Top- k -Convolution and the Quest for Near-Linear Output-Sensitive Subset Sum

Karl Bringmann
bringmann@cs.uni-saarland.de
Saarland University and
Max Planck Institute for Informatics
Saarbrücken, Germany

Vasileios Nakos
vnakos@mpi-inf.mpg.de
Saarland University and
Max Planck Institute for Informatics
Saarbrücken, Germany

ABSTRACT

In the classical SUBSETSUM problem we are given a set X and a target t , and the task is to decide whether there exists a subset of X which sums to t . A recent line of research has resulted in $(t \cdot \text{poly}(\log t))$ -time algorithms, which are (near-)optimal under popular complexity-theoretic assumptions. On the other hand, the standard dynamic programming algorithm runs in time $O(n \cdot |S(X, t)|)$, where $S(X, t)$ is the set of all subset sums of X that are smaller than t . All previous pseudopolynomial algorithms actually solve a stronger task, since they actually compute the whole set $S(X, t)$.

As the aforementioned two running times are incomparable, in this paper we ask whether one can achieve the best of both worlds: running time $|S(X, t)| \cdot \text{poly}(\log t)$. In particular, we ask whether $S(X, t)$ can be computed in near-linear time in the output-size. Using a diverse toolkit containing techniques such as color coding, sparse recovery, and sumset estimates, we make considerable progress towards this question and design an algorithm running in time $|S(X, t)|^{4/3} \cdot \text{poly}(\log t)$.

Central to our approach is the study of *Top- k -Convolution*, a natural problem of independent interest: given degree- d sparse polynomials with non-negative coefficients, compute the lowest k non-zero monomials of their product. We design an algorithm running in time $k^{4/3} \text{poly}(\log d)$, by a combination of sparse convolution and sumset estimates considered in Additive Combinatorics. Moreover, we provide evidence that going beyond some of the barriers we have faced requires either an algorithmic breakthrough or possibly new techniques from Additive Combinatorics on how to pass from information on restricted sumsets to information on unrestricted sumsets.

CCS CONCEPTS

• **Theory of computation** $\rightarrow$ *Dynamic programming*.

*This work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 850979).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '20, June 22–26, 2020, Chicago, IL, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6979-4/20/06...\$15.00

<https://doi.org/10.1145/3357713.3384308>

KEYWORDS

Subset Sum, pseudopolynomial, output-sensitive, convolution, restricted sumset

ACM Reference Format:

Karl Bringmann and Vasileios Nakos. 2020. Top- k -Convolution and the Quest for Near-Linear Output-Sensitive Subset Sum. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC '20)*, June 22–26, 2020, Chicago, IL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3357713.3384308>

1 INTRODUCTION

1.1 Subset Sum

SUBSETSUM is a fundamental problem at the intersection of computer science, mathematical optimization, and operations research. In this problem, given a set X of n integers and a target t , the task is to decide whether there exists a subset of X that sums to t . The problem belongs to Karp's initial list of NP-complete problems [28], and it has given rise to a plethora of algorithmic techniques, see, e.g., the monographs [29, 33]. Apart from being a cornerstone in algorithm design, SUBSETSUM draws its importance from being a special case of many other problems, like KNAPSACK or INTEGER PROGRAMMING. It has also played a role in cryptography, as Merkle and Hellman [34] based their cryptosystem on this problem, see also [13, 18, 26, 39, 42].

Several classic algorithms for SUBSETSUM are typically taught in undergraduate courses, including the meet-in-the-middle algorithm running in time $O(2^{n/2})$ [25] and Bellman's dynamic programming algorithm running in pseudopolynomial time $O(n \cdot t)$ [12].

Surprisingly, after decades of research, major algorithmic advances were still discovered in the last 10 years, e.g., [1, 7–11, 14, 21, 24, 27, 30–32, 37, 38]. Among these developments, the most relevant for this paper are improvements over Bellman's $O(n \cdot t)$ algorithm: Koiliaris and Xu [30] designed a deterministic algorithm running in time¹ $\tilde{O}(\min\{\sqrt{n} \cdot t, t^{4/3}\})$, and Bringmann [14] devised a randomized algorithm running in time $\tilde{O}(t)$ (which was improved in terms of log factors in [27]). The running time $\tilde{O}(t)$ of the randomized algorithms is optimal under the Strong Exponential Time Hypothesis [1] as well as under the SETCOVER Hypothesis [20].

Thus, research on pseudopolynomial algorithms for SUBSETSUM with respect to parameter t is more or less finished. However, it remains to study whether the recent improvements generalize to other parameters as well as to variants of SUBSETSUM. For instance, this has been done for the MODULARSUBSETSUM problem in [10].

¹By $\tilde{O}(T)$ we hide factors of the form $\text{polylog}(T)$ as well as factors $\text{polylog}(u)$, where u is the universe size, and $\text{polylog}(t)$, where t is the target.

In this paper, we start from the observation that Bellman’s classic dynamic programming algorithm can be implemented to run in time $O(n \cdot |S(X, t)|)$, where $S(X, t)$ is the set of all subset sums of X below t . Since $|S(X, t)|$ can be much smaller than t , so far the running times $\tilde{O}(t)$ and $O(n \cdot |S(X, t)|)$ are incomparable. Thus, despite the running time $\tilde{O}(t)$ being matched by a conditional lower bound, in situations where $|S(X, t)|$ is small Bellman’s algorithm can outperform the recent improved algorithms. To obtain the best of both worlds, it would thus be desirable to consider $|S(X, t)|$, rather than t , as the parameter to measure the computational complexity of the problem, and to similarly shave off the factor n from the running time of Bellman’s algorithm. In particular, since all previous pseudopolynomial algorithms for SUBSETSUM produce all attainable subset sums smaller than t , a natural question is whether one can design a *near-linear output-sensitive* algorithm.

QUESTION 1.1. *Is there an algorithm that computes $S(X, t)$ in time $\tilde{O}(|S(X, t)|)$?*

Our work struggles to make progress towards understanding SUBSETSUM under this new computational perspective, and it lead us to study a new type of sparse convolution problem.

1.2 Top- k -Convolution

Convolution and Boolean convolution are fundamental computational primitives that frequently arise in algorithm design, e.g., when combining solutions of two subproblems. The *Boolean convolution* $f \oplus g$ of vectors $f, g \in \{0, 1\}^u$ is the vector with entries $(f \oplus g)_k = \bigvee_{0 \leq i \leq k} f_i \wedge g_{k-i}$ for $0 \leq k < 2u$. This arises when we split a problem into two subproblems, so that the whole problem has a solution of size k if and only if for some i the left subproblem has a solution of size i and the right subproblem has a solution of size $k - i$. Moreover, Boolean convolution is equivalent to *sumset computation*, where we are given sets $A, B \subseteq \{0, 1, \dots, u-1\}$ and the task is to compute $A + B$, the set of all sums $a + b$ with $a \in A, b \in B$.

The *convolution* $f \star g$ of vectors $f, g \in \mathbb{R}^u$ is the vector with entries $(f \star g)_k = \sum_{i=0}^k f_i \cdot g_{k-i}$ for $0 \leq k < 2u$. When we split a problem into two subproblems, and f_i and g_i count the number of size- i solutions of the left and right subproblem, then $(f \star g)_k$ counts the number of size- k solutions of the whole problem. Moreover, convolution is equivalent to polynomial multiplication, where we are given the coefficients of two polynomials and want to compute the coefficients of their product.

Boolean convolution can be solved via convolution, and convolution can be solved in time $O(u \log u)$ using Fast Fourier Transform (FFT). However, when the input vectors are sparse, we can ask for algorithms that compute convolutions much faster than performing FFT. Specifically, the ultimate goal is an algorithm running in near-linear output-sensitive time, i.e., in near-linear time in terms of the number of non-zero entries of $f \star g$. This practically and theoretically relevant problem is called sparse convolution (or sparse polynomial multiplication) and has been studied, e.g., in [6, 15, 17, 19, 35, 36, 40, 41, 46]. The ultimate goal of near-linear output-sensitive time has been achieved by Cole and Hariharan [19] for non-negative vectors by a Las Vegas algorithm, in [36] for general vectors by a Monte Carlo algorithm, and in [15]

for non-negative vectors by an almost linear time deterministic algorithm.

In this paper we study a natural variant that we call *Top- k -Convolution*: Given vectors $f, g \in \mathbb{R}^u$, compute the k lowest non-zero entries of $f \star g$. Formally, denoting by $i_1 < i_2 < \dots < i_\ell$ all indices of non-zero entries of $f \star g$, our goal is to compute the pairs $(i_1, (f \star g)_{i_1}), \dots, (i_k, (f \star g)_{i_k})$. *Boolean Top- k -Convolution* is defined analogously, i.e., the task is to compute the lowest k indices $i_1 < \dots < i_k$ of 1-entries of $f \oplus g$.

“Top- k ” problems, that ask for the k best solutions, are well motivated from a practical perspective, e.g., for displaying search results. Note that in the setting where we split a problem into two subproblems, Boolean Top- k -Convolution asks for the k smallest solution sizes. In the polynomial multiplication setting, Top- k -Convolution asks for the k lowest-degree monomials in the product of two given sparse polynomials. The problem is equivalent if we replace “lowest” by “highest” (by reversing f and g), and thus Top- k -Convolution is also equivalent to computing the k highest-degree monomials in the product of two given sparse polynomials. As an additional major application, we present a connection to SUBSETSUM in this paper.

Therefore, sparse Top- k -Convolution is a well-motivated problem with several applications, and it is surprising that to the best of our knowledge it has not been explicitly studied before. Throughout the paper we will assume that f, g are non-negative vectors, since this is satisfied in many applications and the possible cancelations resulting from negative entries make the problem considerably harder (indeed, sparse convolution with negative entries has been solved much later [36] than on non-negative vectors [19]). Note that Top- k -Convolution on non-negative vectors can be solved naively in time $O(k^2)$: The lowest k non-zeros of $f \star g$ are among the combinations of the lowest k non-zeros of f and the lowest k non-zeros of g . This allows us to assume that f and g are k -sparse, and we can compute $f \star g$ naively in time $O(k^2)$ to obtain its lowest k non-zeros.

1.3 Our Contribution

We initiate the study of SUBSETSUM with respect to the parameter $|S(X, t)|$, making considerable progress in this direction. We show an “output-sensitivity preserving” reduction from SUBSETSUM to Top- k -Convolution, such that if the latter can be solved in near-linear output-sensitive time so can the former. Then, we investigate upper and lower bounds for Top- k -Convolution on non-negative vectors (as well as related restricted variants of sumset computation and sparse convolution). In particular, we present a randomized $\tilde{O}(k^{4/3})$ -time algorithm for Top- k -Convolution, resulting in time $\tilde{O}(|S(X, t)|^{4/3})$ for SUBSETSUM. Our algorithms fall into a natural class that we call “rectangle covering algorithms”, for which we show that they cannot yield near-linear time. Our technical machinery draws from a wide range of techniques such as color coding, sparse convolutions, and sumset estimates from Additive Combinatorics.

2 RESULTS AND TECHNIQUES

2.1 Preliminaries

We write $\mathbb{N} = \{0, 1, \dots\}$ and $[n] = \{0, 1, \dots, n\}$ for any $n \in \mathbb{N}$. For sets $A, B \subseteq \mathbb{N}$ define $A + B = \{a + b \mid a \in A, b \in B\}$. We also define

$\max(A) = \max_{x \in A} x$, and similarly $\min(A)$. For any set $X \subseteq \mathbb{N}$ we define $\Sigma(X) = \sum_{x \in X} x$. For $t \in \mathbb{N}$ we define the set of all subset sums of X below t as

$$\mathcal{S}(X, t) = \{\Sigma(Y) \mid Y \subseteq X, \Sigma(Y) \leq t\}.$$

For integers a, b we write $[a, b] = \{a, \dots, b\}$, $(a, b) = \{a+1, \dots, b-1\}$, and similarly for $(a, b]$ and $[a, b)$.

For a vector $f \in \mathbb{R}^d$ we let $\|f\|_0$ be its number of non-zero entries. For a set $S \subseteq [d]$ we denote by f_S the vector that is zeroed out outside of S , i.e., $(f_S)_i = f_i$ for $i \in S$ and $(f_S)_i = 0$ otherwise. We shall use 0-indexed vectors throughout the paper. We also use a non-standard notation of $\tilde{O}(T)$ throughout the paper that hides factors of the form $\text{polylog}(T)$ as well as $\text{polylog}(u)$, where u is the universe size, or $\text{polylog}(t)$, where t is the target.

We shall need the following result by Cole and Hariharan and an immediate corollary.

