms of addition, is also computable. The subgroups of $G$ as $\mathbb{Z}$-modules are finitely generated and torsion-free, thus free (admit $\mathbb{Z}$ bases), and their ranks as $\mathbb{Z}$-modules coincide with their ranks as Abelian groups, so we will not distinguish the two types of rank notationally. For computational reasons, we prefer to work in coordinates: in particular, we assume the standard ordered $\mathbb{Z}$-basis on $\mathbb{Z}^{r}$, meaning that each element of $G$, an $r$-tuple over $\mathbb{Z}$, coincides with its coordinate representation, a column $r$-vector over $\mathbb{Z}$. Each rank- $h$ subgroup $H$ of $G$ is (non-uniquely) represented by a $\mathbb{Z}$-basis matrix, an $r$-by- $h$ matrix over $\mathbb{Z}$ that horizontally concatenates the column vectors of some ordered $\mathbb{Z}$-basis of $H$; we will often use the same symbol $H$ to denote both a subgroup and any of its $\mathbb{Z}$-basis matrices.

We represent the homomorphisms $\phi_{j}$ by their kernels $K_{j} \leq G$, which we assume are provided in $\mathbb{Z}$-basis matrix representation. To see that the codomains $G_{j}$ are truly irrelevant to the task at hand, rewrite (4.8) in the present notation, $\operatorname{rank}(H) \leq \sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)=$ $\sum_{j} s_{j}\left(\operatorname{rank}(H)-\operatorname{rank}\left(H \cap K_{j}\right)\right)$.

We will compute the intersection $H=H_{1} \cap H_{2}$ of any two subgroups by solving the system of $\mathbb{Z}$-linear equations $H_{1} x_{1}=H_{2} x_{2}$ over $\mathbb{Z}$. In more detail, letting $h_{1}, h_{2}, h$ be the ranks of $H_{1}, H_{2}, H$, reduce the matrix $\left[H_{1},-H_{2}\right]=T U$ to (column-style) Hermite normal form, where $T \in \mathbb{Z}^{r \times\left(h_{1}+h_{2}\right)}$ is rank- $\left(h_{1}+h_{2}-h\right)$ and zero outside its first $h_{1}+h_{2}-h$
columns and $U \in \mathbb{Z}^{\left(h_{1}+h_{2}\right) \times\left(h_{1}+h_{2}\right)}$ is unimodular, represented as a sequence of elementary column operations. Compute $U^{-1}$ by applying the reverse sequence of elementary operations to an identity matrix and let $F \in \mathbb{Z}^{\left(h_{1}+h_{2}\right) \times h}$ be the last $h$ columns of $U^{-1}$; it follows that $\left[H_{1},-H_{2}\right] F=0_{r \times h}$. Split $F=\left[F_{1}^{T}, F_{2}^{T}\right]^{T}$ so $F_{1}$ is the top $h_{1}$ rows and $F_{2}$ is the bottom $h_{2}$ rows of $F$, and finally set $H=H_{1} F_{1}=H_{2} F_{2} \in \mathbb{Z}^{r \times h}$.

We now introduce some notation for inequalities of the form $h \leq \sum_{j} s_{j} h_{j}$, defined by parameters $h, h_{j} \in \mathbb{Z}$, which constrain particular $\mathbb{Z}$-linear combinations of the unknowns $s_{j} \in \mathbb{R}$. We encode such a constraint as $\left(h,\left(h_{j}\right)_{j}\right) \in \mathbb{Z} \times \mathbb{Z}^{m}$.

Lemma 4.49. For any $m \in\{1,2, \ldots\}$, given a closed, convex subset $\mathcal{P} \subseteq \mathbb{R}^{m}$ defined by a finite set of constraints on $s \in \mathbb{R}^{m}$, each of the form $h \leq \sum_{j} s_{j} h_{j}$ and encoded as $\left(h,\left(h_{j}\right)_{j}\right)$, there exists an algorithm that returns the set of $(\mathbb{Q}$-valued) extreme points of the convex polytope $\mathcal{P} \cap[0,1]^{m}$.

Proof. The following is a variation on a standard linear programming problem.
To enforce that $s \in[0,1]^{m}$, we augment the given set with the $2 m$ inequalities, $\left(0, e_{j}\right)$ and $\left(-1,-e_{j}\right)$ for each $j$. Now index the inequalities by $\alpha \in A$, yielding the nonempty, finite $\operatorname{set}\left\{\left(r_{\alpha},\left(r_{\alpha, j}\right)_{j}\right) \mid \alpha \in A\right\}$.

Any $t \in \mathbb{R}^{m}$ is an extreme point of $\mathcal{P} \cap[0,1]^{m}$ if and only if (1) $r_{\alpha} \leq \sum_{j} t_{j} r_{\alpha, j}$ for all $\alpha \in A$ and (2) there exists $B \subseteq A$ of size $|B|=m$ such that the $m$-tuples $\left(r_{\beta, j}\right)_{j}, \beta \in B$, are $\mathbb{Z}$-linearly independent and $r_{\beta}=\sum_{j} t_{j} r_{\beta, j}$ for all $\beta \in B$. (Condition (1) enforces that $t \in \mathcal{P}$ and condition (2) enforces that $t$ lies on $m \mathbb{Z}$-linearly independent bounding hyperplanes.) It follows that for each extreme point $t$ of $\mathcal{P} \cap[0,1]^{m}$, each component $t_{j} \in \mathbb{Q}$.

Identify the maximal subset $\mathcal{B}$ of the finite set $\{X \subset A||X|=m\}$ such that for each $B \in \mathcal{B}$, the $m$-tuples $\left(r_{\beta, j}\right)_{j}, \beta \in B$, are $\mathbb{Z}$-linearly independent. We test $\mathbb{Z}$-linear independence by defining $R \in \mathbb{Z}^{m \times m}$ with entries $R_{j, \beta}=r_{\beta, j}$, reducing $R=U T$ to (rowstyle) Hermite normal form where $U \in \mathbb{Z}^{m \times m}$ is unimodular and the number of leading nonzero rows of $T \in \mathbb{Z}^{m \times m}$ equals $m$ if and only if $\left(r_{\beta, j}\right)_{j}$ are $\mathbb{Z}$-linearly independent. For each $B \in \mathcal{B}$, the system of $m \mathbb{Z}$-linear equations $r_{\beta}=\sum_{j} t_{B, j} r_{\beta, j}, \beta \in B$ in the $m$ unknowns $\left(t_{B, j}\right)_{j}$ thus has a unique solution over $\mathbb{Q}$, which we compute by inverting $R$ over $\mathbb{Q}$ using Gauss-Jordan elimination and applying $R^{-1}$ to the column vector $\left(r_{\beta}\right)_{\beta}$. We then check whether, for all $\alpha \in A \backslash B, r_{\alpha} \leq \sum_{j} t_{B, j} r_{\alpha, j}$, by further rational arithmetic; if so, augment the set of extreme points, initially $\emptyset$, with $t=\left(t_{B, j}\right)_{j}$.

Another task that arises is enumerating the subgroups of $G=\mathbb{Z}^{r}$.
Lemma 4.50. Given $r, n \in \mathbb{N}$, there exists an algorithm that outputs the first $n$ elements of a fixed enumeration (independent of $n$, repetitions allowed) of the subgroups of $\mathbb{Z}^{r}$, represented as basis matrices.

Proof. If $n=0$, terminate with the empty tuple (). Otherwise, if $r=0$, terminate with the singleton tuple $(\{0\})$, where $\{0\}$ is represented as the empty matrix []. Otherwise, consider any (fixed) computable surjective function $A:\{1,2, \ldots\} \rightarrow \mathbb{Z}^{r \times r}$ : at least one such function (algorithm) exists since $\mathbb{Z}^{r \times r}$ is countable. Consider any $i \in\{1,2, \ldots\}$, not necessarily $i \leq n$. The matrix $A_{i}$ is rank- $r_{i}$ and its columns, interpreted as elements of $G$, span (over $\mathbb{Z}$ ) a rank- $r_{i}$ subgroup $H_{i} \leq G$. Moreover, $H_{i}$ is a surjective function from $i \in\{1,2, \ldots\}$ to
the subgroups $H_{i} \leq G$, i.e., every subgroup of $G$ is represented in the enumeration $\left(A_{i}\right)_{i}$. Returning $\left(H_{1}, \ldots, H_{n}\right)$ thus satisfies the conclusion.

It remains to show how to reduce each $A_{i}$ to a $\mathbb{Z}$-basis matrix of the associated subgroup $H_{i}$ : reduce each $A_{i}=T U$ to (column-style) Hermite normal form where $T \in \mathbb{Z}^{r \times r}$ is rank- $r_{i}$ and zero past its first $r_{i}$ columns and $U \in \mathbb{Z}^{r \times r}$ is unimodular, and let be the $H_{i} \in \mathbb{Z}^{r \times r_{i}}$ be the first $r_{i}$ columns of $T$ (the nonzero ones).

Note that a subgroup can appear more than once in the enumeration. Also note that this algorithm is generally non-halting, outputting subgroups indefinitely. However, Lemma 4.9 tells us that after some finite number of steps, the output will contain some (finite) subset $\mathcal{H}$ of subgroups such that $\mathcal{P}(\mathcal{G})=\mathcal{P}(\mathcal{G}, \mathcal{H})$. Thus, our strategy is to enumerate subgroups one-by-one, pausing after each step $i$ to check whether the output, a finite set of subgroups $\mathcal{H}_{i}=\left\{H_{1}, \ldots, H_{i}\right\}$, contains such an $\mathcal{H}$.

With this machinery at hand, we now show that $\mathcal{P}(\mathcal{G})$ is computable under the present assumption that $G=\mathbb{Z}^{r}$ for some $r<\infty$; recall that $\mathcal{H}^{*}(G)$ denotes the set of all subgroups of $G$, which here are all finitely generated, torsion free, and encoded as $\mathbb{Z}$-basis matrices. A similar algorithm to the following one was sketched by Valdimarsson following the proof of [35. Theorem 1.8], in the context of computing the hypotheses of [9, Theorem 2.1], which define convex set like $\mathcal{P}(\mathcal{G})$. Valdimarsson's approach avoids enumerating all subgroups (technically, subspaces), and instead only enumerates the lattice generated by $\left(K_{j}\right)_{j}$; our future work will show that it suffices to search such a lattice in the present setting as well.

Lemma 4.51. There exists an algorithm that returns a finite $\mathcal{H} \subseteq \mathcal{H}^{*}(G)$ such that $\mathcal{P}(\mathcal{G})=$ $\mathcal{P}(\mathcal{G}, \mathcal{H})$.

Proof. By Lemma 4.8, if $\operatorname{rank}(G)=0$, then $\mathcal{P}(\mathcal{G})=[0, \infty)^{m}$, so we return $\mathcal{H}=\emptyset$. Otherwise, $\operatorname{rank}(G) \geq 1$. By Lemma 4.11, letting $K=\bigcap_{j} K_{j}$, computed via pairwise intersections, if $\operatorname{rank}(K)>0$ then $\mathcal{P}(\mathcal{G})=\emptyset$, thus we terminate with $\mathcal{H}=\{K\} \subset \mathcal{H}^{*}(G)$. Otherwise, by Lemma 4.10, $\operatorname{rank}(K)=0$ and $(1, \ldots, 1) \in \mathcal{P}(\mathcal{G})$. We will now construct a finite $\mathcal{H} \subseteq \mathcal{H}^{*}(G)$ such that $\mathcal{P}(\mathcal{G}) \cap[0,1]^{m}=\mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]^{m}$ : the conclusion $\mathcal{P}(\mathcal{G})=\mathcal{P}(\mathcal{G}, \mathcal{H})$ follows from Lemma 4.5.

Consider the base case $m=1$ : Since $\operatorname{rank}\left(K_{1}\right)=0$, for all $H \in \mathcal{H}^{*}(G)$, $\operatorname{rank}(H)=$ $\operatorname{rank}\left(\phi_{1}(H)\right)$, so taking $\mathcal{H}=\{H\}$ for any $H \in \mathcal{H}^{*}(G)$ with $\operatorname{rank}(H)>0, \mathcal{P}(\mathcal{G}, \mathcal{H}) \cap[0,1]=$ $\mathcal{P}(\mathcal{G}) \cap[0,1]=\{1\}$. (We can also confirm this with Lemma 4.12; $s_{1}=\sum_{j} s_{j} \geq 1$.)

Now consider the inductive case $m \geq 2$ and $\operatorname{rank}(G) \geq 1$, meaning that we suppose the conclusion holds for all $\mathcal{G}^{\prime}$ such that $m^{\prime}<m$ or $\operatorname{rank}\left(G^{\prime}\right)<\operatorname{rank}(G)$. We run the following subroutine, which recursively invokes this lemma. Run the algorithm in Lemma 4.50 step-by-step, pausing after each step $i \in\{1,2, \ldots\}$ to examine the finite $\mathcal{H}_{i} \subseteq \mathcal{H}^{*}(G)$ computed thus far. We have that

$$
\emptyset \neq \mathcal{H}_{1} \subseteq \cdots \subseteq \mathcal{H}_{i} \subseteq \cdots \subseteq \mathcal{H}^{*}(G)
$$

so by Lemma 4.6,

$$
[0, \infty)^{m} \supseteq \mathcal{P}\left(\mathcal{G}, \mathcal{H}_{1}\right) \supseteq \cdots \supseteq \mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \supseteq \cdots \supseteq \mathcal{P}(\mathcal{G})
$$

Moreover, by Lemma 4.9, there exists a $l \in\{1,2, \ldots\}$ such that $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right)=\mathcal{P}(\mathcal{G})$ for all $i \geq l$, so we only need to take finitely many steps $i$ (however, we don't yet have bounds
on how big $i$ has to be). Determine the set $\left\{\left(\operatorname{rank}(H),\left(\operatorname{rank}\left(\phi_{j}(H)\right)\right)_{j}\right) \mid H \in \mathcal{H}_{i}\right\}$ of linear inequalities that $\mathcal{H}_{i}$ generates: note that each subgroup's rank is stored as part of the $\mathbb{Z}$-basis matrix representation, and we can compute $\operatorname{rank}\left(\phi_{j}(H)\right)=\operatorname{rank}(H)-\operatorname{rank}\left(H \cap K_{j}\right)$ from $H$ and $K_{j}$ using the algorithm for computing subgroup intersections given earlier. We then apply Lemma 4.49 to obtain the finite set of extreme points of $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \cap[0,1]^{m}$, each of which is $\mathbb{Q}$-valued.

We will analyze the finitely many extreme points of $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \cap[0,1]^{m}$ in turn. Each extreme point $s$ of $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \cap[0,1]^{m}$ falls into one of three cases, and depending on the outcome, the algorithm either terminates or moves on to the next extreme point.

In the first case, $s \in(0,1)^{m}$ - note that we can decide inequality over $\mathbb{Q}$, and testing this particular case involves $2 m$ inequality decisions. We can adapt the algorithm in Lemma 4.49 to exploit Lemma 4.17 and return $H \in \mathcal{H}_{i}$ such that $0<\operatorname{rank}(H)=\sum_{j} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)<$ $\operatorname{rank}(G)$.

A complication now arises: even though $G$ and $H$ are torsion-free, the quotient group $G / H$ need not be, and thus may not have a well-defined basis-matrix representation. As in the proof of Theorem 4.2, we apply Lemma 4.1 to avoid torsion during the factorization step. We now give an algorithm that implements Lemma 4.1, i.e., computing a subgroup $H \leq K \leq G$ such that $G / K$ is torsion-free and $\operatorname{rank}(H)=\operatorname{rank}(K)$. Additionally, since $K$ has a complement in $G$, we will represent $G / K$ by this subgroup. In case $h=\operatorname{rank}(H)=0$, it follows that $H=\{0\}$, thus we simply return $K=H$, and represent $G / K$ by $G$, i.e., the $r$-by- $r$ identity matrix. In case $h=r=\operatorname{rank}(G)$, we simply return $K=G$ and represent $G / K$ by $\{0\}$, i.e., the empty matrix []. Otherwise $0<h<r$. Given the $r$-by- $h$ matrix $H$, reduce $H$ to (row-style) Hermite normal form $H=U T$ where $U \in \mathbb{Z}^{r \times r}$ is unimodular and $T \in \mathbb{Z}^{r \times h}$ is full rank and zero below its top $h$ rows. Represent $K$ by the leading $r$-by- $h$ submatrix of $U$ and represent $G / K$ by the trailing $r$-by- $(r-h)$ submatrix of $U$.

