

Space-Efficient Approximations for Subset Sum

ANNA GÁL, University of Texas at Austin

JING-TANG JANG, Google, Mountain View

NUTAN LIMAYE, IIT Bombay

MEENA MAHAJAN and KARTEEK SREENIVASIAH, Institute of Mathematical Sciences

SUBSETSUM is a well-known NP-complete problem: given $t \in \mathbb{Z}^+$ and a set S of m positive integers, output YES if and only if there is a subset $S' \subseteq S$ such that the sum of all numbers in S' equals t . The problem and its search and optimization versions are known to be solvable in pseudopolynomial time in general.

We develop a one-pass deterministic streaming algorithm that uses space $O(\frac{\log t}{\epsilon})$ and decides if some subset of the input stream adds up to a value in the range $\{(1 \pm \epsilon)t\}$. Using this algorithm, we design space-efficient fully polynomial-time approximation schemes (FPTAS) solving the search and optimization versions of SUBSETSUM. Our algorithms run in $O(\frac{1}{\epsilon}m^2)$ time and $O(\frac{1}{\epsilon})$ space on unit-cost RAMs, where $1 + \epsilon$ is the approximation factor. This implies constant space quadratic time FPTAS on unit-cost RAMs when ϵ is a constant. Previous FPTAS used space linear in m .

In addition, we show that on certain inputs, when a solution is located within a short prefix of the input sequence, our algorithms may run in sublinear time. We apply our techniques to the problem of finding balanced separators, and we extend our results to some other variants of the more general knapsack problem.

When the input numbers are encoded in unary, the decision version has been known to be in log space. We give streaming space lower and upper bounds for unary SUBSETSUM (USS). If the input length is N when the numbers are encoded in unary, we show that randomized s -pass streaming algorithms for exact SUBSETSUM need space $\Omega(\frac{\sqrt{N}}{s})$ and give a simple deterministic two-pass streaming algorithm using $O(\sqrt{N} \log N)$ space.

Finally, we formulate an encoding under which USS is monotone and show that the exact and approximate versions in this formulation have monotone $O(\log^2 t)$ depth Boolean circuits. We also show that any circuit using ϵ -approximator gates for SUBSETSUM under this encoding needs $\Omega(n/\log n)$ gates to compute the disjointness function.

CCS Concepts: • **Theory of computation** → **Streaming, sublinear and near linear time algorithms; Approximation algorithms analysis**

Additional Key Words and Phrases: One-pass streaming algorithm, fully polynomial-time approximations schemes FPTAS, knapsack problem, monotone circuits

ACM Reference Format:

Anna Gál, Jing-Tang Jang, Nutan Limaye, Meena Mahajan, and KartEEK SreenivasaiAH. 2016. Space-efficient approximations for Subset Sum. *ACM Trans. Comput. Theory* 8, 4, Article 16 (June 2016), 28 pages. DOI: <http://dx.doi.org/10.1145/2894843>

This work was supported in part by NSF grant CCF-1018060.

A preliminary version of part of this article appeared in COCOON '12 [Limaye et al. 2012].

Authors' addresses: A. Gál, Department of Computer Science, University of Texas at Austin, 2317 Speedway, Stop D9500, Austin, TX 78712, USA; email: panni@cs.utexas.edu; J.-T. Jang, Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA; email: keithjtjang@google.com; N. Limaye, Department of Computer Science and Engineering, Indian Institute of Technology, Bombay, Powai 400076, Mumbai India; email: nutan@cse.iitb.ac.in; M. Mahajan, The Institute of Mathematical Sciences, CIT Campus, Taramani, Chennai 600113, India; email: meena@imsc.res.in; K. SreenivasaiAH (current address), Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany; email: kartEEK@mpi-inf.mpg.de.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 1942-3454/2016/06-ART16 \$15.00

DOI: <http://dx.doi.org/10.1145/2894843>

1. INTRODUCTION

The SUBSETSUM problem was one of Karp's first 21 NP-complete problems [Karp 1972]. Its decision, search, and optimization versions are defined as follows.

Definition 1.1 (SUBSETSUM).

- Input*: A finite set $S = \{a_1, \dots, a_m\}$ of m positive integers and $t \in \mathbb{Z}^+$.
- Decision*: Decide if there exists a subset $S' \subseteq S$ such that the sum of all numbers in S' equals t .
- Search*: Output such a subset if one exists.
- Optimization*: Output a subset $S' \subseteq S$ such that the sum of all numbers in S' is as large as possible but not larger than t . (If all numbers exceed t , output the empty subset.)

(S and S' can even be multisets; the numbers are not required to be distinct. In that case, we ask for a subset $T \subseteq [m]$ such that $\sum_{i \in T} a_i$ equals t . We use the terms *set* and *subset* only to keep notation simple.)

Approximation algorithms for the optimization version of the SUBSETSUM problem have been extensively studied. Approximation versions can also be defined for the decision and search questions. The approximation problems are framed as follows.

Definition 1.2 (ϵ -approximate SUBSETSUM).

- Input*: A finite set S of positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$.
- Decision*: Decide if there exists a subset $S' \subseteq S$ such that

$$(1 - \epsilon)t \leq \sum_{a_i \in S'} a_i \leq (1 + \epsilon)t.$$

- Search*: Output such a subset if it exists.
- Optimization*: Let t^* be the largest possible sum of a subset of S without exceeding t . Output a subset $S' \subseteq S$ such that the sum of all numbers in S' is at least $(1 - \epsilon)t^*$ but not larger than t .