THEOREM 2.1 ([19]). *Given non-negative vectors f, g of length d , we can compute $f \star g$ in expected time $O(\|f \star g\|_0 \cdot \log^2 d)$.*

THEOREM 2.2. *Given $A, B \subseteq [u]$, we can compute $A + B$ in expected time $O(|A + B| \cdot \log^2 u)$.*

PROOF. Let f be the indicator vector of A , g be the indicator vector of B , $d = 2u + 1$ and use Theorem 2.1. $\square$

2.1.1 Problem Definitions. Our work is concerned with the following problems.

DEFINITION 2.3 (SUBSET SUM). *Given a set $X \subseteq \mathbb{N}$ and a number t , compute $\mathcal{S}(X, t) = \{\Sigma(Y) \mid Y \subseteq X, \Sigma(Y) \leq t\}$. We measure the running time in terms of $|\mathcal{S}(X, t)|$. We write $n = |X|$.*

DEFINITION 2.4 (TOP-k-CONVOLUTION). *Given two vectors $f, g \in \mathbb{R}^u$ and a parameter k , compute the first k non-zero entries of $f \star g$. In other words, find $(i_1, (f \star g)_{i_1}), (i_2, (f \star g)_{i_2}), \dots, (i_k, (f \star g)_{i_k})$ where $i_1 < i_2 < \dots < i_k$ are the smallest k indices on which $f \star g$ is non-zero. We measure the running time in terms of k . We write $n = \|f\|_0$ and $m = \|g\|_0$.*

Top-k-Convolution is equivalent to the following problem, as we will show in Lemma 3.1.

DEFINITION 2.5 (PREFIX-RESTRICTED CONVOLUTION). *Given a positive integer u and two vectors $f, g \in \mathbb{R}^u$, compute $(f \star g)_{[u]}$, i.e., compute (a sparse representation of) the first u entries of $f \star g$. We measure the running time in terms of the output-size $\text{out} = \|(f \star g)_{[u]}\|_0$. We write $n = \|f\|_0$ and $m = \|g\|_0$.*

The following problem is a Boolean version of prefix-restricted convolution.

DEFINITION 2.6 (PREFIX-RESTRICTED SUMSET COMPUTATION). *Given $u \in \mathbb{N}$ and $A, B \subseteq [u]$, compute $(A + B) \cap [u]$. We measure the running time in terms of the output-size $\text{out} = |(A + B) \cap [u]|$. We write $n = |A|$ and $m = |B|$.*

2.1.2 Covering of Prefix-Restricted Sumsets. For a set $A \subseteq \mathbb{N}$ of size n , we implicitly assume that A is sorted and we denote its elements in sorted order as $A_1 < A_2 < \dots < A_n$. Moreover, for $I \subseteq \{1, \dots, n\}$ we write A_I for $\{A_i \mid i \in I\}$. We also let $[n] = \{1, 2, \dots, n\}$.

We define the notion of a covering of a restricted sumset $(A, B, [u])$. Intuitively, we want to cover $(A + B) \cap [u]$ by sumsets of the form $A_I + B_J$, that is, we want to have $(A + B) \cap [u] \subseteq \bigcup_{(I, J) \in C} A_I + B_J$.

DEFINITION 2.7. *Let $u \in \mathbb{N}$ and $A, B \subseteq [u]$ with $n = |A|$, $m = |B|$. A **covering** of $(A, B, [u])$ is a family C such that:*

- (1) *C consists of pairs (I, J) where $I \subseteq [n]$ and $J \subseteq [m]$.*
- (2) *For any $1 \leq i \leq n$, $1 \leq j \leq m$ with $A_i + B_j \in [u]$ there exists $(I, J) \in C$ with $(i, j) \in I \times J$.*

We call a covering C **unique** if the pair in property 2. is unique, i.e., for any $1 \leq i \leq n$, $1 \leq j \leq m$ with $A_i + B_j \in [u]$ there exists a unique pair $(I, J) \in C$ such that $(i, j) \in I \times J$.

We call a covering C a **rectangle covering** if for any $(I, J) \in C$ the sets I and J are intervals, meaning that they consists of contiguous elements in the sorted order of A and B , respectively.

The **cost** of a covering C is

$$\sum_{(I, J) \in C} |A_I + B_J|.$$

This notion is useful because of the following fact.

OBSERVATION 2.8. *Given a covering C of $(A, B, [u])$ of cost c , we can compute the Prefix-Restricted sumset $(A + B) \cap [u]$ in expected time $\tilde{O}(c)$.*

PROOF. Using output-sensitive sumset computation (Theorem 2.2), we can compute $A_I + B_J$ in expected time $O(|A_I + B_J| \log^2 u)$. Thus, we can compute $R := \bigcup_{(I, J) \in C} A_I + B_J$ in time $O(c \log^2 u)$. By the covering property, we can simply return $R \cap [u] = (A + B) \cap [u]$. $\square$

We refer to an algorithm making use of Observation 2.8 as a **covering algorithm**. We call it a **unique-rectangle-covering algorithm** if the used covering is a unique rectangle covering. This is a natural class of algorithms, as we also explain in Section 2.3.2. All algorithms presented in this paper are unique-rectangle-covering algorithms.

2.2 Formal Statement of Results

2.2.1 Main Results. As the technical core of our paper, we present an efficient construction of low-cost coverings.

THEOREM 2.9 (COVERING CONSTRUCTION, SECTION 7). *Given $A, B \subseteq [u]$, in expected time $\tilde{O}(\text{out}^{4/3})$ we can compute a unique rectangle covering of $(A, B, [u])$ of cost $\tilde{O}(\text{out}^{4/3})$, where $\text{out} = |(A + B) \cap [u]|$.*

By Observation 2.8, this yields an $\tilde{O}(\text{out}^{4/3})$ Las Vegas algorithm for Prefix-Restricted sumset computation. Via simple reductions, we obtain similar algorithms for our convolution problems.

COROLLARY 2.10 (TOP-k-CONVOLUTION AND RELATED PROBLEMS, SECTION 3). *Top-k-Convolution on non-negative vectors can be solved in expected time $\tilde{O}(k^{4/3})$. Prefix-restricted sumset computation and Prefix-Restricted convolution on non-negative vectors can be solved in expected time $\tilde{O}(\text{out}^{4/3})$.*

By carefully adapting a recent pseudopolynomial $\tilde{O}(n + t)$ -time SUBSETSUM algorithm [14] to use Prefix-Restricted sumset computation as a subroutine, we obtain our main result for SUBSETSUM. This can be seen as a reduction from SUBSETSUM to Prefix-Restricted sumset computation.

THEOREM 2.11 (SUBSETSUM). *Given $X \subseteq \mathbb{N}$ and $t \in \mathbb{N}$, we can compute the set $\mathcal{S}(X, t)$ in time $\tilde{O}(|\mathcal{S}(X, t)|^{4/3})$ with high probability.*

Since all of these results depend on our technical core (Theorem 2.9), we also study limitations of rectangle-covering algorithms. The following result shows that our approach of using rectangle coverings to solve Top- k -Convolution and SUBSETSUM cannot achieve near-linear output-sensitive algorithms.

THEOREM 2.12 (LOWER BOUND ON RECTANGLE COVERINGS). *There exists an infinite sequence of tuples $(A, B, [u])$, with $u \in \mathbb{N}$ and $A, B \subseteq [u]$, such that any rectangle covering of $(A, B, [u])$ has cost $\Omega(\text{out}^{1.047})$, where $\text{out} = |(A + B) \cap [u]|$.*

We remark that this result crucially uses *rectangle* coverings; we do not rule out the existence of non-rectangle coverings of near-linear cost.

In the remainder of Section 2.2 we present additional related results.

2.2.2 Additional Results: Relaxing the Upper Bound. In our results so far we have relaxed the ultimate goal of algorithms running in time $\tilde{O}(\text{out})$ to time $\tilde{O}(\text{out}^{4/3})$. We can alternatively relax the goal by measuring the running time in terms of a slightly larger upper bound. Specifically, for Prefix-Restricted sumset computation so far we measured time in terms of the output-size $\text{out} = |(A + B) \cap [u]|$. Now we will measure the running time in terms of $|(A + B) \cap [(1 + \zeta)u]|$ for some $\zeta > 0$, while the task still is to compute the set $(A + B) \cap [u]$. Since one could expect that for “realistic” instances (A, B, u) there are not many sums in the boundary region $(A + B) \cap [u, (1 + \zeta)u]$, algorithms that perform well with respect to this new measure might perform well in practice. We show the following.

THEOREM 2.13 (RELAXED UPPER BOUND, SECTION 6). *Prefix-Restricted sumset computation can be solved in expected time $\tilde{O}(\zeta^{-1} \cdot |(A + B) \cap [(1 + \zeta)u]| + \zeta^{-2})$. Prefix-Restricted convolution on non-negative vectors can be solved in expected time $\tilde{O}(\zeta^{-1} \cdot \|(f \star g)_{[(1 + \zeta)u]}\|_0 + \zeta^{-2})$. SUBSETSUM can be solved in time $\tilde{O}(\zeta^{-1} \cdot |\mathcal{S}(X, (1 + \zeta)t)| + \zeta^{-2})$ with high probability.*

Note that $|\mathcal{S}(X, (1 + \zeta)t)| \leq (1 + \zeta)t$, and thus plugging in any constant $\zeta > 0$ yields an algorithm running in time $\tilde{O}(t)$, which is as good as [14, 27].

2.2.3 Additional Results: Interval-Restricted Sumsets and Convolution. So far we studied the problem of computing the first u output values of sumset computation and convolution. More generally, we could ask to compute all output values at given positions $P \subseteq \mathbb{N}$. This is a well-motivated generalization, as it corresponds to computing specific coefficients of the product of two polynomials, which comes up in counting problems.

However, here we show that even when the desired positions P form an *interval* then near-linear output-sensitive algorithms are unlikely to exist. Specifically, we study the following problems.

DEFINITION 2.14 (INTERVAL-RESTRICTED SUMSET COMPUTATION). *Given positive integers $\ell \leq u$ and $A, B \subseteq [u]$, compute $(A + B) \cap [\ell, u]$. We denote the output-size by $\text{out} = |(A + B) \cap [\ell, u]|$. We write $n = |A|$ and $m = |B|$.*

DEFINITION 2.15 (INTERVAL-RESTRICTED CONVOLUTION). *Given positive integers $\ell \leq u$ and vectors $f, g \in \mathbb{R}^u$, compute $(f \star g)_{[\ell, u]}$,*

i.e., compute (a sparse representation of) the entries of $(f \star g)_i$ for $\ell \leq i \leq u$. We denote the output-size by $\text{out} = \|(f \star g)_{[\ell, u]}\|_0$. We write $n = \|f\|_0$ and $m = \|g\|_0$.

We design algorithms for these problems with the following running time. This uses a natural generalization of coverings, where we simply replace the set $[u]$ by $[\ell, u]$.

THEOREM 2.16 (ALGORITHM FOR INTERVAL-RESTRICTED). *Interval-Restricted sumset computation and Interval-Restricted non-negative convolution can be solved in expected time $\tilde{O}(n + m + \sqrt{nm \text{out}})$.*

We present conditional lower bounds based on two classical problems: Boolean matrix multiplication and sliding window Hamming distance. In Boolean matrix multiplication, we are asked to multiply two Boolean $n \times n$ -matrices. By a reduction to multiplying real-valued matrices, Boolean matrix multiplication can be solved in time $O(n^\omega)$, where $\omega \leq 2.373$ is the exponent of fast matrix multiplication [22, 47]. Limitations of fast matrix multiplication have been recently studied, and it was shown that known techniques cannot facilitate $\omega < 2 + \frac{1}{6}$ [3–5]. In particular, $\omega > 2$ is a barrier that known techniques cannot overcome, also for Boolean matrix multiplication.

In the sliding window Hamming distance problem, we are given a text of length $2n$ and a pattern of length n , both over a (possibly large) alphabet Σ , and want to find the Hamming distance between the pattern and every length- n substring of the text. This problem admits a $\tilde{O}(n^{3/2})$ -time algorithm, and it is a major open question in string algorithms whether a faster algorithm exists [2, 23].

THEOREM 2.17 (HARDNESS FOR INTERVAL-RESTRICTED). *Let $\delta > 0$. If interval-Restricted sumset computation is in time $\tilde{O}((n + m + \text{out})^{\omega/2 - \delta})$, then Boolean matrix multiplication is in time $\tilde{O}(n^{\omega - 2\delta})$.*

If interval-Restricted convolution on non-negative vectors is in time $\tilde{O}(n + m + (nm \text{out})^{1/2 - \delta})$, then sliding window Hamming distance is in time $\tilde{O}(n^{3/2 - 3\delta})$.

2.3 Overview of Techniques

2.3.1 Reduction from Subset Sum to Prefix-Restricted Sumset Computation. To reduce SUBSETSUM to Prefix-Restricted sumset computation, our starting point is the pseudopolynomial $\tilde{O}(t)$ -time algorithm of [14], which uses a combination of color coding and FFT. We observe that in this algorithm all uses of FFT in fact compute Prefix-Restricted sumsets, so we can replace FFT by an algorithm for Prefix-Restricted sumset computation. This directly yields a reduction from SUBSETSUM to Prefix-Restricted sumset computation, but it does not directly give the desired guarantees. Our goal is to make this reduction “output-sensitivity preserving”, that is, if we can compute Prefix-Restricted sumsets in near-linear output-sensitive time then the reduction yields an algorithm for SUBSETSUM that runs in near-linear output-sensitive time.