It follows from Lemma 4.3 that $\operatorname{rank}\left(\phi_{j}(H)\right)=\operatorname{rank}\left(\phi_{j}(K)\right)$ for each $j$ as well, so combining Lemmas 4.24 and $4.25, s \in \mathcal{P}(\mathcal{G})$ if and only if $s \in \mathcal{P}\left(\left.\mathcal{G}\right|_{K}\right) \cap \mathcal{P}\left(\left.\mathcal{G}\right|^{K}\right)$.

We now show how to replace $\mathcal{G}$ by $\left.\mathcal{G}\right|_{K}$ and $\left.\mathcal{G}\right|^{K}$; let $r=\operatorname{rank}(G)$ and $k=\operatorname{rank}(K)$. For each $j$, compute $K_{j}^{\prime}=K \cap K_{j}$ and $K_{j}^{\prime \prime}=G / K \cap K_{j}$ (recalling that $G / K$ is represented as a complement of $K$ in $G): K_{j}^{\prime}$ and $K_{j}^{\prime \prime}$ represent $\operatorname{ker}\left(\chi_{j}\right)$ and $\operatorname{ker}\left(\psi_{j}\right)$, respectively, as subgroups of $G$; however, to apply the induction hypotheses, we need $\operatorname{ker}\left(\chi_{j}\right)$ to be represented as a subgroup of $K$, and we need $\operatorname{ker}\left(\psi_{j}\right)$ to be represented as a subgroup of $G / K$. We will change coordinates according the unimodular matrix $U$ obtained above: letting $V_{1}$ and $V_{2}$ denote the top $k$ and bottom $r-k$ rows of $U^{-1}$, computed by applying the reverse sequence of elementary operations to an identity matrix, we represent each $\operatorname{ker}\left(\chi_{j}\right)$ by $V_{1} K_{j}^{\prime}$ and each $\operatorname{ker}\left(\psi_{j}\right)$ by $V_{2} K_{j}^{\prime \prime}$.

Since $\operatorname{rank}(K), \operatorname{rank}(G / K)<\operatorname{rank}(G)$, we can apply the induction hypotheses (i.e., apply the algorithm recursively) to compute finite sets $\mathcal{H}^{\prime} \subseteq \mathcal{H}^{*}\left(\left.\mathcal{G}\right|_{K}\right)$ and $\mathcal{H}^{\prime \prime} \subseteq \mathcal{H}^{*}\left(\left.\mathcal{G}\right|^{K}\right)$ such that $\mathcal{P}\left(\left.\mathcal{G}\right|_{K}\right) \cap[0,1]^{m}=\mathcal{P}\left(\left.\mathcal{G}\right|_{K}, \mathcal{H}^{\prime}\right) \cap[0,1]^{m}$ and $\mathcal{P}\left(\left.\mathcal{G}\right|^{K}\right) \cap[0,1]^{m}=\mathcal{P}\left(\left.\mathcal{G}\right|^{K}, \mathcal{H}^{\prime \prime}\right) \cap[0,1]^{m}$. We can then test whether $s \in \mathcal{P}(\mathcal{G})$ by testing whether, for all $K^{\prime} \in \mathcal{H}^{\prime}$ and $K^{\prime \prime} \in \mathcal{H}^{\prime \prime}$, $\operatorname{rank}\left(K^{\prime}\right) \leq$ $\sum_{j} s_{j} \operatorname{rank}\left(\chi_{j}\left(K^{\prime}\right)\right)$ and $\operatorname{rank}\left(K^{\prime \prime}\right) \leq \sum_{j} s_{j} \operatorname{rank}\left(\psi_{j}\left(K^{\prime \prime}\right)\right)$, using rational arithmetic. If any of these tests fail, then we proceed with taking the next step $i+1$ of the algorithm in Lemma 4.50, and repeat the preceding steps with respect to $\mathcal{H}_{i+1}$.

In the second case, some component $s_{k}=1$; we can reduce $\mathcal{G}$ to the HBL datum $\mathcal{G}^{\prime}=$ $\left(K_{k},\left(G_{j}\right)_{j \neq k},\left(\left.\phi_{j}\right|_{K_{k}}\right)_{j \neq k}\right)$ by dropping the $k$-th homomorphism and, for each $j \neq k$, replacing
$K_{j}$ by $V_{1} K_{j}$, where $V_{1}$ is computed from $K_{k}$ as it was from the matrix $K$ in the preceding case. By Lemma 4.26, we can apply this algorithm recursively to $\mathcal{G}^{\prime}$, which has one less homomorphism, obtaining a finite set $\mathcal{H}^{\prime} \subseteq \mathcal{H}^{*}(G)$ such that $\left(s_{j}\right)_{j \neq k} \in \mathcal{P}\left(\mathcal{G}^{\prime}, \mathcal{H}^{\prime}\right)$ if and only if $s \in \mathcal{P}(\mathcal{G})$. To check whether $s \in \mathcal{P}(\mathcal{G})$, we test whether $\operatorname{rank}(H) \leq \sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$ for all $H \in \mathcal{H}^{\prime}$. If any of these tests fail, then we continue with the next step $i+1$ of the algorithm in Lemma 4.50, and repeat the preceding steps with respect to $\mathcal{H}_{i+1}$.

In the third case, some component $s_{k}=0$; by Lemma 4.13, we can apply this algorithm recursively to the $\operatorname{HBL}$ datum $\mathcal{G}^{\prime}=\left(G,\left(G_{j}\right)_{j \neq k},\left(\phi_{j}\right)_{j \neq k}\right)$, which has one less homomorphism (i.e., we omit $K_{k}$ ), obtaining a finite set $\mathcal{H}^{\prime} \subseteq \mathcal{H}_{\mathrm{fg}}^{*}(\mathcal{G})$ such that $\left(s_{j}\right)_{j \neq k} \in \mathcal{P}\left(\mathcal{G}^{\prime}, \mathcal{H}^{\prime}\right)$ if and only if $s \in \mathcal{P}(\mathcal{G})$. To check whether $s \in \mathcal{P}(\mathcal{G})$, we test whether $\operatorname{rank}(H) \leq \sum_{j \neq k} s_{j} \operatorname{rank}\left(\phi_{j}(H)\right)$ for all $H \in \mathcal{H}^{\prime}$, using rational arithmetic. If any of these tests fail, then continue with the next step $i+1$ of the algorithm in Lemma 4.50, and repeat the preceding steps with respect to $\mathcal{H}_{i+1}$.

If every extreme point $s \in \mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \cap[0,1]^{m}$ is contained in $\mathcal{P}(\mathcal{G})$, which we have already confirmed must happen within finitely many steps, then because any closed and bounded convex subset of $\mathbb{R}^{m}$ equals the convex hull of its extreme points, $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \cap[0,1]^{m} \subseteq$ $\mathcal{P}(\mathcal{G})$. We already concluded the converse inclusion $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right) \supseteq \mathcal{P}(\mathcal{G})$ by Lemma 4.6, thus $\mathcal{P}\left(\mathcal{G}, \mathcal{H}_{i}\right)=\mathcal{P}(\mathcal{G})$ and we return the finite set $\mathcal{H}=\mathcal{H}_{i}$.

Note that the output of the algorithm in Lemma 4.51 can be then used to decide membership in $\mathcal{P}(\mathcal{G})$ via checking each of a finite list of $\mathbb{Z}$-linear inequalities (this task is already performed within the algorithm, in the inductive step). Since $\mathcal{P}(\mathcal{G})$ is a subset of $\mathbb{R}^{m}$, cocountably many points of which are uncomputable, membership can only be decided for the computable elements of $\mathbb{R}^{m}$.

Lemma 4.51 nearly completes the proof of Theorem 4.3. except for one caveat that was glossed over in the theorem statement.

Proof of Theorem 4.3. Lemma 4.51 gives the the desired conclusion in the case $G=\mathbb{Z}^{r}$, assuming that $\mathcal{G}$ is represented as an $m$-tuple $\left(K_{j}\right)_{j}$ of $\mathbb{Z}$-basis matrices, as defined above. We showed earlier that when $G$ is a finite rank Abelian group, the membership problem reduces to the study of the case $G=\mathbb{Z}^{r}$. Whether or not this reduction itself is computable depends on the given representation of $\mathcal{G}$ : we have sidestepped this potential obstacle by assuming a particular representation.

We briefly comment on the most general case, where $\operatorname{rank}(G)=\infty$. An immediate obstacle facing our approach is that $\mathcal{H}_{\mathbb{N}}^{*}(G)$ need not be countable, so even if there exists a finite $\mathcal{H} \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$ such that $\mathcal{P}(\mathcal{G})=\mathcal{P}(\mathcal{G}, \mathcal{H})$, it may not appear after finitely many steps of the algorithm in Lemma 4.50. However, since $\mathcal{P}(\mathcal{G})$ is defined by a countable (not necessarily finite) intersection of half-spaces of $\mathbb{R}^{m}$, there exists a countable analogue of Lemma 4.9. Therefore, supposing first that we can resolve the issues of representing $\mathcal{G}$ and computing $G$ 's elements and action, given an algorithm (like the one in Lemma 4.50) that enumerates a countable $\mathcal{H} \subseteq \mathcal{H}_{\mathbb{N}}^{*}(G)$ such that $\mathcal{P}(\mathcal{G})=\mathcal{P}(\mathcal{G}, \mathcal{H})$, we can extend the algorithm in Lemma 4.51 into a nondeterministic one that halts whenever $s \notin \mathcal{P}(\mathcal{G})$, and otherwise potentially runs forever. In other words, membership in $\mathcal{P}(\mathcal{G})$ may still be cosemidecidable in the case $\operatorname{rank}(G)=\infty$. These comments are motivated by Valdimarsson's result 35 , Theorem 1.8] regarding the hypotheses of [9, Theorem 2.1], demonstrating in the continuum
setting, where $G$ is a finite-dimensional $\mathbb{R}$-vector space, that it suffices to search a countable set of subspaces - the lattice generated by $\left(K_{j}\right)_{j}$ - despite $G$ having (as an Abelian group) both infinite rank and uncountably many finitely generated subgroups.

### 4.7.1 Computing the Inequalities Defining $\mathcal{P}(\mathcal{G})$

Our last investigation regarding computability of $\mathcal{P}(\mathcal{G})$ concerns the question of whether a simpler algorithm would suffice. In particular, in the case $G=\mathbb{Z}^{r}$ for $r<\infty$, given arbitrary $h, h_{j} \in\{0, \ldots, r\}$, can we decide whether there exists $H \leq G$ such that $\operatorname{rank}(H)=h$ and each $\operatorname{rank}\left(\phi_{j}(H)\right)=h_{j}$ ? (Actually computing $H$ is not required.) An algorithm that decides this problem could be used to iterate over the $(r+1)^{m+1}$ possible values of $h, h_{j}$, outputting a $\mathbb{Z}$-linear inequality $h \leq \sum_{j} s_{j} h_{j}$ over $s \in \mathbb{R}^{m}$, encoded as $\left(h,\left(h_{j}\right)_{j}\right)$, whenever the answer is affirmative. It then remains to test whether the (computable) $m$-tuple $s \in \mathbb{R}^{m}$, satisfies a finite set of inequalities, which we have already shown is decidable. However, we have so far failed to devise such an algorithm: we will show that this question of decidability is equivalent to Hilbert's Tenth Problem for $\mathbb{Q}$, a longstanding open problem.

For the following discussion, we will change our notation yet again. In particular, we now suppose that the homomorphisms $\phi_{j}$ are endomorphisms of $G=\mathbb{Z}^{r}$ : we can always satisfy this assumption by identifying each homomorphic image $\phi_{j}(G)$ with an isomorphic subgroup of $G$. We now represent each $\phi_{j}$ as an $r$-by- $r$ matrix over $\mathbb{Z}$ with respect to the standard ordered $\mathbb{Z}$-basis on $\mathbb{Z}^{r}$ (rather than by its kernel $K_{j}$ ), and use the symbol $\phi_{j}$ to refer to both the endomorphism and its matrix. We also introduce the notation $\phi=\left(\phi_{j}\right)_{j} \in\left(\mathbb{Z}^{r \times r}\right)^{m}$ to denote an $m$-tuple over $\mathbb{Z}^{r \times r}$.

In the following definitions, we will replace $\mathbb{Z}$ by a more general integral domain, denoted $\mathbb{D}$; we are primarily interested in the integral domains $\mathbb{Z}$ and $\mathbb{Q}$. For any integral domain $\mathbb{D}, m, r \in\{1,2, \ldots\}$, and $h, h_{j} \in \mathbb{N}$, define (using the notations $\leq_{\mathbb{D}}$ and $\mathrm{rank}_{\mathbb{D}}$ introduced in Section 4.4.2,

$$
E_{h,\left(h_{j}\right)_{j}}^{\mathbb{D}, m, r}=\left\{\phi \in\left(\mathbb{D}^{r \times r}\right)^{m} \mid\left(\exists H \leq_{\mathbb{D}} \mathbb{D}^{r}\right) \operatorname{rank}_{\mathbb{D}}(H)=h \wedge(\forall j) \operatorname{rank}_{\mathbb{D}}\left(\phi_{j}(H)\right)=h_{j}\right\} ;
$$

in our application, $\mathbb{D}=\mathbb{Z}, m, r=\operatorname{rank}(G)$, and $\phi=\left(\phi_{j}\right)_{j}$ are specified by $\mathcal{G}$, while $h, h_{j}$ are specified by the particular inequality $\left(h,\left(h_{j}\right)_{j}\right)$ whose existence in $\mathcal{P}(\mathcal{G})$ we are currently probing; thus, our problem is deciding whether $\phi \in E_{h,\left(h_{j}\right)_{j}}^{\mathbb{D}, m, r}$. This problem reduces to one about solutions of polynomial equations, not over $\mathbb{D}$, but over $\mathbb{D}$ 's field of fractions. We demonstrate this below in Lemma 4.53 in the case $\mathbb{D}=\mathbb{Z}$, with field of fractions $\mathbb{Q}$.

First we pause to introduce some notation for polynomials, as well as introduce Hilbert's Tenth Problem. For any integral domain $\mathbb{D}$ and $t \in \mathbb{N}$, let $\mathbb{D}\left[x_{1}, \ldots, x_{t}\right]$ denote the ring of polynomials over $\mathbb{D}$ in the variables $x_{1}, \ldots, x_{t}$. For any $S \subseteq \mathbb{D}\left[x_{1}, \ldots, x_{t}\right]$, define

$$
Z(S)=\left\{a \in \mathbb{D}^{t} \mid(\forall f \in S) f(a)=0\right\}
$$

with $Z(f)$ being shorthand for $Z(\{f\})$ when $f \in \mathbb{D}\left[x_{1}, \ldots, x_{t}\right]$. Hilbert's Tenth Problem for $\mathbb{D}$ is the question of whether $Z(S) \neq 0$ is decidable for all $t \in \mathbb{N}$ and all finite $S \subseteq \mathbb{D}\left[x_{1}, \ldots, x_{t}\right]$. Matiyasevich-Davis-Putnam-Robinson's theorem (see, e.g., [27]) famously resolved Hilbert's Tenth Problem (for $\mathbb{Z}$, the version Hilbert posed) with a negative answer: there exists no general algorithm for testing solvability over $\mathbb{Z}$. Hilbert's Tenth Problem for the computable
real numbers, however, was resolved with an affirmative answer by Tarski [32]; the asymptotic complexity of Tarski's original algorithm has been improved many times, with Collins' algorithm [17], called cylindrical algebraic decomposition, being a milestone. Hilbert's Tenth Problem for $\mathbb{Q}$, on the other hand, remains open (see, e.g., [28]).

We already saw in Section 4.4.2 that $\mathcal{P}(\mathcal{G})=\mathcal{P}_{\mathbb{Q}}\left(\mathcal{G}^{\mathbb{Q}}\right)$, where $\mathcal{G}^{\mathbb{Q}}=\left(G^{\mathbb{Q}},\left(G_{j}^{\mathbb{Q}}\right)_{j},\left(\phi_{j}^{\mathbb{Q}}\right)_{j}\right)$ is obtained from $\mathcal{G}$ by extending scalars from $\mathbb{Z}$ to $\mathbb{Q}$, and $\mathcal{P}_{\mathbb{Q}}$ is obtained from $\mathcal{P}$ by replacing Abelian group ranks with $\mathbb{Q}$-module ranks in its definition, (4.8).

Lemma 4.52. $E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Z}, m, r}=E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m, r} \cap\left(\mathbb{Z}^{r \times r}\right)^{m}$.
Proof. This follows from specializing Lemma 4.4 to the inclusion $\mathbb{Z} \subset \mathbb{Q}$.
Lemma 4.53. Each instance of the problem of whether $\phi \in E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m, r}$, given $\phi, m, r, h, h_{j}$, computably reduces to an instance of the problem of whether $Z(f) \neq \emptyset$ for some $f \in \mathbb{Q}\left[x_{1}, \ldots, x_{t}\right]$ and $t \in \mathbb{N}$.