To this end, we need to take a closer look at the algorithm in [14]. Given (X, t) , this algorithm splits X into the set $X^{(L)}$ of numbers that are larger than $t/\text{polylog}(n)$ and the remaining small numbers $X^{(S)}$. Any subset summing to at most t uses at most $\text{polylog}(n)$ large numbers, which enables a color-coding-based algorithm in [14]. We observe that, without any problems, replacing the usage of FFT in that part of the algorithm by near-linear output-sensitive

Prefix-Restricted sumset computation yields a near-linear output-sensitive algorithm for handling the large numbers $X^{(L)}$. The small numbers $X^{(S)}$ are then further split at random into two subsets $X^{(1)}, X^{(2)}$. By concentration inequalities, this splits the subset Y into two subsets $Y^{(i)} = Y \cap X^{(i)}$ satisfying with high probability $\Sigma(Y^{(i)}) \leq (1 + \epsilon)t/2$. We can therefore find $\Sigma(Y^{(1)})$ and $\Sigma(Y^{(2)})$ by recursive calls on $(X^{(1)}, (1 + \epsilon)t/2)$ and $(X^{(2)}, (1 + \epsilon)t/2)$. Since we recurse on two subproblems and the target bound t is roughly split in half, one can argue that the total running time is $\tilde{O}(t)$ [14].

In this paper, we instead measure the running time in terms of $|\mathcal{S}(X, t)|$. Hence, it does not suffice that the target bound t is roughly split in half; we also need that the measure $|\mathcal{S}(X, t)|$ is roughly split in half. Indeed, if we can show this, then we again obtain two subproblems and the measure $|\mathcal{S}(X, t)|$ is roughly split in half, so we can argue that the running time is $\tilde{O}(|\mathcal{S}(X, t)|)$. This requires novel insights into the set of all subset sums, to obtain inequalities relating the number of subset sums of X to that of its subsets. Specifically, we show that $X^{(1)}$ and $X^{(2)}$ satisfy

$$|\mathcal{S}(X^{(1)}, (1+\epsilon)t/2)| + |\mathcal{S}(X^{(2)}, (1+\epsilon)t/2)| \leq (1+O(\epsilon))|\mathcal{S}(X, t)| + O(1).$$

This inequality only holds if every element of $X^{(1)}, X^{(2)}$ is sufficiently smaller than the target t . Indeed, in the proof of this inequality we use that all numbers in $X^{(1)}, X^{(2)}$ are small, which justifies why we removed the large numbers $X^{(L)}$.

2.3.2 Restricted Sumset Computation. We now give an overview of our techniques for restricted sumset computation, and discuss the relation to classical results from Additive Combinatorics.

Consider the Prefix-Restricted sumset computation of $(A+B) \cap [u]$. Note that we can assume $A, B \subseteq [u]$, since larger numbers cannot form sums of at most u . We also always implicitly assume that for every number $a \in A$ there is a number $b \in B$ with $a + b \leq u$, since otherwise we can remove a from A without changing $(A+B) \cap [u]$. In other words, we assume $\max(A) + \min(B) \leq u$, and symmetrically $\min(A) + \max(B) \leq u$, which can be ensured after a $O(|A| + |B|)$ -time preprocessing.

Standard sumset computation does not suffice. Standard sparse convolution (Theorem 2.2) allows us to compute $A+B$ in near-linear output-sensitive time $\tilde{O}(|A+B|)$. In particular, since $A+B \subseteq [2u]$, we can compute $(A+B) \cap [2u]$ in near-linear output-sensitive time. Although this might seem close at first glance, careful inspection reveals that it can be much larger than the goal $|(A+B) \cap [u]|$. Indeed, we can construct sets $A, B \subseteq [u]$ with $|A| = |B| = n$ satisfying $|(A+B) \cap [2u]| = \Omega(n^2)$ and $|(A+B) \cap [u]| = O(n)$. Hence, in general we cannot afford to compute $(A+B) \cap [2u]$ when our goal is to compute $(A+B) \cap [u]$ in near-linear output-sensitive time.

For this construction, assume that u is a sufficiently large even integer and let

$$\begin{aligned} X &:= [n] = \{0, 1, 2, \dots, n\}, & Y &:= n[n] = \{0, n, 2n, \dots, n^2\}, \\ A &:= \{0\} \cup \{\frac{u}{2}\} + X, & B &:= \{0\} \cup \{\frac{u}{2}\} + Y. \end{aligned}$$

One can check that $\max(A) + \min(B), \min(A) + \max(B) \leq \frac{u}{2} + n^2 \leq u$. Since for any $x \in X, y \in Y$ the sum $(\frac{u}{2} + x) + (\frac{u}{2} + y)$ is larger than u , we obtain $|(A+B) \cap [u]| \leq |\{0\}| + |X| + |Y| = 2n+1$. Moreover, one can check that $X+Y = [n^2+n]$, and thus $|(A+B) \cap [2u]| = |A+B| \geq n^2$.

A natural class of algorithms. We investigate a natural and safe algorithmic approach: To compute $(A+B) \cap [u]$, we make sure that every sum $A_i + B_j \leq u$ is computed, where $1 \leq i \leq n, 1 \leq j \leq m$. Since we can use as a subroutine that standard sumset computation is in near-linear output-sensitive time, it is natural to compute sumsets $A_I + B_J$ for some $I \subseteq \{1, \dots, n\}$ and $J \subseteq \{1, \dots, m\}$. Combining these two arguments, we arrive at our notion of *covering*. In particular, we want to compute a covering C of $(A, B, [u])$, and then use Observation 2.8 to compute $(\bigcup_{(I,J) \in C} A_I + B_J) \cap [u] = (A+B) \cap [u]$ in time proportional to the cost of C . Hence, covering algorithms are a natural class of algorithms for Prefix-Restricted sumset computation.

We are particularly interested in *unique* coverings, because they allow us to compute Prefix-Restricted convolutions: A pair $A_i + B_j$ being covered exactly once corresponds to adding the product $f_i \cdot g_j$ exactly once to the output value $(f \star g)_{i+j}$.

Moreover, we are interested in *rectangle* coverings, i.e., the case of I, J being intervals, because (i) they have a more geometric flavour and are thus easier to argue about, (ii) all constructions of non-rectangular coverings that we came up with could be adapted to become rectangular, so we are not aware of any better non-rectangular coverings, and (iii) for rectangular coverings we can show (unconditional) lower bounds.

Designing covering algorithms. In what follows, for ease of exposition we shall assume that $|A| = |B| = n$. Our first algorithmic step is to reduce to a promise version of the problem, where we know $\text{out} = |(A+B) \cap [u]|$ up to a constant factor. (In fact, we could even assume to know a superset of $(A+B) \cap [u]$ whose size is at most a constant factor larger, but we do not know how to exploit such a set in this context.)

The easiest possible covering algorithm partitions $\{1, \dots, n\}$ into k consecutive intervals $I_1, \dots, I_k$ for A , and similarly into $J_1, \dots, J_k$ for B . We may ignore pairs (x, y) with $\min(A_{I_x}) + \min(B_{J_y}) > u$, since they do not contain any sum $A_i + B_j \leq u$. For all remaining (x, y) , we construct the pairs (I_x, J_y) to form a covering C . In the following we discuss two ways of choosing $I_1, \dots, I_k$ and $J_1, \dots, J_k$.

Choosing the interval partitioning such that it splits the universe $[u]$ into parts of length u/k ensures that any sum $s \in A_I + B_J$ for $(I, J) \in C$ satisfies $s \leq (1 + \frac{2}{k})u$. Moreover, we show that any $s \in \mathbb{N}$ appears in at most k sumsets $A_I + B_J, (I, J) \in C$ (this holds because every “diagonal” of subproblems contains distinct sums). In total, we can bound the cost of C by $k \cdot |(A+B) \cap [(1 + \frac{2}{k})u]|$, which for $\zeta = 2/k$ yields one result of Theorem 2.13, see Section 6.

Alternatively, we can choose the interval partitioning such that it splits A and B into parts of size n/k . There are $O(k)$ pairs $(I, J) \in C$ along the “boundary”, where $\min(A_I + B_J) \leq u < \max(A_I + B_J)$. For each such pair we bound the cost trivially by $(\frac{n}{k})^2$. The remaining pairs only cover sums in $(A+B) \cap [u]$, and each such sum is covered at most k times, as in the previous paragraph. In total, this yields cost at most $k \cdot \text{out} + \frac{n^2}{k}$, which is $O(n \cdot \text{out}^{1/2})$ by optimizing over k . Since $n \leq \text{out}$, in particular the cost is at most $O(\text{out}^{3/2})$. Generalizing this idea to the Interval-Restricted case yields Theorem 2.16, see Section 5. For the Interval-Restricted case, this cost bound turns out to be (conditionally) optimal, see Theorem 2.17.

So far, we used simple charging arguments to obtain a covering of cost $O(\text{out}^{3/2})$. We show that for Prefix-Restricted sumset computation we can bypass the hardness results for the Interval-Restricted case and reduce the exponent further to $4/3$. This requires insights from Additive Combinatorics. Specifically, to obtain an $\tilde{O}(\text{out}^{4/3})$ algorithm, we make use of Ruzsa's triangle inequality. This inequality from Additive Combinatorics is a charging argument that implies

$$|A_I + B_J| \leq \frac{|A_I + B_{J'}| |A_{I'} + B_J| |A_{J'} + B_J|}{|I'| |J'|},$$

for any sets $I, I', J, J' \subseteq \{1, \dots, n\}$. In particular, if we choose I' and J' such that the three sumsets on the right hand side are all subsets of $[u]$, then the size of every sumset on the right hand side is bounded from above by $\text{out} = |(A + B) \cap [u]|$. With some care, this inequality can be used to show

$$|A_I + B_J| \leq \frac{\text{out}^3}{\min(I) \min(J)}.$$

This allows us to bound the number of “bad” sums in $A_I + B_J$, namely the sums in $(A_I + B_J) \setminus [u]$, in terms of the output-size and the position of the rectangle (I, J) . However, it turns out that this bound is not good enough, essentially because it yields a worst-case bound that holds for all rectangles (I, J) . We thus replace it with an improved bound that holds for “most” rectangles (I, J) .

Combining this bound with several tricks from our previous covering constructions allows us to construct a covering of cost $\tilde{O}(\text{out}^{4/3})$, see Section 7.

Connection to the Balog-Szemerédi-Gowers Theorem. Probably the most important result on restricted sumsets is the Balog-Szemerédi-Gowers (BSG) theorem [44, Theorem 2.13]. Given sets $A, B \subseteq \mathbb{Z}$ and a set of pairs $G \subseteq A \times B$, the BSG theorem is used to pass from information about the *restricted* sumset $A +_G B = \{a + b \mid (a, b) \in G\}$ to information about an *unrestricted* sumset $A_I + B_J$, for some $I, J \subseteq \{1, \dots, n\}$ (where we continue to assume that $|A| = |B| = n$). Specifically, the BSG theorem states that if $|G| = \delta n^2$ and $|A +_G B| = cn$, then there exist sets $I, J \subseteq \{1, \dots, n\}$ of size $\Omega(\delta^2 n)$ with $|A_I + B_J| \leq O((c/\delta)^5 n)$.

In a breakthrough paper, Chan and Lewenstein [17] algorithmically exploited the BSG theorem in order to solve several problems with additive structure. On a high level, their approach uses the BSG theorem repeatedly, constructing a sequence $(I^{(1)}, J^{(1)}), \dots, (I^{(k)}, J^{(k)})$ of subsets $I^{(i)}, J^{(i)} \subseteq \{1, \dots, n\}$ and a remainder $R \subseteq A \times B$ such that

$$A +_G B = \left(\bigcup_{i=1}^k (A_{I^{(i)}} + B_{J^{(i)}}) \right) \cup \{a + b : (a, b) \in R\} \quad (1)$$

$$(2)$$

For $|A +_G B| = cn$, this decomposition satisfies that all sumsets are small, i.e., $|A_{I^{(i)}} + B_{J^{(i)}}| = O((kc)^5 n)$, and the remainder is not too large, i.e., $|R| = O(n^2/k)$. Using this decomposition, one can compute (a superset of) $A +_G B$ by computing the unrestricted sumsets $A_{I^{(i)}} + B_{J^{(i)}}$ and then iterating over all elements of R . Using near-linear output-sensitive sumset computation, this takes total time $\tilde{O}(k^6 c^5 n + n^2/k)$. For several applications considered in [17], this approach yields non-trivial improvements. We note, however,

that constructing the sequence and the remainder in their setting requires quadratic time in n .

A careful reader could notice that our notion of a *covering* is very similar to the decomposition computed by Chan and Lewenstein [17]. In our situation, we are interested in the set $G = \{(a, b) \in A \times B \mid a + b \leq u\}$, since our goal is to compute $A +_G B = (A + B) \cap [u]$. A covering C satisfies $A +_G B \subseteq \bigcup_{(I, J) \in C} A_I + B_J =: T$, from which we can compute $A +_G B$ as $T \cap [u]$. Except for the remainder set R , this is analogous to Chan and Lewenstein's decomposition. Thus, both notions produce a certain form of covering of $A +_G B$.² Unfortunately, in our situation using their decomposition as a black box is worse than the simple covering of cost $O(\text{out}^{3/2})$ for two reasons. The first is that the time needed to construct their decomposition is $\Theta(|A||B|)$, which could be up to out^2 . The second reason is that the bound they get on the cost of the covering they produce is strictly worse than $\text{out}^{3/2}$. Thus, even if their decomposition was given to us for free, it would be worse than the easy approach which uses elementary charging arguments.

Our approach in this paper can be seen as using the arguments of the proof of the BSG theorem in an ad-hoc manner for our specific set G . Indeed, at the heart of the BSG theorem lies a charging argument that is analogous to Ruzsa's triangle inequality (see [17, Section 7.3]), and careful inspection of the algorithmic version of the BSG theorem (see [17, Sections 2 and 7]) shows that it can be modified to produce a rectangle covering for our set G with no loss in parameters. We present a highly optimized variant of this ad-hoc construction in this paper.

In summary, our algorithms can be viewed as ad-hoc BSG-type theorems for a special choice of the restriction set G , obtaining bounds that seem unreachable in the more general setting.

Connection to Freiman-type Theorems. Another celebrated result among combinatorialists is Freiman's theorem [44, Theorem 5.33]: If $A \subseteq \mathbb{N}$ satisfies $|A + A| \leq c|A|$, then A is contained in an $f(c)$ -dimensional arithmetic progression of length $g(c) \cdot |A|$. Such a result could potentially be useful in our situation, since for a rectangle with $A_I + B_J \subseteq [u]$ we have $|A_I + B_J| \leq \text{out}$, and thus a small output-size could imply that A_I and B_J are essentially arithmetic progressions, which is easy to exploit. However, the state-of-the-art bounds on $f(\cdot), g(\cdot)$ are exponential and thus not useful in applications. This is why much of Additive Combinatorics has focused on bypassing the need for Freiman's theorem, to prove robust theorems with polynomial parameter dependence.

Recently, [43] announced such a robust extension of another theorem by Freiman, called the $3k-4$ theorem, which roughly states that if $|A +_G B|$ is “small” and $|G|$ is “large” ($|G| \geq (1 - \epsilon)|A||B|$), then one can extract highly non-trivial information about A and B . However, the “small” and “large” conditions are too restrictive for our algorithmic applications, and it is unclear to us at the moment to what extent this result can be exploited algorithmically. We expect that further interaction with the field of Additive Combinatorics yields more insights that lead to new algorithmic machinery.

²We remark that Chan and Lewenstein [17] construct their decomposition for a somewhat different reason, since they are not (immediately) concerned with output-sensitivity.

Lower Bound on Rectangle Coverings. Let us also briefly discuss the construction of instances $(A, B, [u])$ on which any rectangle covering has superlinear cost (see Theorem 2.12). We build such instances in two steps.

First, we construct “many” sets $X^{(1)}, \dots, X^{(g)}$ and $Y^{(1)}, \dots, Y^{(g)}$ such that $|X^{(\ell)} + Y^{(\ell)}|$ is “large”, while $|X^{(\ell)} + Y^{(\ell')}|$ is “small” for any $\ell \neq \ell'$. To this end, we pick appropriate parameters m, t , and greedily choose an error-correcting code over $\{0, 1\}^t$ where every codeword has Hamming weight exactly $t/2$. We define $X^{(\ell)}$ to be the set of all t -digit numbers in the m -ary number system with the constraint that their r -th digit is 0 if the ℓ -th codeword has a 1 at position r . Similarly, we define $Y^{(\ell)}$ to be the set of all t -digit numbers in the m -ary system with the constraint that their r -th digit is 0 if the ℓ -th codeword has a 0 at position r .

In the second step, we construct the set A as a union over several shifted copies of each $X^{(\ell)}$. The set B is constructed similarly from the sets $Y^{(\ell)}$. In this construction, we ensure that $(A + B) \cap [u]$ contains no sumset $X^{(\ell)} + Y^{(\ell)}$, so the output-size is “small”. We also ensure that any rectangle (I, J) , that covers many parts of $(A + B) \cap [u]$ at once, needs to cover a sumset $X^{(\ell)} + Y^{(\ell)}$, which results in a “large” sumset $|A_I + B_J|$. Therefore, any covering either contains a rectangle of “large” cost or consists of “many” rectangles; in both cases we obtain a superlinear lower bound on the cost in terms of the output-size $|(A + B) \cap [u]|$.

2.4 Organization

In Section 3 we reduce Top- k -Convolution and related problems to finding a covering, proving Corollary 2.10. In Section 4 we show how to learn the output-size up to a constant factor for restricted convolution problems. In Section 5 we study Interval-Restricted convolution, proving Theorems 2.16 and 2.17. In Section 6, we relax the upper bound, proving part of Theorem 2.13. In Section 7 we describe and analyze our $\tilde{O}(\text{out}^{4/3})$ -cost covering for Prefix-Restricted sumset computation (Theorem 2.9) and we prove a lower bound for rectangle coverings (Theorem 2.12).

3 TOP- k -CONVOLUTION

In this section, we present some easy reductions among Top- k -Convolution and related problems, proving Corollary 2.10. We start by proving that Top- k -Convolution is equivalent to Prefix-Restricted convolution.

LEMMA 3.1. *If Top- k -Convolution on non-negative vectors can be solved in time $\tilde{O}(k^\alpha)$, then Prefix-Restricted convolution on non-negative vectors can be solved in time $\tilde{O}(\text{out}^\alpha)$, and vice versa.*

PROOF. To solve Prefix-Restricted sparse convolution, we run Top- k -Convolution for $k = 2^0, 2^1, 2^2, \dots$ until we reach a value of k so that the k -th non-zero entry lies above u . Then we have found all elements of the Prefix-Restricted convolution, and we only computed at most twice as many non-zero entries as necessary.

To solve Top- k -Convolution on vectors $f, g \in \mathbb{R}_{\geq 0}^d$, we perform binary search over $\{0, 1, \dots, 2d\}$, to find the smallest u such that Prefix-Restricted convolution on $[u]$ returns a k -sparse vector. To obtain the desired running time we add a small twist: if the execution time of Prefix-Restricted sumset computation on $[u]$ is more

than $k^\alpha \cdot \text{polylog } d$, i.e., more than what it would be for output size k , then we abort and look for a smaller u .

Both reductions only add a log-factor to the running time. $\square$

We next show that coverings not only allow us to compute Prefix-Restricted sumsets, but they even enable the computation of Prefix-Restricted convolutions, if the covering is *unique*.

LEMMA 3.2. *Suppose that, given (A, B, u) , we can compute a unique covering of $(A, B, [u])$ with cost $\tilde{O}(\text{out}^\alpha)$ in expected time $\tilde{O}(\text{out}^\alpha)$, where $\text{out} := |(A + B) \cap [u]|$. Then Prefix-Restricted convolution on non-negative vectors can be solved in expected time $\tilde{O}(\text{out}^\alpha)$.*

PROOF. For a vector f we denote by f_S the same vector where every entry outside of S is zeroed out. Given non-negative vectors f, g , let $A := \text{supp}(f)$ and $B := \text{supp}(g)$. The output-size of Prefix-Restricted convolution on f, g is the number of non-zero entries of $(f \star g)_{[u]}$. This is the same as the output-size of Prefix-Restricted sumset computation on A, B , namely $|(A + B) \cap [u]|$. Thus, the two output-sizes coincide and we denote both by “out”. In particular, we can afford to compute a unique covering C of $(A, B, [u])$.

Using output-sensitive convolution of non-negative vectors (Theorem 2.1, [19]), we can compute $f_S \star g_T$ in time $\tilde{O}(\|f_S \star g_T\|_0)$. Therefore, we can compute

$$h := \sum_{(I, J) \in C} f_{A_I} \star g_{B_J},$$

in time proportional to the cost of the covering, up to log factors. By the properties of a unique covering, any non-zero product $f_j \cdot g_{i-j}$ appears exactly once in the definition of h , for any $0 \leq j \leq i \leq u$. We thus have $h_{[u]} = (f \star g)_{[u]}$. $\square$

Corollary 2.10 now follows.

PROOF OF COROLLARY 2.10. To solve Prefix-Restricted sumset computation in time $\tilde{O}(\text{out}^{4/3})$, we combine our covering construction (Theorem 2.9) with Observation 2.8. To solve Prefix-Restricted convolution on non-negative vectors in time $\tilde{O}(\text{out}^{4/3})$, we combine our covering construction (Theorem 2.9) with Lemma 3.2. To solve Top- k -Convolution on non-negative vectors in time $\tilde{O}(k^{4/3})$, we combine the result for Prefix-Restricted convolution with the equivalence shown in Lemma 3.1. $\square$

4 RESTRICTED SUMSET COMPUTATION: LEARNING THE OUTPUT-SIZE

In this section, we show that we can assume to know the output-size up to a constant factor. We will later use this for Prefix-Restricted and for Interval-Restricted sumset computation. Here we will discuss the more general setting of Interval-Restricted sumset computation; the same construction works for the Prefix-Restricted case.

Suppose that we are given sets A, B and integers ℓ, u and we want to compute their Interval-Restricted sumset $(A + B) \cap [\ell, u]$. We show that we can assume to know $|(A + B) \cap [\ell, u]|$ up to a constant factor. In fact, we reduce Interval-Restricted sumset computation to a seemingly much easier variant, where additionally we are given a set T of size $\Theta(|(A + B) \cap [\ell, u]|)$ which contains $(A + B) \cap [\ell, u]$.

DEFINITION 4.1 (INTERVAL-RESTRICTED SUMSET COMPUTATION WITH A PROMISE (IR-SMP)). *Given sets $A, B \subseteq \mathbb{Z}$, numbers $\ell \leq u$, and a set T of size $|T| \leq 6|(A+B) \cap [\ell, u]| + 9$, such that*

$$(A+B) \cap [\ell, u] \subseteq T,$$

compute the set

$$(A+B) \cap [\ell, u].$$

We present a reduction from the problem variant without promise to the variant with promise.

LEMMA 4.2. *We can reduce a given instance (A, B, ℓ, u) of Interval-Restricted sumset computation to $O(\log(u-\ell))$ instances of IR-SMP, each of the form (A', B', ℓ', u', T') with $|A'| \leq |A|$, $|B'| \leq |B|$, $\ell' \leq \ell$, $u' \leq u$, $u' - \ell' \leq u - \ell$, and $|T'| = O(|(A+B) \cap [\ell, u]|)$. The reduction runs in linear time in its output-size.*

PROOF. In what follows, for integers x, y we define $x \div y = \lfloor \frac{x}{y} \rfloor$. For a set Z we write $Z \div x = \{z \div x \mid z \in Z\}$, and we write $xZ = \{x \cdot z \mid z \in Z\}$.

Let (A, B, ℓ, u) be a given instance of Interval-Restricted sumset computation. We compute better and better approximations of the desired output $(A+B) \cap [\ell, u]$ by computing the sets

$$S^{(i)} := ((A \div 2^i) + (B \div 2^i)) \cap [\ell \div 2^i, u \div 2^i],$$

for $i = r, \dots, 0$, where $r := \lceil \log(u-\ell) \rceil$. See Algorithm 1 for pseudocode. Note that $S^{(0)}$ is the desired output $(A+B) \cap [\ell, u]$. We compute each set $S^{(i)}$ by calling an IR-SMP oracle on a suitable promise $T^{(i)}$, determined as follows.

Algorithm 1 Reduction of Interval-Restricted Sumset Computation to its Promise Version

```

1: procedure COMPUTEINTERVALRESTRICTEDSUMSET( $A, B, \ell, u$ )
2:    $\triangleright A, B \subseteq [u], \ell \leq u$ 
    $\triangleright$  We assume oracle access to Interval-Restricted sumset
       computation with a promise (IR-SMP)
3:    $r \leftarrow \lceil \log(u-\ell) \rceil$ 
4:    $T^{(r)} \leftarrow \{\ell \div 2^r, \dots, u \div 2^r\}$ 
5:    $S^{(r)} \leftarrow \text{IR-SMP}(A \div 2^r, B \div 2^r, \ell \div 2^r, u \div 2^r, T^{(r)})$ 
6:    $\triangleright$  Call to the promise problem
7:   for  $i = r-1$  down to  $0$  do
8:      $T^{(i)} \leftarrow (2S^{(i+1)} + \{0, 1, 2\})$ 
9:      $T^{(i)} \leftarrow T^{(i)} \cup \{(\ell \div 2^i), (\ell \div 2^i) + 1, (u \div 2^i)\}$ 
10:     $S^{(i)} \leftarrow \text{IR-SMP}(A \div 2^i, B \div 2^i, \ell \div 2^i, u \div 2^i, T^{(i)})$ 
11:     $\triangleright$  Call to the promise problem
12:   Return  $S^{(0)}$ 

```

For $i = r$, the interval $[\ell \div 2^i, u \div 2^i]$ has length at most 2. Thus, the set $T^{(r)} := \{\ell \div 2^r, \dots, u \div 2^r\}$ is a valid promise from which our assumed IR-SMP oracle can compute the set $S^{(r)}$.