Proof. Recall that if $M$ is a matrix over $\mathbb{Q}, \operatorname{rank}_{\mathbb{Q}}(M)=s$ if and only if all $(s+1)$-by- $(s+1)$ minors of $M$ are zero and at least one $s$-by-s minor is nonzero. Since for any $x, y \in \mathbb{Q}$, $x^{2}+y^{2}=0$ if and only if $x=y=0$, we can reduce any finite set of polynomial equations to a single equation. In particular, the condition that all $(s+1)$-by- $(s+1)$ minors of $M$ are zero reduces to an equality of the form $\sum_{f} f^{2}=0$, and the condition that not all $s$-by- $s$ minors are zero reduces to an inequality of the form $\sum_{g} g^{2} \neq 0$. The latter expression can be converted to an equality, since for any $g \in \mathbb{Q}[X], g \neq 0$ is equivalent to $f=g y-1=0$ where $y \notin X$ and $f \in \mathbb{Q}[X \cup\{y\}]$. Now the remaining two polynomial equations can be combined into a single equation.

Let $B$ be a $r$-by- $h$ matrix of variables and let $\phi=\left(\phi_{j}\right)_{j} \in\left(\mathbb{Q}^{r \times r}\right)^{m}$ as given. Now using the preceding argument, define $m+1$ polynomial equations for the matrices $M=B$ (with $s=h)$ and $M=\phi_{j} B\left(\right.$ with $\left.s=h_{j}\right)$ for each $j$, and then finally combine them into a single equation.

If Hilbert's Tenth Problem for $\mathbb{Q}$ has an affirmative solution, then we can use the associated algorithm to directly test each candidate inequality $\left(h,\left(h_{j}\right)\right)$.

Now we develop a stronger converse result than we need to show equivalence with Hilbert's Tenth Problem for $\mathbb{Q}$, that the ability to decide membership in $E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m, r}$ enables deciding membership in arbitrary Diophantine sets (over $\mathbb{Q}$ ). Given $p, q \in \mathbb{N}$ and $S \subseteq$ $\mathbb{Q}\left[x_{1}, \ldots, x_{p} ; y_{1}, \ldots, y_{q}\right]$, the parameter triple $(p, q, S)$ define the Diophantine set (over $\mathbb{Q}$ ),

$$
D(p, q, S)=\left\{a \in \mathbb{Q}^{p} \mid\left(\exists b \in \mathbb{Q}^{q}\right)(a, b) \in Z(S)\right\}
$$

Our task above, deciding whether $Z(S) \neq \emptyset$ for some $S \subseteq \mathbb{Q}\left[x_{1}, \ldots, x_{s}\right]$ and $s \in \mathbb{N}$, is equivalent to deciding whether the empty tuple ()$\in D(0, s, S) \subseteq \mathbb{Q}^{0}=\{()\}$; a different task, when additionally given some $a \in \mathbb{Q}^{s}$, is deciding whether $a \in Z(S)$, which is equivalent to deciding whether $a \in D(s, 0, S) \subseteq \mathbb{Q}^{s}$.

At this point, we observe that $\mathbb{Q}$ is a Noetherian ring, so by Hilbert's basis theorem, for any $S \subseteq \mathbb{Q}\left[x_{1}, \ldots, x_{s}\right]$, there exists a finite $R \subseteq \mathbb{Q}\left[x_{1}, \ldots, x_{s}\right]$ such that $Z(R)=Z(S)$. In particular, for any Diophantine set $D(p, q, S)$, where $p, q \in \mathbb{N}$ such that $s=p+q$, there is no loss of generality to suppose $S$ is finite. From here on, we will suppose that all sets of polynomials are finite.

Lemma 4.54. Each instance of the problem of whether $a \in D(p, q, S)$, given $a, p, q, S$, computably reduces to the problem of $\phi \in E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m, r}$, for some $\phi, m, r, h, h_{j}$.

Proof. Our first step is to simplify the presentation of $S$ by (algorithmically) constructing $T$, a set of simpler polynomials in a larger set of variables, such that $Z(S)$ equals $Z(T)$ on a subset of $T$ 's variables, thus computably reducing the membership problem.

Let $s=p+q$ and substitute the variables $\left(z_{1}, \ldots, z_{s}\right)$ for $\left(x_{1}, \ldots, x_{p} ; y_{1}, \ldots, y_{q}\right)$, rewriting $S$ as a (finite) subset of $\mathbb{Q}\left[z_{1}, \ldots, z_{s}\right]$; we will dedicate the subscript $i$ to range over $\{1, \ldots, s\}$. Compute $d=\max _{f \in S} \max _{i} \operatorname{deg}_{z_{i}}(f)$ and define the index set $A=\{0, \ldots, d\}^{s}$ and $(d+1)^{s}$ variables $v_{\alpha}$ indexed by $\alpha=\left(\alpha_{i}\right)_{i} \in A$. Construct a finite $T_{1} \subseteq \mathbb{Q}\left[\left\{v_{\alpha}\right\}_{\alpha \in A}\right]$, initially $T_{1}=\emptyset$, by examining each $f \in S$, which can be represented as $f=\sum_{\alpha \in A} c_{\alpha} \prod_{i} z_{i}^{\alpha_{i}}$ where each $c_{\alpha} \in \mathbb{Q}$ is uniquely defined by $f$. For each $f \in S$, augment $T_{1}$ with the polynomial $\sum_{\alpha \in A} c_{\alpha} v_{\alpha}$. Also construct a finite $T_{2} \subseteq \mathbb{Q}\left[\left\{v_{\alpha}\right\}_{\alpha \in A}\right]$, initially $T_{2}=\emptyset$, by examining each $\alpha \in A$ : if $\alpha=(0, \ldots, 0)$, augment $T_{2}$ with the polynomial $v_{(0, \ldots, 0)}-1$; otherwise, keeping $\alpha$ fixed, examine each $\beta \in A$ : if $\beta \neq(0, \ldots, 0)$ and $\alpha+\beta \in A$, augment $T_{2}$ with the polynomial $v_{\alpha+\beta}-v_{\alpha} v_{\beta}$ (if not already in $T_{2}$ ). Now let $T=T_{1} \cup T_{2}$.

We now show that $Z(T)$ equals $Z(S)$ when projecting the former set onto to its coordinates corresponding to the $s$ variables $v_{e_{i}}$, where each $e_{i} \in A$ is zero except for its unit $i$-th entry; this projection, $\pi: \mathbb{Q}^{A} \rightarrow \mathbb{Q}^{s}:\left(v_{\alpha}\right)_{\alpha \in A} \mapsto\left(v_{e_{i}}\right)_{i}$, is computable. Now define $\chi: \mathbb{Q}^{s} \rightarrow \mathbb{Q}^{A}:\left(z_{i}\right)_{i} \mapsto\left(\prod_{i} z_{i}^{\alpha_{i}}\right)_{\alpha \in A}$, also computable. By construction, $\chi\left(\mathbb{Q}^{s}\right)=Z\left(T_{2}\right) \subseteq \mathbb{Q}^{A} ;$ defining $\psi$ to be $\chi$ with its codomain restricted to $Z\left(T_{2}\right)$, we see that $\psi$ is a bijection with $\psi^{-1}=\left.\pi\right|_{Z\left(T_{2}\right)}$. Moreover, for any $a \in \mathbb{Q}^{s}, a \in Z(S)$ if and only if $\chi(a) \in Z(T)$. Therefore, $Z(S)=\pi(Z(T))$.

Thus, given $a \in \mathbb{Q}^{p}$, deciding whether $a \in D(p, q, S)$ is equivalent to deciding whether $a \in D\left(p,(d+1)^{p+q}-p, T\right)$, linearizing the indices $\alpha \in A$ so that $T \subseteq \mathbb{Q}\left[v_{e_{1}}, \ldots, v_{e_{p}} ; \ldots\right]$, where the second ellipsis indicates the variables' order is arbitrary after the first $p$. Moreover, each instance of the former problem (with $S$ ) computably reduces to an instance of the latter problem (with $T$ ).

In light of this reduction, we suppose that the given $S$ has the form of $T$. That is, $S \subseteq \mathbb{Q}\left[z_{1}, \ldots, z_{s}\right]$ with $s=p+q$, and there exist $k, l \in \mathbb{N}$ such that $|S|=k+l$ and the elements of $S$ can be expressed as

$$
S=\left\{z_{i_{1, j}} z_{i_{2, j}}-z_{i_{3, j}}\right\}_{j=1}^{k} \cup\left\{\lambda_{0, j}+\sum_{i} \lambda_{i, j} z_{i}\right\}_{j=1}^{l}
$$

where the subscript $i$ continues to range over $\{1, \ldots, s\}$, and the double subscripts $i_{1, j}, i_{2, j}, i_{3, j} \in$ $\{1, \ldots, s\}$ are distinct for each $j$ but not necessarily between different $j$. We suppose that each $\lambda_{0, j}, \lambda_{i, j} \in \mathbb{Z}$ : this can be computably enforced, scaling by integers to clear denominators, without affecting $Z(T)$.

Now define $m=4+s+k+p+l, r=2 s+2, h=2$, and $h_{j}=1$ for $j \leq 4+s+k$ and $h_{j}=0$ otherwise ( $j$ continues to range over $\{1, \ldots, m\}$ unless specified otherwise). We will use the coordinates $(u ; v)=\left(u_{0}, \ldots, u_{s} ; v_{0}, \ldots, v_{s}\right)$ on $\mathbb{Q}^{r}$. Instead of working with a fixed $m$-tuple $\left(\phi_{j}\right)_{j} \in\left(\mathbb{Q}^{r \times r}\right)^{m}$, we will parameterize $\left(\phi_{j}^{(a)}\right)_{j}$ by the input $a \in \mathbb{Q}^{p}$ to the Diophantine membership problem, where we assume that $a=\left(\alpha_{1} / \beta_{1}, \ldots, \alpha_{p} / \beta_{p}\right)$ is given in
lowest terms. We then compute $\left(\phi_{j}^{(a)}\right)_{j}$ componentwise from $a$ as follows:

$$
\begin{array}{rlrl}
\phi_{1}^{(a)}:(u ; v) & \mapsto\left(u_{0}, 0, \ldots 0 ; 0, \ldots, 0\right) & & \\
\phi_{2}^{(a)}:(u ; v) & \mapsto\left(v_{0}, 0, \ldots 0 ; 0, \ldots, 0\right) & & \\
\phi_{3}^{(a)}:(u ; v) & \mapsto\left(u_{0}, \ldots, u_{s} ; 0, \ldots, 0\right) & & \\
\phi_{4}^{(a)}:(u ; v) & \mapsto\left(v_{0}, \ldots, v_{s} ; 0, \ldots, 0\right) & & \\
\phi_{4+j}^{(a)}:(u ; v) & \mapsto\left(u_{0}-v_{0}, u_{j}-v_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right) & & \\
\phi_{4+s+j}^{(a)} & :(u ; v) & \mapsto\left(u_{i_{1, j}}+v_{0}, u_{i_{3, j}}+v_{i_{2, j}}, 0, \ldots, 0 ; 0, \ldots, 0\right) & j \in\{1, \ldots, k\} \\
\phi_{4+s+k+j}^{(a)}: & :(u ; v) & \mapsto\left(\alpha_{j} u_{0}-\beta_{j} u_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right) & j \in\{1, \ldots, p\} \\
\phi_{4+s+k+p+j}^{(a)}:(u ; v) & \mapsto\left(\lambda_{0, j} u_{0}+\sum_{i} \lambda_{i, j} u_{i}, 0, \ldots, 0 ; 0, \ldots, 0\right) & j \in\{1, \ldots, l\} ;
\end{array}
$$

note that only $\phi_{4+s+k+1}^{(a)}, \ldots, \phi_{4+s+k+p}^{(a)}$ actually depend on the $p$-tuple $a$, and none of these matrices exist in the case $p=0$. It remains to show that $a \in D(p, q, S)$ if and only if $\left(\phi_{j}^{(a)}\right)_{j} \in E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m}$. We will show sufficiency and then necessity.

Suppose we are given $a=\left(\alpha_{1} / \beta_{1}, \ldots, \alpha_{p} / \beta_{p}\right) \in \mathbb{Q}^{p}$ (in lowest terms) such that $a \in$ $D(p, q, S)$, i.e., there exists $b \in \mathbb{Q}^{q}$ such that $(a, b) \in Z(S)$. Introduce the notation $\left(c_{1}, \ldots, c_{s}\right)=$ $\left(a_{1}, \ldots, a_{p}, b_{1}, \ldots, b_{q}\right)$ and define

$$
V=\mathbb{Q}\left(1, c_{1}, \ldots, c_{s} ; 0, \ldots, 0, \ldots, 0\right)+\mathbb{Q}\left(0, \ldots, 0 ; 1, c_{1}, \ldots, c_{s}\right) .
$$

Each element of $V$ has the form $\left(\mu, \mu c_{1}, \ldots, \mu c_{s}, \nu, \nu c_{1}, \ldots, \nu c_{s}\right)$ for $(\mu, \nu) \in \mathbb{Q}^{2}$; we thus identify elements of $V$ by their coordinates $(\mu, \nu)$. Clearly $\operatorname{rank}_{\mathbb{Q}}(V)=2=h$, as requested. We now check that for each component $j, \operatorname{rank}_{\mathbb{Q}}\left(\phi_{j}^{(a)}(V)\right)=h_{j}$. First and second, $\phi_{1}^{(a)}(V)=$ $\mathbb{Q}(1,0, \ldots, 0 ; 0, \ldots, 0)=\phi_{2}^{(a)}(V)$, confirming that

$$
h_{1}=\operatorname{rank}_{\mathbb{Q}}\left(\phi_{1}^{(a)}(V)\right)=1=\operatorname{rank}_{\mathbb{Q}}\left(\phi_{2}^{(a)}(V)\right)=h_{2} .
$$

Third and fourth, $\phi_{3}^{(a)}(V)=\mathbb{Q}\left(1, c_{1}, \ldots, c_{s} ; 0, \ldots, 0\right)=\phi_{4}^{(a)}(V)$, confirming that

$$
h_{3}=\operatorname{rank}_{\mathbb{Q}}\left(\phi_{3}^{(a)}(V)\right)=1=\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4}^{(a)}(V)\right)=h_{4} .
$$

For the next $j \in\{1, \ldots, s\}$, each $x \in \phi_{4+j}^{(a)}(V)$ has the form

$$
x=\left(\mu-\nu, \mu c_{j}-\nu c_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right)=(\mu-\nu)\left(1, c_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right)
$$

for some $(\mu, \nu) \in \mathbb{Q}^{2}$, thus $\phi_{4+j}^{(a)}(V)=\mathbb{Q}\left(1, c_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right)$, and so each $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+j}^{(a)}(V)\right)=$ $1=h_{4+j}$. For the next $j \in\{1, \ldots, k\}$, since $c_{i_{3, j}}=c_{i_{1, j}} c_{i_{2, j}}$, each $x \in \phi_{4+s+j}^{(a)}(V)$ has the form

$$
\begin{aligned}
x & =\left(\mu c_{i_{1, j}}+\nu, \mu c_{i_{3, j}}+\nu c_{i_{2, j}}, 0, \ldots, 0 ; 0, \ldots, 0\right) \\
& =\left(\mu c_{i_{1, j}}+\nu, \mu c_{i_{1, j}} c_{i_{2, j}}+\nu c_{i_{2, j}}, 0, \ldots, 0 ; 0, \ldots, 0\right) \\
& =\left(\mu c_{i_{1, j}}+\nu\right)\left(1, c_{i_{2, j}}, 0, \ldots, 0 ; 0, \ldots, 0\right),
\end{aligned}
$$

so as before $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+s+j}^{(a)}(V)\right)=1=h_{4+s+j}$. For the next $j \in\{1, \ldots, p\}$, i.e., the components which depend on $a=\left(\alpha_{1} / \beta_{1}, \ldots, \alpha_{p} / \beta_{p}\right)$, each $x \in \phi_{4+s+k+j}^{(a)}(V)$ has the form

$$
x=\left(\alpha_{j} \mu-\beta_{j} \mu a_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right)=(0, \ldots, 0 ; 0, \ldots, 0),
$$

confirming that each $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+s+k+j}^{(a)}(V)\right)=0=h_{4+s+k+j}$. Finally, for the last $j \in\{1, \ldots, l\}$, each $x \in \phi_{4+s+k+p+j}^{(a)}(V)$ has the form

$$
x=\left(\lambda_{0, j} \mu+\sum_{i} \lambda_{i, j} \mu c_{i}, 0, \ldots, 0 ; 0, \ldots, 0\right)=(0, \ldots, 0 ; 0, \ldots, 0),
$$

confirming that each $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+s+k+p+j}^{(a)}(V)\right)=0=h_{4+s+k+p+j}$. Thus we conclude that $a \in D(p, q, S)$ implies $\left(\phi_{j}^{(a)}\right)_{j} \in E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m, r}$.