For $i < r$, we claim that a valid promise for the instance $(A \div 2^i, B \div 2^i, \ell \div 2^i, u \div 2^i)$ is given by

$$T^{(i)} := (2S^{(i+1)} + \{0, 1, 2\}) \cup \{(\ell \div 2^i), (\ell \div 2^i) + 1, (u \div 2^i)\}.$$

If this claim holds, then from $S^{(r)}$ we can compute $S^{(r-1)}, S^{(r-2)}, \dots, S^{(0)}$, finishing our reduction.

Let us also claim that any set $S^{(i)}$ has size $|S^{(i)}| = O(|(A+B) \cap [\ell, u]|)$. Then in particular the set $T^{(i)}$ has size $|T^{(i)}| = O(|S^{(i)}|) =$

$O(|(A+B) \cap [\ell, u]|)$. Hence, all constructed instances satisfy the claimed properties, and we reduced the given instance (A, B, ℓ, u) to $O(\log(u-\ell))$ calls to an IS-SMP oracle. This finishes the proof, except that it remains to prove the two claims.

CLAIM 4.3. *The set $S^{(i)}$ has size $|S^{(i)}| \leq 2|(A+B) \cap [\ell, u]| + 2$.*

PROOF. It is easy to check that for any integers a, b we have

$$2^i((a \div 2^i) + (b \div 2^i)) \leq a + b \leq 2^i((a \div 2^i) + (b \div 2^i)) + 2^i, \quad (3)$$

or, equivalently,

$$\frac{a+b}{2^i} - 1 \leq (a \div 2^i) + (b \div 2^i) \leq \frac{a+b}{2^i}. \quad (4)$$

Any number in $s \in S^{(i)}$ can be expressed as $s = (a \div 2^i) + (b \div 2^i)$ with

$$(\ell \div 2^i) \leq s \leq (u \div 2^i).$$

Suppose that we have $s \in [(\ell \div 2^i) + 1, (u \div 2^i) - 1]$. Then by inequalities (3), it follows that

$$\ell \leq 2^i((\ell \div 2^i) + 1) \leq 2^i s \leq a + b \leq 2^i s + 2^i \leq 2^i(u \div 2^i) \leq u.$$

Thus, we obtain that the number $s' := a + b$ lies in $(A+B) \cap [\ell, u]$.

We say that s charges s' . From inequalities (4), we infer that any number $s' \in (A+B) \cap [\ell, u]$ can only be charged by numbers in $[s'/2^i - 1, s'/2^i]$, so it is charged by at most two numbers $s \in S^{(i)}$. This charging scheme proves

$$|S^{(i)} \cap [(\ell \div 2^i) + 1, (u \div 2^i) - 1]| \leq 2 \cdot |(A+B) \cap [\ell, u]|.$$

We add 2 to account for the numbers $\ell \div 2^i$ and $u \div 2^i$. This yields the claim. $\square$

CLAIM 4.4. *The set $T^{(i)}$ is a valid promise for the instance $(A \div 2^i, B \div 2^i, \ell \div 2^i, u \div 2^i)$. That is, we have $T^{(i)} \supseteq S^{(i)}$ and $|T^{(i)}| \leq 6|S^{(i)}| + 9$.*

PROOF. Since $x \div 2^{i+1} = (x \div 2^i) \div 2$, it suffices to prove the claim for $i = 0$. The same proof then also works for larger i after replacing A by $A \div 2^i$, B by $B \div 2^i$, ℓ by $\ell \div 2^i$, and u by $u \div 2^i$.

We first show that $|T^{(0)}| \leq 6|S^{(0)}| + 9$. Recall that in Claim 4.3 we proved that

$$|S^{(1)}| \leq 2 \cdot |(A+B) \cap [\ell, u]| + 2 = 2|S^{(0)}| + 2.$$

We combine this with the inequality $|T^{(0)}| \leq 3|S^{(1)}| + 3$ that we obtain from the construction of $T^{(0)}$. This yields the claimed inequality $|T^{(0)}| \leq 6|S^{(0)}| + 9$.

Next we prove $T^{(0)} \supseteq S^{(0)}$. Consider any $a \in A, b \in B$ with $a + b \in [\ell + 2, u - 1]$. Using inequalities (4) we obtain

$$\begin{aligned} \ell \div 2 &\leq \frac{\ell}{2} \leq \frac{a+b}{2} - 1 \leq (a \div 2) + (b \div 2) \\ &\leq \frac{a+b}{2} \leq \frac{u-1}{2} \leq u \div 2. \end{aligned}$$

Therefore, the sum $(a \div 2) + (b \div 2)$ is in $S^{(1)}$. Now, since $a + b$ can be found among the three numbers

$$2((a \div 2) + (b \div 2)), \quad 2((a \div 2) + (b \div 2)) + 1, \quad 2((a \div 2) + (b \div 2)) + 2,$$

it follows that $a + b \in (2S^{(1)} + \{0, 1, 2\}) \subseteq T^{(0)}$. Recall that here we assumed $a + b \in [\ell + 2, u - 1]$. The boundary numbers $\ell, \ell + 1$,

and u are handled by explicitly adding them to the set $T^{(0)}$ in its construction. We thus obtain $S^{(0)} \subseteq T^{(0)}$. $\square$

These claims finish the proof of Lemma 4.2. $\square$

We obtain the following easy corollary where we essentially ignore the superset T and only keep its size $|T|$, which is a constant-factor approximation of the output-size.

LEMMA 4.5 (INTERVAL-RESTRICTED SUMSET COMPUTATION WITH APPROXIMATE OUTPUT-SIZE). *Suppose that given (A, B, ℓ, u) and an additional input $\widetilde{\text{out}}$ satisfying $\text{out} \leq \widetilde{\text{out}} \leq 6\text{out} + 9$, we can compute a covering of $(A, B, [\ell, u])$ of cost $O(c)$ in time $O(T)$, where c and T are monotone functions of $|A|, |B|, u, \text{out}$. Then Interval-Restricted sumset computation can be solved in time $\widetilde{O}(c + T)$.*

An analogous statement holds for Prefix-Restricted sumset computation.

PROOF. Given an instance (A, B, ℓ, u) of Interval-Restricted sumset computation, we run Lemma 4.2 to reduce to $O(\log(u - \ell))$ promise instances of the form (A', B', ℓ', u', T') . By the properties of T' , for $\widetilde{\text{out}}' := |T'|$ we have $\text{out}' \leq \widetilde{\text{out}}' \leq 6\text{out}' + 9$, where $\text{out}' := |(A' + B') \cap [\ell', u']|$. Hence, we can use our assumed algorithm to compute a covering of $(A', B', [\ell', u'])$ of cost $O(c)$ in time $O(T)$. By Observation 2.8, we can thus solve the instance $(A', B', [\ell', u'])$ in time $\widetilde{O}(c + T)$. Over $O(\log(u - \ell))$ many constructed instances, we obtain the same time bound up to log-factors. $\square$

5 INTERVAL-RESTRICTED SUMSET COMPUTATION

In this subsection we prove Theorems 2.16 and 2.17.

5.1 An $\widetilde{O}(\sqrt{mn} \cdot \text{out})$ -Time Algorithm

In this subsection we prove Theorem 2.16.

PROOF. Given A, B, ℓ, u , we will compute the Interval-Restricted sumset $(A + B) \cap [\ell, u]$ in time $\widetilde{O}(n + m + \sqrt{mn} \cdot \text{out})$. More precisely, we will compute a covering of $(A, B, [\ell, u])$ of cost $O(\sqrt{mn} \cdot \text{out})$ in time $O(n + m)$. By Lemma 4.5, we can assume to know an approximation $\widetilde{\text{out}}$ of the output-size, satisfying $\text{out} \leq \widetilde{\text{out}} \leq O(\text{out})$.

Set $q := \lceil (nm/\widetilde{\text{out}})^{1/2} \rceil$. Note that since $n, m \leq \text{out} \leq n \cdot m$ we have $1 \leq q = O(\min\{n, m\})$. We assume that n and m are divisible by q , which we can ensure by duplicating at most q elements in A and B (with the slight abuse of transitioning to multi-sets).

We split A and B into q subsets by defining for any $1 \leq i, j \leq q$:

$$\begin{aligned} I_i &:= \{(i-1) \cdot n/q + 1, \dots, i \cdot n/q\}, \\ J_j &:= \{(j-1) \cdot m/q + 1, \dots, j \cdot m/q\}, \\ A^{(i)} &:= A_{I_i}, \\ B^{(j)} &:= B_{J_j}. \end{aligned}$$

We let C be the set of all pairs (I_i, J_j) with

$$\min(A^{(i)}) + \min(B^{(j)}) \leq u \quad \text{and} \quad \max(A^{(i)}) + \max(B^{(j)}) \geq \ell.$$

Observe that C is a unique-rectangle-covering of $(A, B, [\ell, u])$. We can easily compute C by brute force, by iterating over all $1 \leq i, j \leq q$

and testing whether to put (I_i, J_j) into C in time $O(1)$, see Algorithm 2. This takes total time $O(q^2) = O(nm/\text{out}) = O(\min\{n, m\})$, assuming that we have random access to A and B .

Algorithm 2 Covering for Interval-Restricted Sumset Computation

```

1: procedure FINDCOVERING( $A, B, \ell, u, \widetilde{\text{out}}$ )
2:    $q \leftarrow \lceil (nm/\widetilde{\text{out}})^{1/2} \rceil$ 
3:    $C \leftarrow \emptyset$ 
4:   for  $i = 1$  to  $q$  do
5:     for  $j = 1$  to  $q$  do
6:        $I \leftarrow \{(i-1) \cdot n/q + 1, \dots, i \cdot n/q\}$ 
7:        $J \leftarrow \{(j-1) \cdot m/q + 1, \dots, j \cdot m/q\}$ 
8:       if  $\min(A_{I_i}) + \min(B_{J_j}) \leq u$  then
9:         if  $\max(A_{I_i}) + \max(B_{J_j}) \geq \ell$  then
10:            $C \leftarrow C \cup \{(I, J)\}$ 
11:   Return  $C$ 

```

Recall that the cost of a covering C is

$$\sum_{(I, J) \in C} |A_I + B_J|.$$

We split C into two parts, the rectangles in the interior and the rectangles at the boundary:

$$\begin{aligned} C_{\text{int}} &:= \{(I, J) \in C \mid A_I + B_J \subseteq [\ell, u]\}, \\ C_{\text{bd}} &:= C \setminus C_{\text{int}}. \end{aligned}$$

For the interior rectangles, we split their cost into diagonal sums of the form

$$\sum_{(I_i, J_{i+\Delta}) \in C_{\text{int}}} |A^{(i)} + B^{(i+\Delta)}|,$$

for $-q < \Delta < q$. We claim that each such diagonal sum is bounded from above by out . Indeed, for consecutive terms along a diagonal we have

$$\max(A^{(i)}) + \max(B^{(i+\Delta)}) < \min(A^{(i+1)}) + \min(B^{(i+\Delta+1)}).$$

Therefore, the output-sizes $|A^{(i)} + B^{(i+\Delta)}|$ are disjoint contributions to out , for fixed Δ and ranging over all i . It follows that a diagonal sum is bounded by out , and since there are $2q - 1$ diagonals, we obtain

$$\sum_{(I, J) \in C_{\text{bd}}} |A_I + B_J| \leq (2q - 1)\text{out} = O(\sqrt{nm} \cdot \text{out}).$$

We argue geometrically about the boundary. Observe that C_{bd} contains at most two rectangles per diagonal, which yields $|C_{\text{bd}}| \leq 4q$. For each $(I, J) \in C_{\text{bd}}$ we use the trivial upper bound $|A_I + B_J| \leq |A_I| \cdot |B_J| = nm/q^2$. In total, this yields cost

$$\sum_{(I, J) \in C_{\text{bd}}} |A_I + B_J| \leq 4q \cdot \frac{nm}{q^2} = O(\sqrt{nm} \cdot \text{out}).$$

The contribution from both parts is the same, so in total we bounded the cost of C by $O(\sqrt{nm} \cdot \text{out})$. This finishes the proof. $\square$

5.2 Hardness Results

In this subsection we prove Theorem 2.17.

PROOF. We want to prove hardness of Interval-Restricted sumset computation. Note that here we analyze running time in terms of $n = |A|$, $m = |B|$, and out , which are all invariant under shifting A by adding a number q to each $a \in A$; similarly for B . Therefore, we may drop the assumption that A, B are sets of positive integers and let them be subsets of $\mathbb{Z}$ instead. This is the case because we can shift A, B and the interval appropriately, so that every number in the input becomes non-negative.

Reduction from Boolean Matrix Multiplication to Interval-Restricted Sumset Computation: In Boolean matrix multiplication we are given $n \times n$ matrices $\bar{A}, \bar{B}$ with entries in $\{0, 1\}$ and want to compute their product C with $C_{ij} = \bigvee_r \bar{A}_{ir} \wedge \bar{B}_{rj}$.

Given matrices $\bar{A}, \bar{B}$, we construct sets A, B as

$$A := \{rM^2 + \bar{A}_{ir} \cdot M + i \mid i, r \in [n]\},$$

$$B := \{-rM^2 + \bar{B}_{rj} \cdot M + nj \mid r, j \in [n]\},$$

where M is any integer greater than $10(n^2 + n)$. We also set

$$\ell := 2M + n + 1, \quad u := 2M + n^2 + n.$$

We observe the following.

- (1) Every integer of the form $(\bar{A}_{ir} + \bar{B}_{rj}) \cdot M + (i + nj)$ with $i, r, j \in [n]$ is contained in $A + B$.
- (2) For $r \neq r'$ any sum $(rM^2 + \bar{A}_{ir} \cdot M + i) + (-r'M^2 + \bar{B}_{r'j} \cdot M + nj)$ is either less than $-M^2 + 2M + n^2 + n < 0$ or at least $M^2 + n + 1$, and hence outside of $[\ell, u]$.
- (3) If $\bar{A}_{ir} \wedge \bar{B}_{rj} = 1$, then $(\bar{A}_{ir} + \bar{B}_{rj})M + (i + nj) = 2M + (i + nj)$.
- (4) If $\bar{A}_{ir} \wedge \bar{B}_{rj} = 0$, then $(\bar{A}_{ir} + \bar{B}_{rj})M + (i + nj) \leq M + (i + nj) < 2M$.

It follows that from $(A + B) \cap [2M + n + 1, 2M + n^2 + n]$ we can infer all entries of the product matrix C .

Note that the output-size is $\text{out} \leq u - \ell + 1 = n^2$. Hence, for any $\delta > 0$, an $O((|A| + |B| + \text{out})^{\omega/2 - \delta})$ -time algorithm for Interval-Restricted sumset computation would yield an $O(n^{2\omega - 2\delta})$ -time algorithm for Boolean matrix multiplication.

Reduction from Sliding Window Hamming Distance to Interval-Restricted Convolution: Note that Interval-Restricted convolution allows us to not only to compute $(A + B) \cap [\ell, u]$, but also the number of ways an element $x \in (A + B) \cap [\ell, u]$ can be written as $x = a + b$ with $a \in A, b \in B$; let us call this number the multiplicity of x . Indeed, for the indicator vectors of A and B , the x -th entry of their convolution is the multiplicity of x in $A + B$.

In the sliding window Hamming distance problem we are given a text t of length $2n$ and a pattern p of length n and want to compute the Hamming distance between the pattern and every length- n substring of the text. Given such an instance t, p , we construct sets A, B as

$$A := \{M \cdot t_i + i \mid 1 \leq i \leq 2n\}, \quad B := \{-M \cdot p_j - j \mid 1 \leq j \leq n\},$$

where $M := 100n$. We also set $\ell := 1$, $u := n$ and compute, as mentioned in the previous paragraph, the Interval-Restricted sumset $(A + B) \cap [\ell, u]$ as well as the multiplicity of every x in this set.

Fix a $1 \leq i \leq n$ and observe that

- (1) If $t_{i+j} = p_j$ then the pair $(i + j, j)$ will contribute 1 to the multiplicity of i .
- (2) If $t_{i+j} \neq p_j$ then the pair $(i + j, j)$ will contribute 1 to the multiplicity of a coefficient outside of the interval $[1, n]$, by the choice of M .
- (3) Every pair (i', j) with $i' \leq j$ contributes 1 to a coefficient outside of the interval $[1, n]$.

It follows that we can read off the Hamming distance between the pattern p and the i -th length- n substring of t from the multiplicity of i in the output. This completes the reduction from sliding window Hamming distance to prefix-restricted convolution. Since $\text{out} \leq n$, any $\tilde{O}(|A| + |B| + (|A| \cdot |B| \cdot \text{out})^{1/2 - \delta})$ -time algorithm for Interval-Restricted convolution would solve sliding window Hamming distance in time $\tilde{O}(n^{3/2 - 3\delta})$. $\square$

6 RELAXED VERSION OF PREFIX-RESTRICTED CONVOLUTION

We show how to solve Prefix-Restricted convolution on any instance (A, B, u) in time $\tilde{O}(\zeta^{-1} |(A + B) \cap [u(1 + \zeta u)]| + \zeta^{-2})$, for any $\zeta \leq 1$, proving Theorem 2.13. More precisely, we prove the following theorem, from which we conclude Theorem 2.13 using Observation 2.8.

THEOREM 6.1. *Given sets A, B and a target u we can compute a unique-rectangle-covering of $(A, B, [u])$ with cost*

$$O(\zeta^{-1} |(A + B) \cap [(1 + \zeta)u]|)$$

in time $O(|A| + |B| + \zeta^{-2})$.

PROOF. Find the smallest ℓ such that $2^{-\ell} \leq \zeta$. It suffices to solve the problem for $\zeta = 2^{-\ell}$, since $(A + B) \cap [u + 2^{-\ell}u] \subseteq (A + B) \cap [u + \zeta u]$, and $2^\ell \leq 2\zeta^{-1}$. Thus, we can assume that $1/\zeta$ is a power of 2. Moreover, by shifting u as well as A by a suitable number we can assume that u is a power of 2. We split the set $[u]$ into the intervals

$$U_r := [(r - 1) \cdot (\zeta u)/2 + 1, r \cdot (\zeta u)/2].$$

Moreover, we define the following sets for $1 \leq i, j \leq 2/\zeta$:

$$A^{(i)} := A \cap U_i,$$

$$B^{(j)} := B \cap U_j,$$

$$I_i := \{i' \in [n] : A_{i'} \in U_i\}$$

$$J_j := \{j' \in [m] : B_{j'} \in U_j\}.$$

Let C be the set of all pairs (I_i, J_j) for which $(A^{(i)} + B^{(j)}) \cap [u]$ is not empty; this condition can be easily checked by testing whether

$$\min(A^{(i)}) + \min(B^{(j)}) \leq u.$$

Note that C can be computed in time $O(\zeta^{-2})$, assuming that we have random access to A and B . Observe that C is indeed a unique-rectangle-covering of $(A + B) \cap [u]$. Recall that the cost of C is

$$\sum_{(I, J) \in C} |A_I + B_J|.$$

Similarly to the argument in the proof of Theorem 2.16, this sum can be decomposed into $4/\zeta$ diagonal sums of the form

$$\sum_i |A^{(i)} + B^{(i+\Delta)}|,$$

where the sum is over all i such that $(I_i, J_{i+\Delta}) \in C$. We again use

$$\max(A^{(i)}) + \max(B^{(j)}) < \min(A^{(i+1)}) + \min(B^{(j+1)}).$$

Moreover, for any $(i, j) \in P$ we now have

$$\begin{aligned} \max(A^{(i)}) + \max(B^{(j)}) &\leq \left(\min(A^{(i)}) + \frac{\zeta u}{2} \right) + \left(\min(B^{(j)}) + \frac{\zeta u}{2} \right) \\ &= \left(\min(A^{(i)}) + \min(B^{(j)}) \right) + \zeta u \\ &\leq u + \zeta u. \end{aligned}$$

It follows that every diagonal sum contributes at most

$$|(A+B) \cap [u + \zeta u]|.$$

Summing over all diagonals, in total we can bound the cost of C by $O(\zeta^{-1} |(A+B) \cap [u + \zeta u]|)$. $\square$

7 CONSTRUCTION OF THE $\widetilde{O}(\text{out}^{4/3})$ -COST COVERING

This section is devoted to proving the technical core of our Prefix-Restricted sumset algorithm, specifically we prove Theorem 2.9.

7.1 An Additive Combinatorics Ingredient: Ruzsa's Triangle Inequality

The following is a classical result from Additive Combinatorics. We present a self-contained proof.

LEMMA 7.1 (RUZSA'S TRIANGLE INEQUALITY, SEE ALSO [16, THEOREM 2]). *For any $A, B, C \subseteq \mathbb{Z}$ we have*

$$|A - B| \leq \frac{|A - C| \cdot |C - B|}{|C|}.$$

PROOF. We associate every $s \in A - B$ with the lexicographically smallest pair $(a, b) \in A \times B$ such that $s = a - b$, and we denote this pair by $(a(s), b(s))$. Consider the mapping

$$\begin{aligned} (A - B) \times C &\rightarrow (A - C) \times (C - B) \\ (s, c) &\mapsto (a(s) - c, c - b(s)) \end{aligned}$$

We claim that this mapping is injective. Indeed, from an image $(x, y) = (a(s) - c, c - b(s))$ we can infer $s = a(s) - b(s) = x + y$. The value s then determines $a(s)$ and $b(s)$, so we can infer $c = y + b(s)$. We thus recovered the corresponding preimage (s, c) .

Since this mapping is injective, we obtain $|A - B| \cdot |C| \leq |A - C| \cdot |C - B|$. $\square$

We will use the following simple corollary.

LEMMA 7.2 (COROLLARY OF RUZSA'S TRIANGLE INEQUALITY). *For any $X, Y, Z, W \subseteq \mathbb{Z}$ we have*

$$|X + Y| \leq \frac{|X + Z| \cdot |Z + W| \cdot |W + Y|}{|Z| \cdot |W|}.$$

PROOF. First use Ruzsa's triangle inequality on $A = X$, $B = -Y$, $C = -Z$ to obtain

$$|X + Y| \leq \frac{|X + Z| \cdot |Z - Y|}{|Z|}.$$

Then use Ruzsa's triangle inequality on $A = Z$, $B = Y$, $C = -W$ to obtain

$$|Z - Y| \leq \frac{|Z + W| \cdot |W + Y|}{|W|}.$$

Plugging the latter into the former proves the claim. $\square$

7.2 Description of the Algorithm

Given $A, B \subseteq [u]$, we write $n = |A|$, $m = |B|$, and $\text{out} = |(A+B) \cap [u]|$. We describe an algorithm that computes a unique rectangle covering C of $(A, B, [u])$ of cost $\widetilde{O}(\text{out}^{4/3})$ in time $\widetilde{O}(\text{out}^{4/3})$. However, as a subroutine we will use output-sensitive sumset computation (Theorem 2.2), which shows that it is hard to completely separate the tasks of computing a covering and computing the Prefix-Restricted sumset itself.

Invoking Lemma 4.5, it suffices to solve the promise problem where we are given a value out guaranteed to satisfy $\text{out} \leq \widetilde{\text{out}} \leq O(\text{out})$. We can assume that out is larger than some absolute constant, since otherwise we have $|A|, |B| \leq \text{out} \leq \widetilde{\text{out}} = O(1)$, so we can compute a trivial covering of cardinality 1 and cost $|A| \cdot |B| = O(1)$.

We maintain families $\mathcal{C}, \mathcal{D}$, initialized to $\mathcal{C} = \emptyset$ and $\mathcal{D} = \{([n], [m])\}$, with the invariant that $\mathcal{C} \cup \mathcal{D}$ is a unique rectangle covering of $(A, B, [u])$. We refer to the rectangles in $\mathcal{D}$ as the *unprocessed subproblems*, or simply *subproblems*. The algorithm is finished when there are no more unprocessed subproblems, i.e., $\mathcal{D} = \emptyset$, and then we return the unique rectangle covering $\mathcal{C}$ as output.

We associate to every subproblem $(I, J) \in \mathcal{D}$ the *type* (x, y) for $x = \lceil \log |I| \rceil$ and $y = \lceil \log |J| \rceil$. Initially, there is exactly one subproblem $([n], [m])$ of type $(\lceil \log n \rceil, \lceil \log m \rceil)$.

We define a total order on types: $(x, y) < (x', y')$ iff $x+y < x'+y'$, or $x+y = x'+y'$ and $x < x'$. Our algorithm processes types in descending order according to this total order. For any type (x, y) , we process all subproblems of type (x, y) in one batch. Upon processing a subproblem, our algorithm may generate further subproblems of strictly smaller type. As we will see below, all subproblems of the same type that we generate are disjoint, i.e., for any subproblems $(I, J), (I', J')$ of the same type we have $I \cap I' = \emptyset$ and $J \cap J' = \emptyset$.