To show necessity, suppose that we are given $a=\left(\alpha_{1} / \beta_{1}, \ldots, \alpha_{p} / \beta_{p}\right) \in \mathbb{Q}^{p}$ (in lowest terms) such that $\left(\phi_{j}^{(a)}\right)_{j} \in E_{h,\left(h_{j}\right)_{j}}^{\mathbb{Q}, m, r}$ with the same parameters $m, r, h, h_{j}$; we will show there exists $b \in \mathbb{Q}^{q}$ such that $(a, b) \in Z(S)$. Pick any $V \leq_{\mathbb{Q}} \mathbb{Q}^{r}$ such that $\operatorname{rank}_{\mathbb{Q}}(V)=h$ and each $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{j}^{(a)}(V)\right)=h_{j}$.

We first show that there exist elements $f, g \in V$ of the forms

$$
f=\left(f_{0}, \ldots, f_{s} ; 0, \ldots, 0\right) \quad g=\left(0, \ldots, 0 ; g_{0}, \ldots, g_{s}\right),
$$

such that $f_{0}=g_{0}=1$ and $\mathbb{Q} f+\mathbb{Q} g=V$. Let $d=\left(d_{0}, \ldots, d_{s} ; d_{0}^{\prime}, \ldots, d_{s}^{\prime}\right)$ and $e=$ $\left(e_{0}, \ldots, e_{s} ; e_{0}^{\prime}, \ldots, e_{s}^{\prime}\right)$ be a $\mathbb{Q}$-basis of $V$. Since $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{1}^{(a)}(V)\right)=h_{1}=1$, we may suppose that $d_{0}=1$ : otherwise, if $0 \neq d_{0} \neq 1$ then we can scale $d$ so that $d_{0}=1$, except if $d_{0}=0$, in which case necessarily $e_{0} \neq 0$, thus we can swap $d$ and $e$ and repeat this argument. Since $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{3}^{(a)}(V)\right)=h_{3}=1$, there exists $\gamma \in \mathbb{Q}$ such that $\gamma\left(d_{0}, \ldots, d_{s}\right)=\left(e_{0}, \ldots, e_{s}\right)$. Let $g^{\prime}=e-\gamma d=\left(0, \ldots, 0 ; g_{0}^{\prime}, \ldots, g_{s}^{\prime}\right)$. Since $g^{\prime}$ and $d$ are $\mathbb{Q}$-linearly independent elements of $V$ and $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4}^{(a)}(V)\right)=h_{4}=1$, there exists $\delta \in \mathbb{Q}$ such that $\delta\left(g_{0}^{\prime}, \ldots, g_{s}^{\prime}\right)=\left(d_{0}^{\prime}, \ldots, d_{s}^{\prime}\right)$. Now define $f=d-\delta g^{\prime}=\left(f_{0}, \ldots, f_{s} ; 0, \ldots, 0\right)$; it has the desired form. Noting that $f=-\delta e+(1+\delta \gamma) d), g^{\prime}$ and $f$ are $\mathbb{Q}$-linearly independent elements of $V$, and since $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{2}^{(a)}(V)\right)=h_{2}=1$, it follows that $g_{0}^{\prime} \neq 0$, so we can define $g=\left(1 / g_{0}^{\prime}\right) g^{\prime}$, which also has the desired form. Additionally, $f_{0}=g_{0}=1$ and $\mathbb{Q} f+\mathbb{Q} g=V$, as desired.

It follows that each element of $V$ can be written as $\mu f+\nu g$ for some $(\mu, \nu) \in \mathbb{Q}^{2}$. For each $j \in\{1, \ldots, s\}$, each $x \in \phi_{4+j}^{(a)}(V)$ has the form
$x=\left(\mu f_{0}-\nu g_{0}, \mu f_{j}-\nu g_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right)=\left((\mu-\nu) 1,(\mu-\nu) f_{j}+\nu\left(f_{j}-g_{j}\right), 0, \ldots, 0 ; 0, \ldots, 0\right) ;$
if $f_{j} \neq g_{j}$, then $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+j}^{(a)}(V)\right)=2 \neq h_{4+j}=1$, a contradiction. Therefore, $f_{j}=g_{j}$ for all $j \in\{1, \ldots, s\}$, as well as $f_{0}=g_{0}$ as shown earlier.

For each $j \in\{1, \ldots, p\}$, given $a=\left(\alpha_{1} / \beta_{1}, \ldots, \alpha_{p} / \beta_{p}\right)$, we have that, for any $(\mu, \nu) \in \mathbb{Q}^{2}$,

$$
\phi_{4+s+k+j}^{(a)}(\mu f+\nu g)=\left(\alpha_{j} \mu f_{0}-\beta_{j} \mu f_{j}, 0, \ldots, 0 ; 0, \ldots, 0\right)
$$

since $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+s+k+j}^{(a)}(V)\right)=h_{4+s+k+j}=0$, it follows that $\beta_{j} \mu f_{j}=\alpha_{j} \mu f_{0}=\alpha_{j} \mu$; now picking $(\mu, \nu) \in \mathbb{Q}^{2}$ such that $\mu \neq 0$, we have that $f_{j}=a_{j}$ for all $j \in\{1, \ldots, p\}$.

Now consider any $F \in S$. If $F=\lambda_{0, j}+\sum_{i} \lambda_{i, j} z_{i}$ for some $j \in\{1, \ldots, l\}$, then because $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+s+k+p+j}^{(a)}(V)\right)=h_{4+s+k+p+j}=0$, then $\lambda_{0, j}+\sum_{i} \lambda_{i, j} f_{i}=0$; that is, $F\left(f_{i}\right)_{i}=0$. Otherwise, $F=z_{i_{1, j}} z_{i_{2, j}}-z_{i_{3, j}}$ for some $j \in\{1, \ldots, k\}$; because $\operatorname{rank}_{\mathbb{Q}}\left(\phi_{4+s+j}^{(a)}(V)\right)=h_{4+s+j}=$ 1 , then there exists $\gamma \in \mathbb{Q}$ such that $\gamma\left(\mu f_{i_{1, j}}+\nu\right)=\mu f_{i_{3, j}}+\nu f_{i_{2, j}}$ for any $(\mu, \nu) \in \mathbb{Q}^{2}$. Considering the case $\nu=0$ and $\mu=1$, it follows that if $f_{i_{1, j}}=0$ then $f_{i_{3, j}}=0$; in this case, $f_{i_{1, j}} f_{i_{2, j}}=f_{i_{3, j}}$. Otherwise, $f_{i_{1, j}}, f_{i_{3, j}} \neq 0$, and we have $\gamma=f_{i_{3, j}} / f_{i_{1, j}}$, which we substitute to obtain $f_{i_{3, j}} / f_{i_{1, j}}=f_{i_{2, j}}$. That is, again, $f_{i_{1, j}} f_{i_{2, j}}=f_{i_{3, j}}$, so $F\left(f_{i}\right)_{i}=0$.

In conclusion, letting $b=\left(f_{p+1}, \ldots, f_{s}\right)$, we have shown that $\left(f_{1}, \ldots, f_{s}\right)=(a, b) \in$ $Z(S)$.

Proof of Theorem 4.4. Together, Lemmas 4.52, 4.53, and 4.54 complete the proof of Theorem 4.4. As in the case of Theorem 4.3, the statement of Theorem4.4 tacitly assumes a particular representation of $\mathcal{G}$.

## Chapter 5

## Communication Bounds for Loop Nests

In this chapter we apply the communication lower bounds framework developed in Section 2.3 to the class of algorithms characterized in Section 3.2, and study the attainability of these bounds. A section-by-section outline is as follows.

- In Section 5.1 we prove Theorem 3.1 (stated in Section 3.2), which yields communication lower bounds via Corollary 2.1 (in Section 2.3).
- In Section 5.2, we discuss attainability of the lower bounds, proposing blocked implementations as promising candidates. In Theorem 5.1, we suggest a class of tiles which in many cases attain equality in the conclusion (4.2) of Theorem 4.1 (in Chapter 4), and can be used to derive blocked implementations in many important cases.
- In Section 5.3, we discuss the infeasible case, the class of HBL interpretations where the hypothesis $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$ of Theorem 3.1 fails. In this case, the linear program (3.2) defining the parameter $\sigma$ is infeasible; this turns out to be promising from the standpoint of minimizing communication, because arbitrary data reuse may be possible.
- In Section 5.4. we discuss the injective case, the class of HBL interpretations where at least one homomorphism $\phi_{j}$ is an injection. In this case, we always have the parameter $\sigma=1$, its minimum value. This case turns out to be the least promising from the standpoint of minimizing communication, because no asymptotic data reuse is possible.
- In Section 5.5, we discuss the product case, mentioned in the context of Lemma 4.30 (in Section 4.4.5). This models nested-loop programs whose array subscripts are subsets of the loop indices, e.g., the matrix multiplication example at the beginning of Chapter 3. which has loop indices $i, j, k$ and array subscripts $(i, k),(k, j),(i, j)$, corresponding to the array elements $A(i, k), B(k, j), C(i, j)$. In the product case, communication lower bounds can be computed by simpler means than using the algorithm in Theorem 4.3 (stated in Section 4.2); moreover, Theorem 5.1 can be specialized to yield a stronger result (see Corollary 5.1). In Sections 5.5.1 5.5 .5 we treat several examples of
the product case including dot products, matrix-vector multiplication, matrix-matrix multiplication, tensor operations, and many-body calculations.


### 5.1 Lower Bounds for Loop Nests

Recall that Theorem 3.1(in Section 3.2) gives bounds on input-/output-path cutsizes $Q_{\text {in }}, Q_{\text {out }}$ for algorithms with HBL interpretations. In particular, each HBL interpretation defines a real number $\sigma \geq 1$ such that $Q_{\text {in }}(J)+Q_{\text {out }}(J) \geq|J|^{1 / \sigma}$ for all nonempty subsets $J$ of some subset $K$ of the algorithm. We can then apply Corollary 2.1 to derive communication lower bounds, as shown following the statement of Theorem 3.1. That is, given an algorithm with an HBL interpretation, in any implementation of that algorithm where some processor $p$, with $M_{p}$ local memory cells, is assigned the $N=|K|$ Operations corresponding to the evaluations $K \subseteq F$, if $N \geq\left(4 M_{p}\right)^{\sigma}$, then by Corollary 2.1, processor $p$ performs at least

$$
\frac{1}{2 \cdot 3^{\sigma}} \cdot \frac{N}{M_{p}^{\sigma-1}}=\Omega\left(\frac{N}{M_{p}^{\sigma-1}}\right)
$$

Loads and Stores between their first and last assigned Operation corresponding to $K$.
We now prove Theorem 3.1. Recall that we are given the HBL datum $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)$, and recall the definition of $\mathcal{P}(\mathcal{G})$ from (4.8).

Proof of Theorem 3.1. Theorem 4.1 tells us that if $s \in \mathcal{P}(\mathcal{G})$, then for all nonempty finite $E \subseteq \mathbb{Z}^{r},|E| \leq \prod_{j}\left|\phi_{j}(E)\right|^{s_{j}}$. Therefore, since $\mathcal{P}(\mathcal{G})$ is a closed subset of the first orthant of $\mathbb{R}^{m}$, and since $\mathcal{P}(\mathcal{G}) \neq \emptyset$ via Lemma 4.10 and the hypothesis $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$, we can take a minimum over $s \in \mathcal{P}(\mathcal{G})$,

$$
\begin{aligned}
|E| & \leq \min _{s \in \mathcal{P}(\mathcal{G})} \prod_{j}\left|\phi_{j}(E)\right|^{s_{j}} \\
& \leq \min _{s \in \mathcal{P}(\mathcal{G})} \prod_{j}\left(\max _{i=1}^{m}\left|\phi_{i}(E)\right|\right)^{s_{j}} \\
& =\min _{s \in \mathcal{P}(\mathcal{G})}\left(\max _{j}\left|\phi_{j}(E)\right|\right)^{\sum_{j} s_{j}} \\
& =\left(\max _{j}\left|\phi_{j}(E)\right|\right)^{\sigma},
\end{aligned}
$$

where

$$
\begin{equation*}
\sigma=\sigma(\mathcal{G})=\min _{s \in \mathcal{P}(\mathcal{G})} \sum_{j} s_{j} . \tag{5.1}
\end{equation*}
$$

We have that $\sigma \geq 1$ by Lemma 4.12 (by hypothesis, $1 \leq r=\operatorname{rank}\left(\mathbb{Z}^{r}\right)$ ); additionally, since $s \in$ $\mathcal{P}(\mathcal{G})$ implies $\left(\min \left(1, s_{j}\right)\right)_{j} \in \mathcal{P}(\mathcal{G})$ by Lemma 4.27. $\min _{s \in \mathcal{P}(\mathcal{G})} \sum_{j} s_{j}=\min _{s \in \mathcal{P}(\mathcal{G}) \cap[0,1]^{m}} \sum_{j} s_{j}$, so $\sigma \leq \sum_{j} 1=m$.

Having established the existence of $\sigma \in[1, m]$ such that $\max _{j}\left|\phi_{j}(E)\right| \geq|E|^{1 / \sigma}$ for all nonempty finite $E \subseteq \mathbb{Z}^{r}$, consider the case $E=Z(J)$ for any nonempty $J \subseteq K$. By definition, $A(J)=\bigcup_{j} A_{j}\left(\phi_{j}(Z(J))\right)$, and for each $j,\left|A_{j}\left(\phi_{j}(Z(J))\right)\right|=\left|\phi_{j}(Z(J))\right|$ by injectivity of $A_{j}$. So

$$
|A(J)|=\left|\bigcup_{j} A_{j}\left(\phi_{j}(Z(J))\right)\right| \geq \max _{j}\left|\phi_{j}(Z(J))\right| \geq|Z(J)|^{1 / \sigma}=|J|^{1 / \sigma}
$$

where the last equality follows from injectivity of $Z$. The conclusion follows from the fact that $Q_{\text {in }}(J)+Q_{\text {out }}(J) \geq|A(J)|$, noted before the statement of Theorem 3.1.

As remarked following the statement of Theorem 3.1, the exponent $\sigma$ is the solution of a linear program, $\min _{s \in \mathcal{P}(\mathcal{G})} \sum_{j} s_{j}$. We have studied the linear inequalities defining $\mathcal{P}(\mathcal{G})$ in depth in Sections 4.4 and 4.7. In particular, since $G=\mathbb{Z}^{r}$ is finitely generated, we can apply Theorem 4.3 to compute a finite set of $\mathbb{Z}$-linear inequalities which define $\mathcal{P}(\mathcal{G})$, and then (computably) solve this linear program for $\sigma$. Moreover, note that $\sigma \in \mathbb{Q}$ : the extreme points of $\mathcal{P}(\mathcal{G})$ are $\mathbb{Q}$-valued and the hyperplane defined by $\sum_{j} s_{j}=\sigma$ must intersect at least one extreme point of $\mathcal{P}(\mathcal{G})$.

We mention several ways in which the computation of $\sigma(\mathcal{G})$ can be simplified. First, by Lemma 4.14, if there exist distinct indices $k, l \in\{1, \ldots, m\}$ such that $\operatorname{rank}\left(\operatorname{ker}\left(\phi_{k}\right) \cap\right.$ $\left.\operatorname{ker}\left(\phi_{l}\right)\right)=\operatorname{rank}\left(\operatorname{ker}\left(\phi_{k}\right)\right)=\operatorname{rank}\left(\operatorname{ker}\left(\phi_{l}\right)\right)$, then letting $\mathcal{G}^{\prime}=\left(G,\left(G_{j}\right)_{j \neq l},\left(\phi_{j}\right)_{j \neq l}\right)$, we have that $\sigma(\mathcal{G})=\sigma\left(\mathcal{G}^{\prime}\right)$ : that is, dropping 'duplicate' homomorphisms does not affect the communication lower bound. Second, by Lemma 4.15, if there exists an index $k \in\{1, \ldots, m\}$ such that $\operatorname{rank}\left(\phi_{k}(G)\right)=0$, letting $\mathcal{G}^{\prime}=\left(G,\left(G_{j}\right)_{j \neq k},\left(\phi_{j}\right)_{j \neq k}\right)$, we have have $\sigma(\mathcal{G})=\sigma\left(\mathcal{G}^{\prime}\right)$ : that is, dropping rank-0 homomorphisms does not affect the communication lower bound. Third, by Lemma 4.16, if there exists an index $k \in\{1, \ldots, m\}$ such that $\operatorname{rank}\left(\phi_{k}(G)\right)=\operatorname{rank}(G)$, then $e_{k} \in \mathcal{P}(\mathcal{G}) ;$ since $\operatorname{rank}(G) \geq 1$, by Lemma 4.12 we conclude that $\sigma(\mathcal{G})=1$ : that is, we need only identify a single corank-0 homomorphism to identify the communication lower bound. Fourth, if the hypotheses of Lemma 4.30 apply, then the set $\mathcal{P}(\mathcal{G})$ has a simpler definition, which we discuss below in Section 5.5.