Since we want to process all subproblems of a particular type (x, y) in one batch, we need to store the family $\mathcal{D}$ in such a way that we can efficiently enumerate all subproblems of a type (x, y) . To this end, we store $\mathcal{D}$ in a standard data structure such as a self-adjusting binary search tree, where subproblems are first compared according to their type and then according to the endpoints of I and J . Note that we can store any subproblem (I, J) using $O(1)$ integers, since I and J are intervals.

We set the parameter $q := \lceil \widetilde{\text{out}}^{1/3} \rceil$.

To finish the description of the algorithm, it remains to describe how we process all subproblems of a particular type (x, y) in one batch. We consider two cases.

Case 1: $2^{x+y} \leq \widetilde{\text{out}}$. Then for each subproblem (I, J) of type (x, y) we move (I, J) from $\mathcal{D}$ to $\mathcal{C}$.

Case 2: $2^{x+y} > \widetilde{\text{out}}$. If $\mathcal{D}$ contains more than q subproblems of type (x, y) , then we do the following. We start computing $A_I + B_J$ for each of these subproblems in parallel (using Theorem 2.2), and we stop once all but q of these calls have finished. For each finished call $A_I + B_J$, we move (I, J) from $\mathcal{D}$ to $\mathcal{C}$.

At this point, we have at most q subproblems of type (x, y) left. We split each such subproblem $(I, J) = ([i_1, i_2], [j_1, j_2])$ as follows. We set $i := \lfloor (i_1 + i_2)/2 \rfloor$ and determine the maximum index $j \in J$ with $A_i + B_j \leq u$. This splits I into $I_1 = [i_1, i]$ and $I_2 = [i+1, i_2]$ and

Algorithm 3

```

1: procedure COVERINGCONSTRUCTION( $A, B, u, \widetilde{\text{out}}$ )
2:    $n \leftarrow |A|, m \leftarrow |B|, q \leftarrow \widetilde{\text{out}}^{1/3}$ 
3:   Initialize  $C \leftarrow \emptyset, \mathcal{D} \leftarrow \{([n], [m])\}$ 
4:   while  $\mathcal{D} \neq \emptyset$  do
5:     Let  $(x, y)$  be the largest type such that  $\mathcal{D}$  contains sub-
     problems of type  $(x, y)$ 
6:     // process all subproblems of type  $(x, y)$  in one batch:
7:     if  $2^{x+y} \leq \widetilde{\text{out}}$  then
8:       for each  $(I, J) \in \mathcal{D}$  of type  $(x, y)$ : Move  $(I, J)$  from
        $\mathcal{D}$  to  $C$ 
9:     else
10:      if  $\mathcal{D}$  contains more than  $q$  subproblems of type
       $(x, y)$  then
11:        Compute  $A_I + B_J$  in parallel for all  $(I, J) \in \mathcal{D}$ 
        of type  $(x, y)$ 
12:        Stop once all but  $q$  of these calls finished
13:        for each finished call  $A_I + B_J$ : Move  $(I, J)$  from
         $\mathcal{D}$  to  $C$ 
14:        for each remaining  $(I, J) \in \mathcal{D}$  of type  $(x, y)$  do
15:          Split  $I = [i_1, i_2]$  at  $i = \lfloor (i_1 + i_2)/2 \rfloor$  into  $I_1$  and  $I_2$ 
16:          Determine the maximum  $j \in J$  with  $A_i + B_j \leq u$ 
17:          Split  $J$  at  $j$  into  $J_1$  and  $J_2$ 
18:          Add  $(I_1, J_1)$  to  $C$ 
19:          Add  $(I_1, J_2), (I_2, J_1)$  to  $\mathcal{D}$ 
20:          Remove  $(I, J)$  from  $\mathcal{D}$ 
21:   return  $C$ 

```

J into $J_1 = [j_1, j]$ and $J_2 = [j + 1, j_2]$. We add (I_1, J_1) to the output C . Moreover, we add the subproblems (I_1, J_2) and (I_2, J_1) to $\mathcal{D}$ and we remove (I, J) from $\mathcal{D}$. Since $u < A_i + B_{j+1} \leq A_{i+1} + B_{j+1}$, we can ignore the subproblem (I_2, J_2) .

This finishes the description of our algorithm, for pseudocode see Algorithm 3.

7.3 Analysis

The correctness of the algorithm, meaning that the output is a unique rectangle covering, follows from the next claim and the fact that the algorithm stops when $\mathcal{D} = \emptyset$.

CLAIM 7.3 (INVARIANT THAT $C \cup \mathcal{D}$ IS A COVERING). *At any point during the algorithm, the family $C \cup \mathcal{D}$ is a unique rectangle covering of $(A, B, [u])$.*

PROOF. Moving (I, J) from $\mathcal{D}$ to C does not change this property. Therefore, the only crucial step is the splitting of I into I_1, I_2 and of J into J_1, J_2 , where we add (I_1, J_1) to C and add $(I_1, J_2), (I_2, J_1)$ to $\mathcal{D}$. Observe that at this point we have $u < A_i + B_{j+1} \leq A_{i+1} + B_{j+1}$, and thus the rectangle (I_2, J_2) does not contain any sum below u , so it is unnecessary for a covering. It follows that $C \cup \mathcal{D}$ remains a covering. Moreover, noting that I_1, I_2, J_1, J_2 are again intervals, it remains rectangular. Finally, since no pair (i, j) is contained in two subproblems $I_b \times J_{b'}$, it remains unique. $\square$

We will need the following claims to analyze the running time and the cost of the covering.

CLAIM 7.4 (SUBPROBLEMS OF TYPE (x, y) FORM A STAIRCASE). *For any distinct subproblems $(I, J), (I', J')$ of the same type (x, y) , we have $\max(I') < \min(I)$ and $\max(J) < \min(J')$ (or vice versa). Moreover, if this holds then $A_{I'} + B_J \subseteq [u]$.*

PROOF. This is true in the beginning since $|\mathcal{D}| = 1$. Moving (I, J) from $\mathcal{D}$ to C cannot violate this property. Therefore, the only crucial step is the splitting of I into I_1, I_2 and of J into J_1, J_2 , where we remove (I, J) from $\mathcal{D}$ and add $(I_1, J_2), (I_2, J_1)$ to $\mathcal{D}$. In this situation, clearly $(I_1, J_2), (I_2, J_1)$ form a staircase, and $A_{I_1} + B_{J_1} \subseteq [u]$. Thus, the property is maintained “locally”. It is not hard to see that the property is also maintained “globally”, when we have distinct subproblems $(I, J), (I', J')$ that satisfy the property and we split both of them into subproblems. $\square$

CLAIM 7.5 (ANY SUBPROBLEM CREATES AT MOST TWO SUBPROBLEMS OF STRICTLY SMALLER TYPE). *Processing a subproblem $(I, J) \in \mathcal{D}$ can cause the insertion of at most two new subproblems into $\mathcal{D}$, both of strictly smaller type.*

PROOF. The only point at which we add new subproblems to $\mathcal{D}$ is the splitting phase. So let (I, J) be a subproblem of type (x, y) and consider the splitting of I into I_1, I_2 and of J into J_1, J_2 , where we remove (I, J) from $\mathcal{D}$ and add $(I_1, J_2), (I_2, J_1)$ to $\mathcal{D}$. Since we split I at the midpoint, we have $|I_r| \leq \lceil |I|/2 \rceil$, for any $r \in \{1, 2\}$. We clearly also have $|J_r| \leq \lfloor |J| \rfloor$. Hence, the new type $(x_r, y_r) = (\lceil \log |I_r| \rceil, \lfloor \log |J_r| \rfloor)$ satisfies³ $x_r < x$ and $y_r \leq y$. As this implies $x_r + y_r < x + y$, the newly added subproblems have a strictly smaller type. $\square$

CLAIM 7.6 (INVARIANT ON THE NUMBER OF SUBPROBLEMS). *At any point during the algorithm, there are at most $2q \log(n) \log(m)$ subproblems of any fixed type (x, y) .*

PROOF. The claim is immediate for $(x, y) = (\lceil \log n \rceil, \lfloor \log m \rfloor)$. Fix a type (x, y) and assume that the claim is true for every type (x', y') with $(x, y) < (x', y')$. Note that for any type (x', y') at most q subproblems of type (x', y') reached the splitting phase, and each such subproblem gave rise to at most two newly added subproblems. Since the number of different types is at most $\log(n) \log(m)$, we obtain the claimed bound on the number of subproblems of type (x, y) . $\square$

The following claim lies at the core of the analysis of our covering construction, as it bounds the running time and added cost of Case 2.

CLAIM 7.7. *Fix a type (x, y) with $2^{x+y} > \widetilde{\text{out}}$. Lines 11-13 of Algorithm 3 take time $\widetilde{O}(\text{out}^{4/3})$ and add rectangles of total cost $\widetilde{O}(\text{out}^{4/3})$ to C .*

PROOF. We denote by $(I_1, J_1), (I_2, J_2), \dots, (I_R, J_R)$ the subproblems of type (x, y) . We can assume $R \geq q$, since otherwise lines 11-13 are not called. Note that $R \leq 2q \log(n) \log(m)$ by Claim 7.6.

³To be precise, here we use that for any integers z, w the inequalities $2^{w-1} < z \leq 2^w$ imply $2^{w-2} < \lceil z/2 \rceil \leq 2^{w-1}$, and thus $\lceil \log(\lceil z/2 \rceil) \rceil \leq \lceil \log z \rceil - 1$.

By Claim 7.4 these subproblems form a staircase, so we can assume that

$$\max(I_1) < \min(I_2), \max(I_2) < \min(I_3), \dots, \max(I_{R-1}) < \min(I_R) \\ \min(J_1) > \max(J_2), \min(J_2) > \max(J_3), \dots, \min(J_{R-1}) > \max(J_R).$$

Claim 7.4 also implies that for any $r < \ell$ we have

$$A_{I_r} + B_{J_\ell} \subseteq [u]. \quad (5)$$

Set $\delta := 1/(6 \log(n) \log(m))$, so that $3\delta R \leq q$. For any $r \in [R]$ set

$$S(r) := \sum_{i=1}^{r-1} \sum_{j=r+1}^R |A_{I_r} + B_{J_j}| + |A_{I_i} + B_{J_j}| + |A_{I_i} + B_{J_r}|.$$

CLAIM 7.8. *There are at least $(1 - \delta)R$ indices r with $S(r) \leq 6R \text{out}/\delta$.*

PROOF. By simple counting how often a summand $|A_{I_r} + B_{J_\ell}|$ can appear, we observe that

$$\sum_{r=1}^R S(r) \leq 3R \sum_{r < \ell} |A_{I_r} + B_{J_\ell}|.$$

Note that we have $|A_{I_r} + B_{J_\ell}| \leq \text{out}$ by equation (5). However, we need a stronger property. We decompose the sum $\sum_{r < \ell} |A_{I_r} + B_{J_\ell}|$ into $2R - 1$ diagonal sums of the form $\sum_r |A_{I_r} + B_{J_{r+\Delta}}|$. Since $\max(A_{I_r} + B_{J_\ell}) < \min(A_{I_{r+1}} + B_{J_{\ell+1}})$, all summands in a diagonal sum are disjoint, and thus each diagonal sum is bounded from above by out. We thus obtain

$$\sum_{r < \ell} |A_{I_r} + B_{J_\ell}| \leq (2R - 1) \text{out},$$

which yields

$$\sum_{r=1}^R S(r) \leq 6R^2 \text{out}.$$

By Markov's inequality, it follows that all but δR indices r satisfy $S(r) \leq 6R \text{out}/\delta$. $\square$

In the remainder we consider only indices $\delta R \leq r \leq (1 - \delta)R$ with $S(r) \leq 6R \text{out}/\delta$; note that there are at least $(1 - 3\delta)R$ such indices r . For any such r , there are at least $\frac{\delta R^2}{2}$ pairs (i, j) with $1 \leq i < r$ and $r < j \leq R$. Hence, there exist i, j with $i < r < j$ and

$$|A_{I_r} + B_{J_j}| + |A_{I_i} + B_{J_j}| + |A_{I_i} + B_{J_r}| \leq \left(\frac{6R \text{out}}{\delta} \right) / \left(\frac{\delta R^2}{2} \right) = \frac{12 \text{out}}{\delta^2 R}.$$

Invoking the corollary of Ruzsa's triangle inequality (Lemma 7.2) now yields

$$\begin{aligned} |A_{I_r} + B_{J_r}| &\leq \\ \frac{|A_{I_r} + B_{J_j}| \cdot |B_{J_j} + A_{I_i}| \cdot |A_{I_i} + B_{J_r}|}{|A_{I_i}| \cdot |B_{J_j}|} &\leq \\ \frac{12^3 \text{out}^3 / (\delta^6 R^3)}{2^{x-1} \cdot 2^{y-1}}. \end{aligned}$$

By the case assumption $2^{x+y} > \widetilde{\text{out}} \geq \text{out}$ and $R \geq q = \lceil \text{out}^{1/3} \rceil \geq \text{out}^{1/3}$, and by our choice of $\delta = 1/(6 \log(n) \log(m))$, we obtain

$$|A_{I_r} + B_{J_r}| \leq O(\text{out} \cdot \log^{12}(nm)).$$

Recall that this holds for at least $(1 - 3\delta)R \geq R - q$ many indices r .