### 5.2 Upper Bounds for Loop Nests

We now turn to the question of the attainability of our communication lower bounds in Section 2.3: given an algorithm with an HBL interpretation, can we find an implementation with communication cost within a constant factor of the lower bound? This question is imprecise because there are actually multiple communication lower bounds in Section 2.3 , and their derivations involved a number of steps that (possibly) weakened the ultimate results; in order to clarify what we are attaining, we first review the main lower bound results of Section 2.3,

The main lower bound results of Section 2.3 - Theorem 2.1 and its special case Corollary 2.1 - are stated for a fixed algorithm with parameters $V, F, x, y, I, O$ (see Section 2.2 for definitions). Both results hypothesize that there exists a nonempty finite $K \subseteq F$ such that, for all nonempty $J \subseteq K$, the sum $Q_{\text {in }}(J)+Q_{\text {out }}(J)$ of the input- and output-path cutsizes of $J$ is bounded below by a nondecreasing function $q(|J|)$. Both results then conclude that in any implementation of this algorithm in which some processor $p$ is assigned the Operations corresponding to $K$, then processor $p$ must perform at least a certain number of Loads and Stores.

We now define a special class of implementations, called blocked implementations, whose definition is motivated by the segmentation argument in the proof of Theorem 2.1. Consider any algorithm with parameters $V, F, x, y, I, O$, and any of its implementations. Identifying $F$ with the corresponding Operations, the assignment of instructions to processors induces
a disjoint cover $F=\bigcup_{p} F_{p}$, and for each $p, F_{p}$ induces a chain in the execution order. We say that the given implementation is blocked it, for each $p$, we can write the chain $F_{p}=B_{p, 1} \cdots B_{p, b_{p}}$ as the concatenation of some $b_{p} \in \mathbb{N}$ nonempty contiguous subchains $B_{p, i}$, called blocks, such that processor $p$ 's chain of instructions is a repetition, for each $i \in\left[b_{p}\right]$, of the three-phase instruction sequence,

$$
\begin{aligned}
\text { Loads, } & \text { for each } v \in \mathcal{X}_{p, i}=X\left(B_{p, i}\right) \backslash Y\left(B_{p, i}\right) \\
\text { Operations, } & \text { for each } f \in B_{p, i} \\
\text { Stores, } & \text { for each } v \in \mathcal{Y}_{p, i}=Y\left(B_{p, i}\right) \cap\left(O \cup \underset{(q, j) \neq(p, i)}{\bigcup} X\left(B_{q, j}\right)\right) .
\end{aligned}
$$

Here we have also introduced the notations $X(J)=\left\{x_{f, 1}, \ldots, x_{f, m_{f}} \mid f \in J\right\}$ and $Y(J)=$ $\left\{y_{f, 1}, \ldots, y_{f, n_{f}} \mid f \in J\right\}$ for any $J \subseteq F$. It follows that any blocked implementation does not perform 'extra' Operations, i.e., there is a bijection between $F$ and the set of Operations in the execution. The key property of a blocked implementation is that the Operations in each block $B_{p, i}$ are performed without intervening Loads and Stores. It follows that, for each block $B_{p, i},\left|\mathcal{X}_{p, i}\right| \leq M_{p}$ and $\left|\mathcal{Y}_{p, i}\right| \leq M_{p}$. Therefore, each processor $p$ 's communication cost is bounded above,

$$
W_{p} \leq 2 M_{p} b_{p}
$$

Note that this upper bound is reminiscent of the lower bound (2.1) (Section 2.3) that we are trying to attain (in the case $l=2 M_{p}$ ). Our approach to deriving communication-reducing implementations, given an algorithm and machine parameters $P$ and $\left(M_{p}\right)_{p}$, is to construct a blocked implementation that attempts to minimize

$$
\max _{p} 2 M_{p} b_{p} \geq \max _{p} W_{p}
$$

We do not claim that blocked implementations minimize $\max _{p} W_{p}$ : for example, we are neglecting the possibility of inter-block data reuse in each processor's local memory (some Loads/Stores could be avoided).

We are particularly interested in the attainability of the (weaker) conclusion of Corollary 2.1. Consider any implementation that minimizes $\max _{p} W_{p}$, and look at any processor $p$ with maximal $W_{p}$. Corollary 2.1 hypothesizes that there exists a nonempty $K \subseteq F_{p}$ and a $\sigma \in[1, \infty)$ such that $|J|^{1 / \sigma} \leq Q_{\text {in }}(J)+Q_{\text {out }}(J)$ for all nonempty $J \subseteq K$. Additionally, if $N=|K| \geq\left(4 M_{p}\right)^{\sigma}$, then Corollary 2.1 gives the (nontrivial) lower bound

$$
W_{p} \geq \frac{1}{2 \cdot 3^{\sigma}} \frac{N}{M_{p}^{\sigma-1}}=2 M_{p} \cdot \frac{N}{4\left(3 M_{p}\right)^{\sigma}} .
$$

This motivates us to seek a blocked implementation where $b_{p} \leq C N / M_{p}^{\sigma}$ where $C$ is independent of the machine parameters $P$ and $\left(M_{p}\right)_{p}$, as well as $N ; C$ will, however, depend on another aspect of the HBL interpretation, its induced HBL datum $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right),\left(\phi_{j}\right)_{j}\right)$.

So far, we have not developed a general framework for deriving blocked implementations (for algorithms with or without HBL interpretations). The examples treated later in this chapter have dependence structures which enable a simpler, geometry-inspired blocking approach which we call tiling. For example, as mentioned in Chapter 3, multiplying
$N$-by- $N$ matrices, $C=A \cdot B$, can be programmed with triply nested loops (assuming $C$ is zero-initialized),

$$
\begin{gathered}
\text { for } i=1: N, \quad \text { for } j=1: N, \quad \text { for } k=1: N, \\
C(i, j)=C(i, j)+A(i, k) * B(k, j) .
\end{gathered}
$$

A blocked version of this code, where the blocking factor $b \in \mathbb{N}_{\geq 1}$ is assumed to divide $N$, is given by

$$
\begin{aligned}
& \text { for } i_{1}=1: N / b, \quad \text { for } j_{1}=1: N / b, \quad \text { for } k_{1}=1: N / b, \\
& \quad \text { for } i_{2}=1: b, \quad \text { for } j_{2}=1: b, \quad \text { for } k_{2}=1: b, \\
& \quad / / \text { reindexing }(i, j, k)=b \cdot\left(i_{1}-1, j_{1}-1, k_{1}-1\right)+\left(i_{2}, j_{2}, k_{2}\right) \\
& \quad C(i, j)=C(i, j)+A(i, k) * B(k, j) .
\end{aligned}
$$

The inner three loops (indexed by $i_{2}, j_{2}, k_{2}$ ) represent blocks, of which there are $(N / b)^{3}$, traversed by the outer three loops (indexed by $i_{1}, j_{1}, k_{1}$ ). Each block requires $3 b^{2}$ local memory cells to execute: the submatrices of $A, B, C, b^{2}$ entries each, must be loaded, and the submatrix of $C$ must be stored, but each updated entry of $C$ can overwrite its previous value. Therefore, we can define a blocked implementation by picking $b$ such that $3 b^{2} \leq M_{p}$ for some processor $p$, and assigning all the work to that processor $p$. A parallel implementation can be defined, e.g., by replacing $M_{p}$ by $\min _{p} M_{p}$, and assigning blocks to processors, ensuring that the execution order preserves the partial order on blocks, in particular, when multiple processors are assigned blocks that update the same submatrix of $C$. (Additionally, more care must be taken when trying to minimize the critical-path communication cost, rather than the per-processor communication costs.) We will revisit this example in Section 5.5.3 more formally (in terms of algorithms), to illustrate the algorithmic dependence structure that permits this code transformation. To foreshadow this discussion, we observe that while the unblocked example suggests a lexicographical order on $(i, j, k) \in[N]^{3}$, the only interiteration dependences are due to the summation order, which impose a lexicographical order on the set $\{(i, j, k) \mid k \in[N]\}$ for each $(i, j) \in[N]^{2}$, i.e., a partial order on $[N]^{3}$, and the blocked code respects this partial order.

This matrix multiplication example demonstrates tiling $\mathbb{Z}^{3}$ (in particular, the cube $[N]^{3}$ ) by translates of the cube $[b]^{3}$. In the remainder of this section, we develop a tool (see Theorem 5.1), to assist in finding tilings, which can use more general parallelotopes than cubes like $[b]^{3}$. This development is motivated by attainability of the upper bounds given by Theorem 4.1 (Section 5.2). That is, given some $s \in \mathcal{P}(\mathcal{G})$, does there exist a nonempty finite $E \subseteq \mathbb{Z}^{r}$ such that $|E|=\prod_{j}\left|\phi_{j}(E)\right|^{s_{j}}$ ? Actually, recalling the invocation of Theorem 4.1 within the proof of Theorem 3.1, we are really concerned with whether $|E|=\max _{j}\left|\phi_{j}(E)\right|^{\sigma}$, where $\sigma=\sigma(\mathcal{G})=\min _{s \in \mathcal{P}(\mathcal{G})} \sum_{j} s_{j} \in[1, m]$ and we assume $r \geq 1$. Of course, either of these equalities are satisfied by taking $E$ to be any singleton - so to get a meaningful answer, we must impose additional constraints.

Let's borrow some notation from Chapter 4; given the HBL datum $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)$ where $r \in\{1,2, \ldots\}$, we let $\mathcal{H}^{*}$ denote the set of all nontrivial subgroups of $\mathbb{Z}^{r}$ and let $\mathcal{E}^{*}$ denote the set of all nonempty finite subsets of $\mathbb{Z}^{r}$. For any $\mathcal{H} \subseteq \mathcal{H}^{*}$, we define $\mathcal{P}(\mathcal{G}, \mathcal{H})$ according to (4.8) in Section 4.4, with the special case $\mathcal{P}(\mathcal{G})=\mathcal{P}\left(\mathcal{G}, \mathcal{H}^{*}\right)$. We also extend
$\sigma=\sigma(\mathcal{G})$ to the more general notation, continuing to assume $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$,

$$
\begin{equation*}
\sigma(\mathcal{G}, \mathcal{H})=\min _{s \in \mathcal{P}(\mathcal{G}, \mathcal{H})} \sum_{j} s_{j}, \tag{5.2}
\end{equation*}
$$

and say $\sigma(\mathcal{G})$ to mean $\sigma\left(\mathcal{G}, \mathcal{H}^{*}\right)$ from now on; as in the case of $\sigma(\mathcal{G}), \sigma(\mathcal{G}, \mathcal{H}) \in[1, \sigma(\mathcal{G})] \subseteq$ $[1, m]$ is well defined because $\mathcal{P}(\mathcal{G}, \mathcal{H})$ is a closed subset of the first orthant of $\mathbb{R}^{m}$.

A nonempty finite subset $\mathcal{H}=\left\{H_{1}, \ldots, H_{k}\right\}$ of $\mathcal{H}^{*}$ is independent if the subgroups' sum is isomorphic to their direct sum, i.e., $H=\sum_{i=1}^{k} H_{i} \cong \bigoplus_{i=1}^{k} H_{i}$, or, equivalently, if we can find a basis $B$ (over $\mathbb{Z}$ ) of $H$ that can be partitioned as $\left\{B_{1}, \ldots, B_{k}\right\}$, where for each $i \in\{1, \ldots, k\}$, $B_{i}$ is a basis of $H_{i}$.

Theorem 5.1. If $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$ and $\mathcal{H} \subseteq \mathcal{H}^{*}$ is independent, then there exists $\mu \in$ $\{1,2, \ldots\}, c \in(0, \infty), \tau \in[\sigma(\mathcal{G}, \mathcal{H}), \sigma(\mathcal{G})]$, and a sequence $\left(E_{1}, E_{2}, \ldots\right)$ over $\mathcal{E}^{*}$ such that, for all $N \in\{\mu, \mu+1, \ldots\}$,

Property (1) $\quad \sum_{j}\left|\phi_{j}\left(E_{N}\right)\right| \leq N$,
Property (2) $\left|E_{N}\right| \geq c N^{\tau}$, and
Property (3) $\mathbb{Z}^{r}$ is partitioned by translates of $E_{N}$.
Proof. We adopt a few notational conveniences. We reserve the index $i$ to range over $\{1, \ldots, k\}$, while $j$ continues to range over $\{1, \ldots, m\}$. We treat the Abelian group $\mathbb{Z}^{d}$ and its subgroups as (sub-) $\mathbb{Z}$-modules, defining (scalar) multiplication in terms of addition. In this case, we assume the standard basis $\left(e_{1}, \ldots, e_{d}\right)$ on $\mathbb{Z}^{d}$ and represent tuples over $\mathbb{Z}$ as column vectors. If $B$ is a column vector, or a set or matrix of column vectors, then $\langle B\rangle$ is the set of all linear combinations of $B$ over $\mathbb{Z}$. If $y$ is an $n$-element column vector, $X, X_{i}$ are sets of $n$-element column vectors, and $A$ is an $n$-column matrix, $y+X$ denotes the set $\{y+x \mid x \in X\}, \sum_{i} X_{i}$ denotes the set $\left\{\sum_{i} x_{i} \mid x \in X_{i} X_{i}\right\}$, and $A \cdot X$ denotes the set $\{A x \mid x \in X\}$.

For each $n \in\{1,2, \ldots\}, h \in\{1, \ldots, r\}$, rank- $h$ matrix $B=\left[b_{1}, \ldots, b_{h}\right] \in \mathbb{Z}^{r \times h}$, and $t \in \mathbb{Z}^{r}$, we define the cube

$$
C(n, h, B, t)=t+B \cdot\{0, \ldots, n-1\}^{h} .
$$

Suppose we have $k$ cubes $C_{i}=C\left(n_{i}, h_{i}, B_{i}, t_{i}\right)$ and we define $h=\sum_{i} h_{i}, B=\left[B_{1}, \ldots, B_{k}\right]$, and $t=\sum_{i} t_{i}$. If $B$ is rank- $h$, then we define the tile

$$
D\left(C_{1}, \ldots, C_{k}\right)=t+B \cdot \underset{i}{X}\left\{0, \ldots, n_{i}-1\right\}^{h_{i}}
$$

For each $i$, let $h_{i}=\operatorname{rank}\left(H_{i}\right)$ and let $B_{i}$ be an arbitrary basis of $H_{i}$ represented as a matrix of column vectors, defining the family of cubes $C\left(N, h_{i}, B_{i}, 0\right)$ for $N \in\{1,2, \ldots\}$. Adapting the argument in the proof of Lemma 4.48 (Section 4.6), for each $j$, there exists $A_{i, j} \in[1, \infty)$ such that $\left|\phi_{j}\left(C\left(N, h_{i}, B_{i}, 0\right)\right)\right| \leq A_{i, j} N^{\operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)}$ for all $N$; in case $\operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)=0$ we set $A_{i, j}=1$. Let $x=\left(x_{1}, \ldots, x_{k}\right)$ be any optimal solution of the linear program

$$
\begin{equation*}
\max \left\{\sum_{i} x_{i} \operatorname{rank}\left(H_{i}\right) \mid x \in[0, \infty)^{k} \wedge(\forall j) 1 \geq \sum_{i} x_{i} \operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)\right\} \tag{5.3}
\end{equation*}
$$

which is the dual of the linear program (5.2) for the given $\mathcal{H}$, whose optimal value $\sigma(\mathcal{G}, \mathcal{H})$ we have already confirmed exists; therefore, we know that $\sum_{i} x_{i} \operatorname{rank}\left(H_{i}\right)=\sigma(\mathcal{G}, \mathcal{H})$. We can also show that any feasible dual solution $x \in[0,1]^{m}$ : supposing towards a contradiction that some $x_{i}>1$, it follows that $\operatorname{rank}\left(\phi_{j}\left(H_{i}\right)\right)=0$ for all $j$, so $H_{i} \leq \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)$, but $\operatorname{rank}\left(H_{i}\right)>$ 0 by hypothesis, contradicting the hypothesis that $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$. Additionally, since $\sigma(\mathcal{G}, \mathcal{H}) \geq 1$, at least one $x_{i}>0$. Consider each index $i$ : if $x_{i}=0$ then we set $n_{i}(N)=1$ and $\mu_{i, j}=1$ for all $j$. Otherwise, if $x_{i}>0$, consider each index $j$; let $\mu_{i}$ be the smallest $N \in$ $\{1,2, \ldots\}$ such that $n_{i}(N)=\left\lfloor N^{x_{i}} /\left(m \prod_{j} A_{i, j}\right)\right\rfloor>0 ; \mu_{i}$ is well defined since $N^{x_{i}} /\left(m \Pi_{j} A_{i, j}\right)$ is an increasing function of $N$; note that $\mu_{i} \geq m^{1 / x_{i}} \geq m$, since $x_{i} \in(0,1]$. Let $\mu=\max _{i} \mu_{i} \geq$ $m$.

For each $N \in\{1, \ldots, \mu-1\}$, let $E_{N}=\{0\}$. For each $N \in\{\mu, \mu+1, \ldots\}$, let $C_{i}(N)=$ $C\left(n_{i}(N), h_{i}, B_{i}, 0\right)$ for each $i$; by independence, $B=\left[B_{1}, \ldots, B_{k}\right]$ has rank $h=\sum_{i} h_{i}$ so we can define the tile $E_{N}=D\left(C_{1}(N), \ldots, C_{k}(N)\right)$. We now verify that the sequence of tiles $E_{1}, E_{2}, \ldots$ has the three desired properties.
Property (1). Consider any $N \in\{\mu, \mu+1, \ldots\}$. For each $j$,

$$
\begin{aligned}
\phi_{j}\left(E_{N}\right)=\phi_{j}\left(D\left(C_{1}(N), \ldots, C_{k}(N)\right)\right) & =\phi_{j}\left(B \cdot \underset{i}{X}\left\{0, \ldots, n_{i}(N)-1\right\}^{h_{i}}\right) \\
& =\sum_{i} \phi_{j}\left(B_{i} \cdot\left\{0, \ldots, n_{i}(N)-1\right\}^{h_{i}}\right) \\
& =\sum_{i} \phi_{j}\left(C_{i}(N)\right)
\end{aligned}
$$

Consider any $k \in\{1, \ldots, m\}$. Let $I=\left\{i \mid x_{i}>0 \wedge \operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)>0\right\}$; if $i \notin I$, then $\left|\phi_{k}\left(C_{i}(N)\right)\right|=1$. So, if $I=\emptyset$, then $\left|\phi_{k}\left(E_{N}\right)\right|=1$. Otherwise, $I \neq \emptyset$, then

$$
\begin{aligned}
\left|\phi_{k}\left(E_{N}\right)\right| & \leq \prod_{i}\left|\phi_{k}\left(C_{i}(N)\right)\right|=\prod_{i \in I}\left|\phi_{k}\left(C_{i}(N)\right)\right| \leq \prod_{i \in I} A_{i, k} n_{i}^{\operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)} \\
& =\prod_{i \in I} A_{i, k}\left|\frac{N^{x_{i}}}{m \prod_{j} A_{i, j}}\right|^{\operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)} \leq \prod_{i \in I} A_{i, k}\left(\frac{N^{x_{i}}}{m \prod_{j} A_{i, j}}\right)^{\operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)} \\
& =\left(\prod_{i \in I} \frac{A_{i, k}}{\left(m \prod_{j} A_{i, j}\right)^{\operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)}}\right) \prod_{i \in I} N^{x_{i} \operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)} \leq \frac{1}{m} \prod_{i \in I} N^{x_{i} \operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)} \\
& =\frac{1}{m} N^{\sum_{i} x_{i} \operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right)} \leq \frac{N^{1}}{m} .
\end{aligned}
$$

where we have exploited the facts that all $A_{i, j} \geq 1$, that $\operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right) \geq 1$ for all $i \in I$, and that $\sum_{i} x_{i} \operatorname{rank}\left(\phi_{k}\left(H_{i}\right)\right) \leq 1$ by the fact that $x$ is a feasible solution of (5.3).

Since $N \geq \mu \geq m$, we have thus shown that for all $j,\left|\phi_{j}\left(E_{N}\right)\right| \leq\lfloor N / m\rfloor$, therefore

$$
\sum_{j}\left|\phi_{j}\left(E_{N}\right)\right| \leq \sum_{j}\left\lfloor\frac{N}{m}\right\rfloor=m\left\lfloor\frac{N}{m}\right\rfloor \leq N
$$

as desired.

Property (2). Consider any $N \in\{\mu, \mu+1, \ldots\}$; by independence (of $\mathcal{H}$ and $B$ ) and definition of $\mu$,

$$
\left|E_{N}\right|=\prod_{i}\left\lfloor\frac{N^{x_{i}}}{m \prod_{j} A_{i, j}}\right]^{h_{i}} \geq \prod_{i}\left(\frac{N^{x_{i}}}{2 m \prod_{j} A_{i, j}}\right)^{h_{i}}=\left(\prod_{i} \frac{1}{\left(2 m \prod_{j} A_{i, j}\right)^{h_{i}}}\right) \prod_{i} N^{x_{i} h_{i}}
$$

letting $c=\Pi_{i}\left(2 m \prod_{j} A_{i, j}\right)^{-h_{i}} \in(0,1]$, we have that

$$
\left|E_{N}\right| \geq c \prod_{i} N^{x_{i} h_{i}}=c N^{\sum_{i} x_{i} \operatorname{rank}\left(H_{i}\right)}=c N^{\sigma(\mathcal{G}, \mathcal{H})} .
$$

We also have that

$$
\left|E_{N}\right| \leq\left(\max _{j}\left|\phi_{j}\left(E_{N}\right)\right|\right)^{\sigma(\mathcal{G})} \leq\left(\sum_{j}\left|\phi_{j}\left(E_{N}\right)\right|\right)^{\sigma(\mathcal{G})} \leq N^{\sigma(\mathcal{G})}
$$

Therefore, there exists $c \in(0, \infty)$ and $\tau \in[\sigma(\mathcal{G}, \mathcal{H}), \sigma(\mathcal{G})]$ such that, for all $N \in\{\mu, \mu+1, \ldots\}$, $\left|E_{N}\right| \geq c N^{\tau}$.
Property (3). For each $N \in\{1,2, \ldots\}$, the cubic structure of $E_{N}$ makes it straightforward to partition $\mathbb{Z}^{r}$ with its translates as follows. If $N<\mu$, then $E_{N}=\{0\}$ by construction and the conclusion is immediate: $\mathbb{Z}^{r}$ is partitioned by the set of all its singletons. Otherwise $N \geq \mu$, and for each $i$, the cube $C_{i}$ can be translated to partition $H_{i}$ by varying $t_{i}$ over the linear combinations of the columns of $B_{i}$ by all integer multiples of $n_{i}$. By partitioning each subgroup $H_{i}$ with translated cubes in this manner, we obtain a partition of $\bigoplus_{i} H_{i}$ and thus of $H=\sum_{i} H_{i}$ by isomorphism, and each part in this induced partition of $H$ is a tile, in particular, $t+E_{N}$ for some $t \in H$. The cosets of $H$ in $\mathbb{Z}^{r}$ (i.e., the elements of the quotient group $\left.\mathbb{Z}^{r} / H\right)$ partition $\mathbb{Z}^{r}$, and each coset $y+H$ is a translation of $H$ by some representative $y \in \mathbb{Z}^{r}$, so the partition of $H$ (i.e., the coset $0+H$ ) into tiles induces partitions into tiles of all the other cosets of $H$ by translation by $y$.

The constant $c$ can be improved in a number of ways. For example, we can weakening Property 1 to bound $\max _{j}\left|\phi_{j}\left(E_{N}\right)\right|$ instead, we can model some instances of aliasing arrays (see Section 3.2. Also, a few simplifications to the dual linear program (5.3) are possible. For example, the constraints corresponding to any rank-0 homomorphisms $\phi_{j}$ have no impact, and could be omitted. Other possible simplifications include combining duplicate homomorphisms $\phi_{k}=\phi_{l}$, as discussed in the preceding section. In the case of injective $\phi_{j}$, since $\sigma(\mathcal{G})=1$, a simpler special case of Theorem 5.1 applies, as discussed below in Section 5.4 .

At this point, we have defined blocked implementations, which enable a simpler communication upper bound analysis, as well as defined a tool, Theorem 5.1, for finding tilings of $\mathbb{Z}^{r}$ that have certain properties with respect to a given HBL datum. It remains future work to develop a general approach for using Theorem 5.1 to define blocked implementations of given algorithms with HBL interpretations. In the remainder of this chapter, we apply Theorem 5.1 to a sequence of examples, demonstrating its practical use.

### 5.3 Infeasible Case

Our first class of examples is called the infeasible case. These are algorithms with HBL interpretations where $\mathcal{P}\left(\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)\right)=\emptyset$, i.e., the linear program 5.1) is infeasible. By Lemma 4.10, this means that $\operatorname{rank}\left(\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)\right)>0$.

As an example of the infeasible case, recall the "path graph" example in Section 3.1.3 (see also Figure 3.3), meant to model a loop like

$$
\text { for } i=1: N, \quad a_{i}=f\left(a_{i-1}\right),
$$

which transform an input $a_{0}$ to an output $a_{N}$; a single processor can perform the corresponding sequence of Ops entirely in their local memory, without any communication other than (perhaps) initially Loading $a_{0}$. From the perspective of our lower-bound analysis, infeasibility is the best-case scenario for minimizing communication. We interpret the preceding code as a family of algorithms whose parameters satisfy

$$
\begin{equation*}
V=\left\{a_{0}, a_{i} \mid i \in[N]\right\}, \quad F=\left\{f_{i} \mid i \in[N]\right\}, \quad x:\left(f_{i}, 1\right) \mapsto a_{i-1}, \quad y:\left(f_{i}, 1\right) \mapsto a_{i} . \tag{5.4}
\end{equation*}
$$

Such algorithms admit the HBL interpretation

$$
\begin{gathered}
m=2, \quad r=1, \quad r_{1}=r_{2}=0, \quad K=F, \quad Z\left(f_{i}\right)=i, \\
\phi_{1}(i)=\phi_{2}(i)=0, \quad A_{1}(0)=a_{0}, \quad A_{2}(0)=a_{N}
\end{gathered}
$$

This induces the $\operatorname{HBL}$ datum $\mathcal{G}=\left(\mathbb{Z}^{1},\left(\mathbb{Z}^{0}\right)^{2},\left(\mathbb{Z}^{1} \rightarrow \mathbb{Z}^{0}\right)^{2}\right)$, and we immediately see that $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\mathbb{Z}^{1}$, implying infeasibility.

Suppose the primal linear program $\min \left\{\sum_{j} s_{j} \mid s \in \mathcal{P}(\mathcal{G})\right\}$ is infeasible. Since $x=$ $(0, \ldots, 0)$ is always a feasible solution of the dual linear program (5.3), the dual must be unbounded. This confirms the intuition in the preceding section, that we can increase the tile edge lengths (e.g., via the dual variables $\left\{x_{i}\right\}$ ) without bound. Note that for arbitrary $\mathcal{H} \subseteq \mathcal{H}^{*}$, possibly $\mathcal{P}(\mathcal{G}, \mathcal{H}) \neq \emptyset$ even if $\mathcal{P}(\mathcal{G})=\emptyset$. That is, the (generalized) primal linear program $\min \left\{\sum_{j} s_{j} \mid s \in \mathcal{P}(\mathcal{G}, \mathcal{H})\right\}$ may be feasible, and the primal/dual solution $\sigma(\mathcal{P}, \mathcal{H}) \in$ $[1, \sigma(\mathcal{P})]$ as otherwise expected. For example, $H \cap \bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$ for all $H \in \mathcal{H}$. If algorithm dependences do not permit a tiling that takes care of infeasibility in the manner discussed next, then it may be worthwhile to study such an $\mathcal{H}$.

Note than in the path graph example above, picking $E_{N}=\{0, \ldots, N-1\}$, we have that $\left|\phi_{1}\left(E_{N}\right)\right|=1 \leq N$ and $\left|E_{N}\right|=N=N^{1}$ for all $N \in\{1,2, \ldots\}$. In the case of more general infeasible HBL data, we can pick $E_{N}$ as a cube in $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)$ and apply similar reasoning.

### 5.4 Injective Case

Our next class of examples is called the injective case. These are algorithms with HBL interpretations which induce HBL data $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)$ where at least one homomorphism $\phi_{j}$ is an injection, and $\sigma(\mathcal{G})=1$.

As an example of the injective case, recall the example in Section 3.1.1 (see also Figure 3.1), meant to model a loop like

$$
\text { for } i=1: N, b_{i}=f\left(a_{i}\right),
$$

which transforms an input sequence $a_{1}, \ldots, a_{N}$ to an output sequence $b_{1}, \ldots, b_{N}$ in a componentwise fashion. If a processor performs sufficiently many of the corresponding Operations (at least their local memory size), then they must perform Loads/Stores to read inputs $a_{i}$ (which cannot all initially be in their local memory) or write outputs $b_{i}$ (which cannot be overwritten by definition of implementation). From the perspective of our lower-bound analysis, injectivity is the worst-case scenario for minimizing communication. We interpret the pr

Consider the HBL datum $\mathcal{G}=\left(\mathbb{Z}^{1},\left(\mathbb{Z}^{1}\right),\left(\phi_{1}: x \mapsto x\right)\right)$. We interpret the preceding code as a family of algorithms whose parameters satisfy

$$
\begin{equation*}
V=\left\{a_{i}, b_{i} \mid i \in[N]\right\}, \quad F=\left\{f_{i} \mid i \in[N]\right\}, \quad x:\left(f_{i}, 1\right) \mapsto a_{i}, \quad y:\left(f_{i}, 1\right) \mapsto b_{i} . \tag{5.5}
\end{equation*}
$$

Such algorithms admit the HBL interpretation

$$
\begin{gathered}
m=2, \quad r=1, \quad r_{1}=r_{2}=1, \quad K=F, \quad Z\left(f_{i}\right)=i, \\
\phi_{1}(i)=\phi_{2}(i)=i, \quad A_{1}(i)=a_{i}, \quad A_{2}(i)=b_{i}
\end{gathered}
$$

This induces the HBL datum $\mathcal{G}=\left(\mathbb{Z}^{1},\left(\mathbb{Z}^{1}\right)^{2},(x \mapsto x)^{2}\right)$, and we confirm that both $\phi_{1}$ and $\phi_{2}$ are injections (in particular, identity functions).

If an algorithm admits an HBL interpretation with $\sigma(\mathcal{G})=1$, it means that there is a set of disjoint input- and output-paths that connect at least one input or output argument variable of every $f \in K$. Conversely, if some nonempty $K \subseteq F$ has this property, then we can record a (distinct) $v(f) \in I \cup O$ for each $f \in K$ and define an HBL interpretation with $m=r=1, Z(K)=\{1, \ldots,|K|\}, \phi_{1}=\mathrm{id}_{\mathbb{Z}}$, and $A_{1}(i)=v\left(Z^{-1}(i)\right)$ for each $i \in\{1, \ldots,|K|\} ;$ we observe that $\sigma\left(\left(\mathbb{Z}, \mathbb{Z}, \mathrm{id}_{\mathbb{Z}}\right)\right)=1$.

From an asymptotic standpoint, we are done the moment we identify an injective homomorphism $\phi_{j}$. Picking $E_{N}=\{0, \ldots, N-1\}$, we have that $\left|\phi_{1}\left(E_{N}\right)\right|=N$ and $\left|E_{N}\right|=N=$ $N^{\sigma(\mathcal{G})}$. In more general injective HBL data, we can pick $E_{N}$ to be any size- $N$ tile in $\mathbb{Z}^{r}$, with obtaining similar results (possibly) with constants depending on $m$.

We will see two more examples of injectivity in Sections 5.5.1 and 5.5.2.