Since we compute all sumsets $A_{I_r} + B_{J_r}$ in parallel and stop once all but q calls are finished, it follows that we stop after spending

Figure 1: In the image, the thick gray squares depict the unprocessed subproblems $(I_1, J_1), (I_2, J_2), \dots$ of a specific type (x, y) . The light gray squares depict pairs (I_r, J_ℓ) for $r < \ell$. The vertices of the blue rectangle correspond to (starting from the top-right corner and going clockwise) $(I_r, J_r), (I_r, J_j), (I_i, J_j), (I_i, J_r)$. From such a 4-tuple, where all sumsets $|A_{I_r} + B_{J_j}|, |A_{I_i} + B_{J_j}|, |A_{I_r} + B_{J_r}|$ are “small”, Ruzsa’s triangle inequality (Lemma 7.2) allows us to bound $|A_{I_r} + B_{J_r}|$ from above.

time $\widetilde{O}(\text{out})$ for each call. Over $R \leq 2q \log(n) \log(m) = \widetilde{O}(\text{out}^{1/3})$ calls, this takes total time $\widetilde{O}(\text{out}^{4/3})$. Moreover, any finished call ran in time $\widetilde{O}(\text{out})$, so it has output-size $\widetilde{O}(\text{out})$, so it contributes cost $\widetilde{O}(\text{out})$. Thus, we add rectangles of total cost $\widetilde{O}(\text{out}^{4/3})$ in every invocation of lines 11-13 of Algorithm 3. $\square$

CLAIM 7.9 (COST BOUND). *For any type (x, y) we add rectangles of total cost $\widetilde{O}(\text{out}^{4/3})$ to C .*

PROOF. If $2^{x+y} > \widetilde{\text{out}}$, then the cost contributed by lines 11-13 of Algorithm 3 is bounded in Claim 7.7. Additionally, in the splitting phase (line 18) we add a rectangle (I_1, J_1) satisfying $A_{I_1} + B_{J_1} \subseteq [u]$, and thus $|A_{I_1} + B_{J_1}| \leq \text{out}$. Since we split at most q subproblems, we add a total cost of $\widetilde{O}(q \cdot \text{out}) = \widetilde{O}(\text{out}^{4/3})$.

If $2^{x+y} \leq \widetilde{\text{out}}$, then by the trivial bound each added rectangle (I, J) has cost at most $|I| \cdot |J| \leq 2^{x+y} \leq \widetilde{\text{out}} = O(\text{out})$. By Claim 7.6 we have $\widetilde{O}(q)$ subproblems, and thus we add a total cost of $\widetilde{O}(q \cdot \text{out}) = \widetilde{O}(\text{out}^{4/3})$. $\square$

Over all types (x, y) , we obtain a total cost of $\widetilde{O}(\text{out}^{4/3})$. The running time bound of $\widetilde{O}(\text{out}^{4/3})$ follows along the same lines. This finishes the proof of Theorem 2.9.

8 CONCLUSION AND FUTURE WORK

We initiated a line of research which strives for a SUBSETSUM algorithm that computes the set $\mathcal{S}(X, t)$, consisting of all subsets sums

of X below t , in near-linear output-sensitive time $\tilde{O}(|S(X, t)|)$. Our approach lead us to studying a new type of convolution problem: In Top- k -Convolution the task is to compute the lowest k monomials in the product of two sparse polynomials. Many open problems are spawned by our work, most notably on understanding the notions of *covering* and designing algorithms that exploit the additive structure in a different way.

ACKNOWLEDGEMENTS

We are grateful to Shachar Lovett for the resolution of an Additive Combinatorics question in an early stage of this work, which gave the core idea for Theorem 2.12, and for allowing us to include his construction in this paper.

REFERENCES

- [1] Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. 2019. SETH-based lower bounds for subset sum and bicriteria path. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 41–57.
- [2] Karl R. Abrahamson. 1987. Generalized String Matching. *SIAM J. Comput.* 16, 6 (1987), 1039–1051.
- [3] Josh Alman. 2019. Limits on the Universal Method for Matrix Multiplication. In *34th Computational Complexity Conference, CCC 2019, July 18–20, 2019, New Brunswick, NJ, USA*. 12:1–12:24.
- [4] Josh Alman and Virginia Vassilevska Williams. 2018. Further Limitations of the Known Approaches for Matrix Multiplication. In *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11–14, 2018, Cambridge, MA, USA*. 25:1–25:15.
- [5] Josh Alman and Virginia Vassilevska Williams. 2018. Limits on All Known (and Some Unknown) Approaches to Matrix Multiplication. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7–9, 2018*. 580–591.
- [6] Andrew Arnold and Daniel S Roche. 2015. Output-sensitive algorithms for subset sum and sparse polynomial multiplication. In *Proceedings of the 2015 ACM on International Symposium on Symbolic and Algebraic Computation*. ACM, 29–36.
- [7] Per Austrin, Petteri Kaski, Mikko Koivisto, and Jussi Määttä. 2013. Space-Time Tradeoffs for Subset Sum: An Improved Worst Case Algorithm. In *Proc. of the 40th International Colloquium on Automata, Languages, and Programming (ICALP)*. 45–56.
- [8] Per Austrin, Petteri Kaski, Mikko Koivisto, and Jesper Nederlof. 2015. Subset Sum in the Absence of Concentration. In *Proc. of the 32nd International Symposium on Theoretical Aspects of Computer Science (STACS)*. 48–61.
- [9] Per Austrin, Petteri Kaski, Mikko Koivisto, and Jesper Nederlof. 2016. Dense Subset Sum May Be the Hardest. In *Proc. of the 33rd Symposium on Theoretical Aspects of Computer Science (STACS)*. 13:1–13:14.
- [10] Kyriakos Axiotis, Arturs Backurs, Ce Jin, Christos Tzamos, and Hongxun Wu. 2019. Fast Modular Subset Sum using Linear Sketching. In *Proc. of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 58–69.
- [11] Nikhil Bansal, Shashwat Garg, Jesper Nederlof, and Nikhil Vyas. 2017. Faster Space-Efficient Algorithms for Subset Sum, k -Sum and Related Problems. In *Proc. of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. 198–209.
- [12] Richard E. Bellman. 1957. *Dynamic programming*. Princeton University Press.
- [13] Ernest F. Brickell and Andrew M. Odlyzko. 1988. Cryptanalysis: A survey of recent results. *Proc. IEEE* 76, 5 (1988), 578–593.
- [14] Karl Bringmann. 2017. A Near-Linear Pseudopolynomial Time Algorithm for Subset Sum. In *Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1073–1084.
- [15] Karl Bringmann and Vasileios Nakos. 2019. Fast n -fold Boolean Convolution via Additive Combinatorics. *Under submission* (2019).
- [16] Boris Bukh. 2018. Walk through Combinatorics: Subset inequalities. http://www.borishbukh.org/DiscreteMath14/notes_additive_inequalities.pdf. (2018).
- [17] Timothy M. Chan and Moshe Lewenstein. 2015. Clustered Integer 3SUM via Additive Combinatorics. In *Proc. of the 47th Annual ACM Symposium on Theory of Computing (STOC)*. 31–40.
- [18] Benny Chor and Ronald R. Rivest. 1988. A knapsack-type public key cryptosystem based on arithmetic in finite fields. *IEEE Transactions on Information Theory* 34, 5 (1988), 901–909.
- [19] Richard Cole and Ramesh Hariharan. 2002. Verifying candidate matches in sparse and wildcard matching. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 592–601.
- [20] Marek Cygan, Holger Dell, Daniel Lokshtanov, Daniel Marx, Jesper Nederlof, Yoshio Okamoto, Ramamohan Paturi, Saket Saurabh, and Magnus Wahlström. 2016. On problems as hard as CNF-SAT. *ACM Transactions on Algorithms* 12, 3 (2016), 41.
- [21] Itai Dinur, Orr Dunkelman, Nathan Keller, and Adi Shamir. 2012. Efficient Dissection of Composite Problems, with Applications to Cryptanalysis, Knapsacks, and Combinatorial Search Problems. In *Proc. of the 32nd Annual Conference on Advances in Cryptology (CRYPTO)*. 719–740.
- [22] François Le Gall. 2014. Powers of tensors and fast matrix multiplication. In *International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23–25, 2014*. 296–303.
- [23] Paweł Gawrychowski and Przemysław Uznanski. 2017. Optimal trade-offs for pattern matching with k mismatches. *arXiv preprint abs/1704.01311* (2017).
- [24] Omer Gold and Micha Sharir. 2017. Improved Bounds for 3SUM, k -SUM, and Linear Degeneracy. In *Proc. of the 25th Annual European Symposium on Algorithms (ESA)*. 42:1–42:13.
- [25] Ellis Horowitz and Sartaj Sahni. 1974. Computing partitions with applications to the knapsack problem. *J. ACM* 21, 2 (1974), 277–292.
- [26] Russell Impagliazzo and Moni Naor. 1996. Efficient cryptographic schemes provably as secure as subset sum. *Journal of Cryptology* 9, 4 (1996), 199–216.
- [27] Ce Jin and Hongxun Wu. 2019. A Simple Near-Linear Pseudopolynomial Time Randomized Algorithm for Subset Sum. In *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019 (OASICS)*, Vol. 69. 17:1–17:6.
- [28] Richard M. Karp. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*. Springer, 85–103.
- [29] Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. *Knapsack problems*. Springer.
- [30] Konstantinos Koiliaris and Chao Xu. 2017. A Faster Pseudopolynomial Time Algorithm for Subset Sum. In *Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1062–1072.
- [31] Andrea Lincoln, Virginia Vassilevska Williams, Joshua R. Wang, and R. Ryan Williams. 2016. Deterministic Time-Space Trade-Offs for k -SUM. In *Proc. of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*. 58:1–58:14.
- [32] Daniel Lokshantov, Daniel Marx, and Saket Saurabh. 2011. Known Algorithms on Graphs on Bounded Treewidth are Probably Optimal. In *Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 777–789.
- [33] Silvano Martello and Paolo Toth. 1990. *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc.
- [34] Ralph Merkle and Martin Hellman. 1978. Hiding information and signatures in trapdoor knapsacks. *IEEE Transactions on Information Theory* 24, 5 (1978), 525–530.
- [35] Michael Monagan and Roman Pearce. 2009. Parallel sparse polynomial multiplication using heaps. In *Proceedings of the 2009 international symposium on Symbolic and algebraic computation*. ACM, 263–270.
- [36] Vasileios Nakos. 2019. Nearly Optimal Sparse Polynomial Multiplication. *arXiv preprint arXiv:1901.09355* (2019).
- [37] Jesper Nederlof. 2017. A short note on Merlin-Arthur protocols for subset sum. *Inform. Process. Lett.* 118 (2017), 15–16.
- [38] Kevin O'Bryant. 2011. Sets of integers that do not contain long arithmetic progressions. *Electronic Journal of Combinatorics* 18, 1 (2011), P59.
- [39] Andrew M. Odlyzko. 1990. The rise and fall of knapsack cryptosystems. *Cryptology and Computational Number Theory* 42 (1990), 75–88.
- [40] Daniel S. Roche. 2008. Adaptive polynomial multiplication. *Proc. Milestones in Computer Algebra (MICA@AZ08)* (2008), 65–72.
- [41] Daniel S. Roche. 2018. What Can (and Can't) We Do with Sparse Polynomials? *arXiv preprint arXiv:1807.08289* (2018).
- [42] Adi Shamir. 1984. A polynomial-time algorithm for breaking the basic Merkle-Hellman cryptosystem. *IEEE Transactions on Information Theory* 30, 5 (1984), 699–704.
- [43] Xuancheng Shao and Wenqiang Xu. 2019. A robust version of Freiman's 3k–4 Theorem and applications. In *Mathematical Proceedings of the Cambridge Philosophical Society*, Vol. 166. Cambridge University Press, 567–581.
- [44] Terence Tao and Van H. Vu. 2006. *Additive combinatorics*. Vol. 105. Cambridge University Press.
- [45] Haim Kaplan, László Kozma, Or Zamir, and Uri Zwick. Selection from heaps, row-sorted matrices, and $X+Y$ using soft heaps. In Jeremy T. Fineman and Michael Mitzenmacher, editors, *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019, January 8–9, 2019 - San Diego, CA, USA*, volume 69 of OASICS, pages 5:1–5:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [46] Joris Van Der Hoeven and Grégoire Lecerc. 2012. On the complexity of multivariate blockwise polynomial multiplication. In *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation*. ACM, 211–218.
- [47] Virginia Vassilevska Williams. 2012. Multiplying matrices faster than Coppersmith-Winograd. In *Proc. of the 44th Symposium on Theory of Computing (STOC)*. 887–898.