### 5.5 Product Case

Our main class of examples is called the product case, introduced in Lemma 4.30 (in Section 4.4.5). These are algorithms with HBL interpretations inducing HBL data $\mathcal{G}$ such that $\mathcal{P}(\mathcal{G})$ is defined by a simpler set of inequalities. We will consider applying Lemma 4.30 using the standard coordinate $\mathbb{Z}$-basis of $\mathbb{Z}^{r}$ : in this case, each $\phi_{j}$ "forgets" a subset of the coordinates $i_{1}, \ldots, i_{r}$. Consider, for example, the HBL datum

$$
\mathcal{G}=\left(\mathbb{Z}^{r},\left\{\mathbb{Z}^{r-1}\right\}^{m},\left(\phi_{j}:\left(i_{1}, \ldots, i_{r}\right) \mapsto\left(i_{1}, \ldots, i_{j-1}, i_{j+1}, \ldots, i_{r}\right)\right)_{j}\right)
$$

This is an HBL interpretation for the $r$-dimensional diamond DAG example in Section 3.1.5, revisited below in Section 5.5.5. We have found a number of practical applications of the product case, which we explore below.

Corollary 5.1. Consider an HBL datum $\mathcal{G}=\left(\mathbb{Z}^{r},\left(\mathbb{Z}^{r_{j}}\right)_{j},\left(\phi_{j}\right)_{j}\right)$, where each $\phi_{j}\left(i_{1}, \ldots, i_{r}\right)=$ $\left(i_{\pi_{j}(1)}, \ldots, i_{\pi_{j}\left(r_{j}\right)}\right)$ for some $\pi_{j}:\left\{1, \ldots, r_{j}\right\} \rightarrow\{1, \ldots, r\}$. Given $\mathcal{H}=\left\{\left\langle e_{1}\right\rangle, \ldots,\left\langle e_{r}\right\rangle\right\}$ and any optimal solution $x$ to the linear program (5.3), the conclusion of Theorem5.1 holds with $\mu=m^{\max \left\{1 / x_{i} \mid x_{i}>0\right\}}, c=(2 m)^{-r}$, and $\tau=\sigma(\mathcal{G}, \mathcal{H})$

Proof. To specialize Theorem 5.1 in this manner, all we need to show is that the constants $A_{i, j}$ appearing in the proof can all be made equal to one. To see this, we return to the proof of Lemma 4.48 (Section 4.6) where these constants originated, and exploit the fact that the matrix for each $\phi_{j}$ is $0 / 1$-valued with at most one nonzero in any row.

In the following, we go through a number of examples of the product case. We will use the following notation about the linear programs related to HBL data. For any HBL datum $\mathcal{G}$ and any subset $\mathcal{H}$ of the finite-rank subgroups of $G$, let $y_{H}=\operatorname{rank}(H)$ and $\Delta_{H, j}=\operatorname{rank}\left(\phi_{j}(H)\right)$, and assemble $y=\left(y_{H}\right)_{H}$ and $\Delta=\left(\Delta_{H, j}\right)_{H, j}$, where $y$ is a column vector and $\Delta$ is an $m$ column matrix both with (possibly infinitely many) rows indexed by $\mathcal{H}$. We also let $1_{m}$ denote the column vector of all ones of length $m$. In this notation, the primal and dual linear programs (5.2) 5.3) are the following.

\[

\]

In the present (product case) setting, $G=\mathbb{Z}^{r}$ and $\mathcal{H}=\left\{\left\langle e_{1}\right\rangle, \ldots,\left\langle e_{r}\right\rangle\right\}$, so $|\mathcal{H}|=r$ and $y=1_{r}$, the column vector of all ones of length $r$. Moreover, $\Delta$ is $0 / 1$-valued: $\Delta_{i, j}=1$ if $e_{i} \in \operatorname{ker}\left(\phi_{j}\right)$ and $\Delta_{i, j}=0$ otherwise.

Primal linear program (product case) Dual linear program (product case)

$$
\min _{s \in[0, \infty)^{m}} 1_{m}^{T} s \quad \text { s.t. } \quad \Delta s \geq 1_{r} \quad \max _{x \in[0, \infty)^{r}} 1_{r}^{T} x \quad \text { s.t. } \quad \Delta^{T} x \leq 1_{m}
$$

We give a number of examples (Sections 5.5.1 5.5.5) of the product case, including several from our earlier work [15]. Many of these examples are related to 'independent evaluations' studied by Hong-Kung [21, Section 6] and expanded in subsequent works [1, 33, 22, 5]. That is, communication lower bounds and, in many cases, matching upper bounds, are already known for several of these examples. Our contribution is to generalize this to more general loop nests.

### 5.5.1 Dot Product

We model computing sums-of-products, expressed mathematically as $\sum_{i} a_{i} b_{i}$, or in pseudocode as

$$
\begin{aligned}
& \text { for } i=1: N, \\
& \quad s=s+a_{i} * b_{i},
\end{aligned}
$$

by a family of algorithms whose parameters satisfy

$$
\begin{gather*}
I=\left\{a_{i}, b_{i} \mid i \in[N]\right\}, \quad O=\{t\}, \quad R=V \backslash I \supseteq\left\{s_{i} \mid i \in[N]\right\}, \quad F \supseteq\left\{z_{i} \mid i \in[N]\right\} \\
\left\{a_{i}, b_{i}\right\} \subseteq \operatorname{im}\left(x\left(z_{i}, \cdot\right)\right) \subseteq\left\{a_{i}, b_{i}\right\} \cup\left(R \backslash\left\{s_{i}\right\}\right), \quad\left\{s_{i}\right\} \subseteq \operatorname{im}\left(y\left(z_{i}, \cdot\right)\right) \subseteq R \backslash\left\{s_{k} \mid i \neq k \in[N]\right\} . \tag{5.6}
\end{gather*}
$$



Figure 5.1: Dot product example (Section 5.5.1) in the case $N=3$ : general form (top); using binary,$+ *$ operations (bottom-left); using "fused multiply-add" $*+$ operation (bottomright).
(Informally, the constraints on $\operatorname{im}\left(x\left(z_{i}, \cdot\right)\right)$ and $\operatorname{im}\left(y\left(z_{i}, \cdot\right)\right)$ mean that additions can be fused with multiplications, but multiplications cannot be fused.) The variable $t \in V$ is any element of $R$ (perhaps one of $s_{1}, \ldots, s_{N}$ ). Algorithms satisfying (5.6) induce DAGs of the form shown in Figure 5.1. Note that the summation is modeled by a more general reduction, introduced in Section 3.1.6. In particular, note that the operations $z_{i}$ are not necessarily independent; the important property is that they each depend on a distinct pair $\left\{a_{i}, b_{i}\right\}$ of inputs.

Algorithms satisfying (5.6) admit the following HBL interpretation.

$$
\begin{gathered}
m=3, \quad r=1, \quad\left(r_{1}, r_{2}, r_{3}\right)=(1,1,0) \\
K=\left\{z_{i} \mid i \in[N]\right\}, \quad Z\left(z_{i}\right)=i \\
\phi_{1}=\phi_{2}=\mathrm{id}_{\mathbb{Z}}, \quad \phi_{3}: \mathbb{Z} \rightarrow\{0\} \\
A_{1}(i)=a_{i}, \quad A_{2}(i)=b_{i}, \quad A_{3}(i)=t
\end{gathered}
$$

We confirm that $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$. We also recognize that $\phi_{1}$ and $\phi_{2}$ are both injections, so the conclusions in Section 5.4 apply. However, we can arrive at the same conclusions by exploiting the machinery of the product case. That is, we observe that

$$
\Delta=\left[\begin{array}{lll}
1 & 1 & 0
\end{array}\right]
$$

so the LPs have value $\sigma=1$, and the dual solution $x=\left(x_{1}\right)=1$.

Although blocking cannot demonstrate an asymptotic benefit (this is the injective case), we can apply Corollary 5.1 to obtain the tiles

$$
E_{M}=\left\{0, \ldots,\left\lfloor M^{x_{1}} / m\right\rfloor-1\right\}=\{0, \ldots,\lfloor M / 3\rfloor-1\},
$$

supposing $M \geq \mu=3$. A blocked version of the preceding code, where the blocking factor $b=\lfloor M / 3\rfloor$ is assumed to divide $N$, is given by

$$
\begin{aligned}
& \text { for } j=1: N / b, \\
& \quad \text { for } k=1: b, \\
& \quad s=s+a_{i} * b_{i} .
\end{aligned} \quad / / \text { reindexing } i=b \cdot(j-1)+k
$$

### 5.5.2 Matrix-Vector Multiplication

Next we model premultiplying a length $-N$ column vector by a $N$-by- $N$ matrix, expressed componentwise mathematically as $\sum_{i_{2}} a_{i_{1}, i_{2}} b_{i_{2}}$, or in pseudocode as

$$
\begin{gathered}
\text { for } i_{1}=1: N, \quad \text { for } i_{2}=1: N, \\
s_{i_{1}}=s_{i_{1}}+a_{i_{1}, i_{2}} * b_{i_{2}},
\end{gathered}
$$

by a family of algorithms whose parameters satisfy

$$
\begin{gather*}
I=\left\{a_{i_{1}, i_{2}}, b_{i_{2}} \mid\left(i_{1}, i_{2}\right) \in[N]^{2}\right\}, \quad O=\left\{t_{i_{1}} \mid i_{1} \in[N]\right\} \\
R=V I \supseteq\left\{s_{i_{1}, i_{2}} \mid\left(i_{1}, i_{2}\right) \in[N]^{2}\right\}, \quad F \supseteq\left\{z_{i_{1}, i_{2}} \mid\left(i_{1}, i_{2}\right) \in[N]^{2}\right\} \\
\left\{a_{i_{1}, i_{2}}, b_{i_{2}}\right\} \subseteq \operatorname{im}\left(x\left(z_{i_{1}, i_{2}}, \cdot\right)\right) \subseteq\left\{a_{i_{1}, i_{2}}, b_{i_{1}}\right\} \cup\left(R \backslash\left\{s_{i_{1}, i_{2}}\right\}\right)  \tag{5.7}\\
\left\{s_{i_{1}, i_{2}}\right\} \subseteq \operatorname{im}\left(y\left(z_{i_{1}, i_{2}}, \cdot\right)\right) \subseteq R \backslash\left\{s_{k, l} \mid\left(i_{1}, i_{2}\right) \neq(k, l) \in[N]^{2}\right\} .
\end{gather*}
$$

(The constraints on $x$ and $y$ are analogous to those in the dot product example.) The variables $t_{1}, \ldots, t_{M} \in V$ can be any $N$ distinct elements of $R$ (perhaps $s_{i_{1}, i_{2}}$ ). Algorithms satisfying (5.7) induce DAGs of the form shown in Figure 5.2. Each element of the output vector is computed by a dot product (see Section 5.5.1) of a distinct matrix row and the input vector.

Algorithms satisfying (5.7) admit the following HBL interpretation.

$$
\begin{gathered}
m=3, \quad r=2, \quad\left(r_{1}, r_{2}, r_{3}\right)=(2,1,0) \\
K=\left\{z_{i_{1}, i_{2}} \mid\left(i_{1}, i_{2}\right) \in[N]^{2}\right\}, \quad Z\left(z_{i_{1}, i_{2}}\right)=\left(i_{1}, i_{2}\right) \\
\phi_{1}\left(i_{1}, i_{2}\right)=\left(i_{1}, i_{2}\right), \quad \phi_{2}\left(i_{1}, i_{2}\right)=i_{2}, \quad \phi_{3}\left(i_{1}, i_{2}\right)=i_{1} \\
A_{1}\left(i_{1}, i_{2}\right)=a_{i_{1}, i_{2}}, \quad A_{2}\left(i_{2}\right)=b_{i_{2}}, \quad A_{3}\left(i_{1}\right)=t_{i_{1}} .
\end{gathered}
$$

We confirm that $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$. We also recognize that $\phi_{1}$ is an injection, so as in the previous example, the conclusions in Section 5.4 apply, but we can arrive at the same conclusions by exploiting the machinery of the product case. That is, we observe that

$$
\Delta=\left[\begin{array}{lll}
1 & 0 & 1 \\
1 & 1 & 0
\end{array}\right]
$$



Figure 5.2: Matrix-vector multiplication example (Section 5.5.2) in the case $N=3$. Pencils are meant to suggest reductions and expansions (see Section 3.1.6).
so the LPs have value $\sigma=1$, and the dual solution $x=\left(x_{1}, x_{2}\right)=(\alpha, 1-\alpha)$ for any $\alpha \in[0,1]$.
Although blocking cannot demonstrate an asymptotic benefit (again, this is the injective case), we can apply Corollary 5.1 to obtain the tiles

$$
\begin{aligned}
E_{M} & =\left\{0, \ldots,\left\lfloor M^{x_{1}} / m\right\rfloor-1\right\} \times\left\{0, \ldots,\left\lfloor M^{x_{2}} / m\right\rfloor-1\right\} \\
& =\left\{0, \ldots,\left\lfloor M^{\alpha} / 3\right\rfloor-1\right\} \times\left\{0, \ldots,\left\lfloor M^{1-\alpha} / 3\right\rfloor-1\right\},
\end{aligned}
$$

supposing $M \geq \mu=3$. A blocked version of the preceding code, where the blocking factors

$$
b_{1}=\left\{\begin{array}{ll}
\left\lfloor M^{\alpha} / 3\right\rfloor & \alpha>0 \\
1 & \text { otherwise },
\end{array} \quad b_{2}= \begin{cases}\left\lfloor M^{1-\alpha} / 3\right\rfloor & \alpha<1 \\
1 & \text { otherwise },\end{cases}\right.
$$

are both assumed to divide $N$, is given by

$$
\begin{aligned}
& \text { for } j_{1}=1: N / b_{1}, \quad \text { for } j_{2}=1: N / b_{2}, \\
& \text { for } k_{1}=1: b_{1}, \quad \text { for } k_{2}=1: b_{2}, \\
& \quad / / \text { reindexing }\left(i_{1}, i_{2}\right)=b \cdot\left(j_{1}-1, j_{2}-1\right)+\left(k_{1}, k_{2}\right) \\
& \quad s_{i_{1}}=s_{i_{1}}+a_{i_{1}, i_{2}} * b_{i_{2}} .
\end{aligned}
$$

### 5.5.3 Matrix-Matrix Multiplication

Next we model multiplying $N$-by- $N$ matrices, componentwise mathematically as $\sum_{i_{3}} a_{i_{1}, i_{3}} b_{i_{3}, i_{2}}$, or in pseudocode as

$$
\begin{gathered}
\text { for } i_{1}=1: N, \quad \text { for } i_{2}=1: N, \quad \text { for } i_{3}=1: N, \\
s_{i_{1}, i_{2}}=s_{i_{1}, i_{2}}+a_{i_{1}, i_{3}} * b_{i_{3}, i_{2}},
\end{gathered}
$$



Figure 5.3: Matrix-matrix multiplication example (Section 5.5.3) in the case $N=3$.
by a family of algorithms whose parameters satisfy

$$
\begin{gather*}
I=\left\{a_{i_{1}, i_{3}}, b_{i_{3}, i_{2}} \mid\left(i_{1}, i_{2}, i_{3}\right) \in[N]^{3}\right\}, \quad O=\left\{t_{i_{1}, i_{2}} \mid\left(i_{1}, i_{2}\right) \in[N]^{2}\right\} \\
R=V \backslash I \supseteq\left\{s_{i_{1}, i_{2}, i_{3}} \mid\left(i_{1}, i_{2}, i_{3}\right) \in[N]^{3}\right\}, \quad F \supseteq\left\{z_{i_{1}, i_{2}, i_{3}} \mid\left(i_{1}, i_{2}, i_{3}\right) \in[N]^{3}\right\} \\
\left\{a_{i_{1}, i_{3}}, b_{i_{3}, i_{2}}\right\} \subseteq \operatorname{im}\left(x\left(z_{i_{1}, i_{2}, i_{3}}, \cdot\right)\right) \subseteq\left\{a_{i_{1}, i_{3}}, b_{i_{3}, i_{2}}\right\} \cup\left(R \backslash\left\{s_{i_{1}, i_{2}, i_{3}}\right\}\right)  \tag{5.8}\\
\left\{s_{i_{1}, i_{2}, i_{3}}\right\} \subseteq \operatorname{im}\left(y\left(z_{i_{1}, i_{2}, i_{3}}, \cdot\right)\right) \subseteq R \backslash\left\{s_{k, l, m} \mid\left(i_{1}, i_{2}, i_{3}\right) \neq(k, l, m) \in[N]^{3}\right\} .
\end{gather*}
$$

The variables $t_{i_{1}, i_{2}} \in V$ can be any $N^{2}$ distinct elements of $R$ (perhaps $s_{i_{1}, i_{2}, i_{3}}$ ). Algorithms satisfying (5.8) induce DAGs of the form shown in Figure 5.3. Each element of the output matrix is computed by a dot product (see Section 5.5.1) of a distinct matrix row and matrix column.

Algorithms satisfying (5.8) admit the following HBL interpretation.

$$
\begin{gathered}
m=3, \quad r=3, \quad\left(r_{1}, r_{2}, r_{3}\right)=(2,2,2) \\
K=\left\{z_{i_{1}, i_{2}, i_{3}} \mid\left(i_{1}, i_{2}, i_{3}\right) \in[N]^{3}\right\}, \quad Z\left(z_{i_{1}, i_{2}, i_{3}}\right)=\left(i_{1}, i_{2}, i_{3}\right) \\
\phi_{1}\left(i_{1}, i_{2}, i_{3}\right)=\left(i_{1}, i_{3}\right), \quad \phi_{2}\left(i_{1}, i_{2}, i_{3}\right)=\left(i_{3}, i_{2}\right), \quad \phi_{3}\left(i_{1}, i_{2}, i_{3}\right)=\left(i_{1}, i_{2}\right) \\
A_{1}\left(i_{1}, i_{3}\right)=a_{i_{1}, i_{3}}, \quad A_{2}\left(i_{3}, i_{2}\right)=b_{i_{3}, i_{2}}, \quad A_{3}\left(i_{1}, i_{2}\right)=t_{i_{1}, i_{2}} .
\end{gathered}
$$

We confirm that $\bigcap_{j} \operatorname{ker}\left(\phi_{j}\right)=\{0\}$. We have the linear system

$$
\Delta=\left[\begin{array}{lll}
1 & 0 & 1 \\
0 & 1 & 1 \\
1 & 1 & 0
\end{array}\right]
$$

so the LPs have value $\sigma=3 / 2$, and the dual solution $x=\left(x_{1}, x_{2}, x_{3}\right)=(1 / 2,1 / 2,1 / 2)$.
In this case, we can apply Corollary 5.1 to obtain the tiles

$$
E_{M}=\underset{i \in[r]}{\chi}\left\{0, \ldots,\left\lfloor M^{x_{i}} / m\right\rfloor-1\right\}=\left\{0, \ldots,\left\lfloor M^{1 / 2} / 3\right\rfloor-1\right\}
$$

supposing $M \geq \mu=3$. A blocked version of the preceding code, where the blocking factor $b=\left\lfloor M^{1 / 2} / 3\right\rfloor$ is assumed to divide $N$, is given by

$$
\begin{aligned}
& \text { for } j_{1}=1: N / b, \quad \text { for } j_{2}=1: N / b, \quad \text { for } j_{3}=1: N / b, \\
& \text { for } k_{1}=1: b, \quad \text { for } k_{2}=1: b, \quad \text { for } k_{3}=1: b \\
& \quad / / \text { reindexing }\left(i_{1}, i_{2}, i_{3}\right)=b \cdot\left(j_{1}-1, j_{2}-1, j_{3}-1\right)+\left(k_{1}, k_{2}, k_{3}\right) \\
& \quad s_{i_{1}, i_{2}}=s_{i_{1}, i_{2}}+a_{i_{1}, i_{3}} * b_{i_{3}, i_{2}} .
\end{aligned}
$$

This example extends to multiplication of rectangular and sparse matrices, as well as tensor contractions, which can be recast as matrix multiplication (see, e.g., [3, Section 2.6.1]).

### 5.5.4 Multiple-Tensor Contractions

Multiplying multiple matrices or tensors is usually done more efficiently by exploiting associativity, i.e., inserting parentheses and multiplying pairs of matrices/tensors; in this case, we can apply the preceding example (matrix multiplication, Section 5.5.3) to each pair, obtaining a communication lower bound based on $\sigma=3 / 2$. Now we consider two examples where this natural optimization is not exploited.

In the first example, we suppose we obtain an HBL datum with $m=r$ and homomorphisms

$$
\phi_{j}\left(i_{1}, \ldots, i_{r}\right)=\left(i_{j \bmod r}, i_{(j+1) \bmod r}, \ldots, i_{(j+k-1) \bmod r}\right)
$$

given some $k \in\{1, \ldots, r\}$. We can solve the corresponding linear program to see that $\sigma=r / k$. The case $r=3$ and $k=2$ is equivalent to the matrix multiplication example.

In the second example, meant to model a "matricized tensor times Khatri-Rao product operation" (given by G. Ballard in personal communication, May 14, 2015), $m=r \geq 2$ and

$$
\begin{aligned}
\phi_{1}\left(i_{1}, \ldots, i_{r}\right) & =\left(i_{1}, \ldots, i_{r-1}\right) \\
\phi_{2}\left(i_{1}, \ldots, i_{r}\right) & =\left(i_{2}, i_{r}\right) \\
& \vdots \\
\phi_{m}\left(i_{1}, \ldots, i_{r}\right) & =\left(i_{r-1}, i_{r}\right) ;
\end{aligned}
$$

we can solve the corresponding linear program to see that $\sigma=2-1 /(r-1)$. Note that when $r \geq 3, \sigma>3 / 2$.

### 5.5.5 $N$-body Simulation

In our last example, we return to the diamond DAG examples introduced in Section 3.1.5. We have that $m=r$ and for each $j$,

$$
\phi_{j}\left(i_{1}, \ldots, i_{r}\right)=\left(i_{1}, \ldots, i_{j-1}, i_{j+1}, \ldots, i_{r}\right) .
$$

We find that $\sigma=r$, and that $x=\left(x_{j}\right)_{j}=(1, \ldots, 1)$ is the optimal solution of the dual linear program, indicating that the natural approach of tiling these computations with cubes is optimal.

## Chapter 6

## Conclusion

Motivated by the fact that communication (moving data) frequently dominates the costs of computation (performing operations on that data), we have studied the communication costs within a class of nested-loop programs.

We started in Chapter 2 by modeling communication in an abstract machine model, and deriving communication lower bounds. The main contribution of Chapter 2 is a new sharedmemory parallel variant of Hong-Kung's sequential communication lower bounds approach [21], based on the notions of input-/output-path cutsize.

In Chapter 3, we demonstrated how algorithms can model a class of nested-loop programs with group-theoretic interpretations. The main contribution of Chapter 3 is the notion of an HBL interpretation, which generalizes previous work.

In Chapter 4, we developed a deeper theory needed show that algorithms with HBL interpretations have particular input-/output-path cutsize properties. The main contribution of Chapter 4 is deriving sharper constants in certain HBL-type inequalities, which in turn give sharper constants in our lower bounds.

Finally, in Chapter 5, we studied the attainability of the lower bounds, showing that they are attainable within constant factors in many practical cases. The main contribution of Chapter 5 is a framework for finding communication-optimal implementations by tiling.

There are a number of different directions for future work.
First, there are a number of possible improvements to our approach for computing the constraints for HBL-type inequalities in Section 4.7. For example, we have identified simpler approaches when all the homomorphisms have rank-1 or corank-1. Additionally, we believe that a simplification proposed by Valdimarsson [34] in the continuum setting extends to the discrete setting; this would simplify computing the constraints in many more cases, for example, whenever there are at most three homomorphisms, but it would also simplify our algorithm for the general case.

Second, there are number of ways to tighten the communication lower bounds. We have suggested changing the segment length as one option, as well as exploiting knowledge about aliasing arrays, but there are a wide variety of options when it comes to picking an HBL interpretation. It would be interesting to characterize different HBL interpretations that a single algorithm admits.

Third, there are other ways to model communication cost other than the ones proposed here. Our analysis was mostly in terms of per-processor communication costs, and it remains
future work to develop tighter bounds on critical-path communication cost. This will involve a more complicated proof approach, since we largely ignored the influence of processors on each other. Another important aspect of communication is overlap: there is often concurrency in the memory system (network, etc.) which is exploited by overlapping the cost of moving single values by moving data in blocks/messages. Yet another issue is distributed versus shared-memory. In this work we model shared-memory parallel machines but in our previous work we modeled distributed-memory machines. We can model distributed memory in the present model by simulating message-passing in shared-memory; it remains to demonstrate this extension, and show how tiling changes to address it.

Fourth, there is much to be done from an upper-bounds standpoint. We would like to give an automated approach to transform a nested-loop program into a communication-optimal version. There are a number of complications both theoretically (algorithm dependences) and practically (program representations), and this remains future work.

## References

[1] A. Aggarwal, A. Chandra, and M. Snir. "Communication complexity of PRAMs". In: Theoretical Computer Science 71 (1 Mar. 1990), pp. 3-28. ISSN: 0304-3975. DOI: 10.1016/0304-3975(90) 90188-N URL: http://portal.acm.org/citation.cfm? id=77831.77836.
[2] G. Ballard. "Avoiding Communication in Dense Linear Algebra". PhD thesis. EECS Department, University of California, Berkeley, Aug. 2013. url: http://www.eecs. berkeley.edu/Pubs/TechRpts/2013/EECS-2013-151.html.
[3] G. Ballard, E. Carson, J. Demmel, M. Hoemmen, N. Knight, and O. Schwartz. "Communication lower bounds and optimal algorithms for numerical linear algebra". In: Acta Numerica 23 (2014), pp. 1-155.
[4] G. Ballard, J. Demmel, O. Holtz, B. Lipshitz, and O. Schwartz. "Brief announcement: strong scaling of matrix multiplication algorithms and memory-independent communication lower bounds". In: Proceedings of the twenty-fourth annual ACM symposium on Parallelism in algorithms and architectures. ACM. 2012, pp. 77-79.
[5] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz. "Minimizing communication in numerical linear algebra". In: SIAM Journal on Matrix Analysis and Applications 32.3 (2011), pp. 866-901.
[6] F. Barthe. "The Brunn-Minkowski theorem and related geometric and functional inequalities". In: International Congress of Mathematicians. Vol. 2. 2006, pp. 1529-1546.
[7] W. Beckner. "Inequalities in Fourier analysis". In: Annals of Mathematics (1975), pp. 159-182.
[8] C. Bennett and R. Sharpley. Interpolation of Operators. Vol. 129. Pure and Applied Mathematics. Academic Press, 1988.
[9] J. Bennett, A. Carbery, M. Christ, and T. Tao. "Finite bounds for Hölder-BrascampLieb multilinear inequalities". In: Mathematical Research Letters 17.4 (2010), pp. 647666.
[10] J. Bennett, A. Carbery, M. Christ, and T. Tao. "The Brascamp-Lieb inequalities: finiteness, structure, and extremals". In: Geometric and Functional Analysis 17.5 (Jan. 1, 2008), pp. 1343-1415.
[11] G. Bilardi, A. Pietracaprina, and P. D'Alberto. "On the space and access complexity of computation DAGs". In: Graph-Theoretic Concepts in Computer Science. 26th International Workshop, WG 2000 Konstanz, Germany, June 15-17, 2000 Proceedings. 26th International Workshop on Graph-Theoretic Concepts in Computer Science. WG 2000 (Konstanz, Germany, June 15-17, 2000). Ed. by U. Brandes and D. Wagner. Vol. 1928. Lecture Notes in Computer Science. Algorithms and Data Structures Group of the Department of Computer and Information Science, University of Konstanz. Springer-Verlag, 2000, pp. 47-58.
[12] H. J. Brascamp and E. H. Lieb. "Best constants in Young's inequality, its converse, and its generalization to more than three functions". In: Advances in Mathematics 20.2 (1976), pp. 151-173.
[13] E. Carson, J. Demmel, L. Grigori, N. Knight, P. Koanantakool, O. Schwartz, and H. V. Simhadri. Write-Avoiding Algorithms. Tech. rep. UCB/EECS-2015-163. EECS Department, University of California, Berkeley, June 2015. URL: http://www. eecs. berkeley.edu/Pubs/TechRpts/2015/EECS-2015-163.html.
[14] M. Christ. "The optimal constants in Hölder-Brascamp-Lieb inequalities for discrete Abelian groups". In: (July 31, 2013). arXiv:1307.8442v1.
[15] M. Christ, J. Demmel, N. Knight, T. Scanlon, and K. Yelick. Communication lower bounds and optimal algorithms for programs that reference arrays (part 1). Tech. rep. UCB/EECS-2013-61. University of California, Berkeley, Dept. of Electrical Engineering and Computer Science, May 14, 2013.
[16] M. Christ, J. Demmel, N. Knight, T. Scanlon, and K. Yelick. "Communication lower bounds for programs with affine dependences". Talk at SIAM Conference on Parallel Processing for Scientific Computing, Portland, Oregon, USA. Feb. 19, 2014. URL: http: //www.cs.berkeley.edu/~knight/cdksy_PP14_slides.pdf.
[17] G. Collins. "Quantifier elimination for real closed fields by cylindrical algebraic decomposition". In: Automata Theory and Formal Languages: 2nd GI Conference. Springer. 1975, pp. 134-183.
[18] V. Elango, F. Rastello, L. Pouchet, J. Ramanujam, and P. Sadayappan. "On Characterizing the Data Access Complexity of Programs". In: Proceedings of the $42 n d$ Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015. 2015, pp. 567-580. DOI: 10.1145/ 2676726.2677010. URL: http://doi.acm.org/10.1145/2676726.2677010.
[19] P. van Emde Boas and J. van Leeuwen. "Move rules and trade-offs in the pebble game". English. In: Theoretical Computer Science 4th GI Conference. Ed. by K. Weihrauch. Vol. 67. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1979, pp. 101112. ISBN: 978-3-540-09118-9. DOI: 10.1007/3-540-09118-1_12, uRL: http://dx. doi.org/10.1007/3-540-09118-1_12.
[20] O. Hölder. "Uber einen Mittelwertssatz". In: Nachr. Akad. Wiss. Göttingen Math.Phys. Kl. (1889), pp. 38-47.
[21] J.-W. Hong and H. T. Kung. "I/O complexity: the red-blue pebble game". In: Proceedings of the 13th annual ACM symposium on theory of computing. STOC '81. (Milwaukee, Wisconsin, USA, 1981). ACM special interest group on algorithms and computation theory (SIGACT). New York, New York, USA: ACM, 1981, pp. 326-333.
[22] D. Irony, S. Toledo, and A. Tiskin. "Communication lower bounds for distributedmemory matrix multiplication". In: Journal of Parallel and Distributed Computing 64.9 (2004), pp. 1017-1026.
[23] S. Lang. Algebra. $3^{\text {rd }}$. Vol. 211. Graduate Texts in Mathematics. Springer, 2002. ISBN: 9780387953854.
[24] L. Loomis and H. Whitney. "An inequality related to the isoperimetric inequality". In: Bulletin of the American Mathematical Society 55 (1949), pp. 961-962.
[25] B. Lowery and J. Langou. "Improving Communication Lower Bounds for MatrixMatrix Multiplication". Talk at the 9th Scheduling for Large Scale Systems Workshop, Lyon, France. July 2, 2014. URL: http://scheduling2014.sciencesconf.org/ conference/scheduling2014/pages/julien_langou.pdf.
[26] L. Maligranda. "Why Hölder's inequality should be called Rogers' inequality". In: Math. Inequal. Appl 1.1 (1998), pp. 69-83.
[27] Y. Matiyasevich. Hilbert's Tenth Problem. MIT Press, 1993.
[28] B. Poonen. "Hilbert's Tenth Problem over rings of number-theoretic interest". Notes from Arizona Winter School, "Logic and Number Theory", March 15-19, 2003, Tuscon, AZ. 2003.
[29] L. J. Rogers. "An extension of a certain theorem in inequalities". In: Messenger of Math 17 (1888), pp. 145-150.
[30] W. Rudin. Real and Complex Analysis. Tata McGraw-Hill Education, 1987.
[31] J. E. Savage. "Extending the Hong-Kung model to memory hierarchies". In: Computing and Combinatorics. First Annual International Conference, COCOON '95 Xi'an, China, August 24-26, 1995 Proceedings. First Annual International Conference on Computing and Combinatorics. COCOON '95 (Xi'an, China, Aug. 24-26, 1995). Ed. by D.-Z. Du and M. Li. Vol. 959. Lecture Notes in Computer Science. Springer, 1995, pp. 270-281.
[32] A. Tarski. A decision method for elementary algebra and geometry. Tech. rep. R-109. RAND Corporation, 1951.
[33] A. Tiskin. "The design and analysis of bulk-synchronous parallel algorithms". PhD thesis. Oxford Univ., 1998.
[34] S. I. Valdimarsson. "Optimizers for the Brascamp-Lieb inequality". In: Israel Journal of Mathematics 168.1 (Nov. 1, 2008), pp. 253-274.
[35] S. I. Valdimarsson. "The Brascamp-Lieb polyhedron". In: Canadian Journal of Mathematics 62.4 (2010), pp. 870-888.
[36] W. Young. "On the multiplication of successions of Fourier constants". In: Proc. Roy. Soc. of London, Ser. A (1912), pp. 331-339.