Space complexity. The focus of this work is the space complexity of SUBSETSUM. As is common in complexity theory (i.e., see Arora and Barak [2009]), to measure the amount of workspace used by an algorithm, we assume that the input is “read only” and the output is “write only,” and we only count toward the space bound the space used in addition to storing the input and the output. In other words, storing the input and output is for free, but the input/output space cannot be used by the computation for any other purpose. Space may be measured in the Turing machine model (bit complexity) or in the unit-cost RAM model. An even more stringent model is the streaming model, where not only is the input read only, but further, the input arrives as a stream, and the algorithm can make a limited number of left-to-right passes over the stream.

We denote by n the length of the input represented in binary—that is, $n = \sum_{i=1}^m \log a_i + \log t$. The Unary SUBSETSUM (USS) problem is the same problem but with the input numbers given in unary (i.e., $1^B 0 1^{a_1} 0 \dots 1^{a_m}$). In this case, the input length N is $B + \sum_{i=1}^m a_i$. (To emphasize the difference, we use B instead of t when considering USS.) Note that $m + B \leq N$, and if the numbers form a set (as opposed to multiset) and are all less than B , then $N \leq (m + 1)B$. However, in general, there is no tighter relation between N and m, B . Our bounds for binary encodings are stated in terms of n, m, t and for unary encodings in terms of N alone.

1.1. Previous Space-Efficient Algorithms

The classical solution for the decision, search, and optimization versions of the SUBSET-SUM problem is to use dynamic programming [Dantzig 1957; Bellman 1957], which runs in time $O(mt)$ and space $O(t)$ in the unit-cost RAM model. This is pseudopolynomial in m —that is, if the value of t is bounded by a polynomial in m , then it runs in polynomial time and space. See Section 4.2 and Garey and Johnson [1979] for details.

With unary encodings of the input, the SUBSETSUM problem USS is known to be in P. Elberfeld et al. [2010] showed a powerful metatheorem for obtaining logspace upper bounds and used it to conclude that USS is even in logspace. Kane [2010] gave a considerably simplified logspace algorithm. Improving upon this further, Elberfeld et al. [2012] showed that under appropriate encodings, USS has polynomial-size formulas (equivalently, polynomial-size log depth circuits) and hence is in the class NC^1 . They also showed that USS has polynomial-size constant-depth formulas using majority and NOT gates and hence is in the class TC^0 . On the other hand, it can be shown easily that majority reduces to computing USS, and hence USS is TC^0 -hard. Thus, USS is TC^0 -complete.

Approximation algorithms for the optimization version of the SUBSETSUM problem have also been extensively studied. An approximation algorithm is said to achieve approximation factor $1 + \epsilon$ if for every input $\max(\frac{t^*}{\hat{t}}, \frac{\hat{t}}{t^*}) \leq 1 + \epsilon$, where t^* is the optimal solution and \hat{t} is the solution provided by the algorithm. In other words, for small ϵ , the computed solution is not too far from the optimal. With this definition, in our setting, if for every input $(1 - \epsilon)t^* \leq \hat{t} \leq t^*$, the approximation algorithm achieves approximation factor $1 + \epsilon'$ with $\epsilon' = \frac{\epsilon}{1 - \epsilon}$. The solution computed may not be optimal, but it is required to be feasible; therefore, for the SUBSETSUM problem, $\hat{t} \leq t^*$. Note that t^* is not known in advance.

If the parameter $\epsilon > 0$ is not fixed but is also provided as input, then the approximation algorithm with approximation factor $1 + \epsilon$ is called an *approximation scheme*. An approximation scheme is a *polynomial-time approximation scheme* (PTAS) if for any $\epsilon > 0$, the running time is polynomial in n , where n is the input length. An approximation scheme is a *fully polynomial-time approximation scheme* (FPTAS) if the running time is polynomial in both $1/\epsilon$ and n . See Vazirani [2001] and Cormen et al. [2009] for detailed discussions.

Ibarra and Kim [1975] gave FPTAS for the SUBSETSUM optimization problem with running time $O(m/\epsilon^2)$ and space $O(m + 1/\epsilon^3)$. Gens and Levner [1979] gave an alternative FPTAS algorithm running in time $O(m/\epsilon + 1/\epsilon^3)$ and space $O(m + 1/\epsilon)$. The results of these two works use dynamic programming (in addition, see the survey by Gens and Levner [1980]). As far as we know, the current best FPTAS algorithm for the SUBSETSUM problem is given by Kellerer et al. [2003], which runs in time $O(\min(m/\epsilon, m + 1/\epsilon^2 \log \frac{1}{\epsilon}))$ and space $O(m + 1/\epsilon)$. These bounds are stated with respect to the unit-cost RAM model. (See Woeginger and Yu [1992] and Bazgan et al. [2002] for further variants and FPTAS for the SUBSETSUM problem.)

Any algorithm solving the decision version of SUBSETSUM can be used to also output an appropriate subset as follows. Fix an arbitrary element a , and run the algorithm for the decision version (with $t - a$ instead of t) to check if there is a subset with the required sum that contains a . Repeat this process until an element is found that can participate in the subset, then repeat again on a smaller collection of elements. This procedure eventually finds a solution, and it runs the algorithm for the decision version no more than n^2 times. However, the input to the decision version needs to be changed for the various runs of the algorithm. If the updated input sets are kept in working space, the procedure described earlier will use linear space.

With a slightly more careful reduction, any space-efficient algorithm solving the decision version of SUBSETSUM can be used to obtain a space-efficient algorithm for solving the search and optimization versions of SUBSETSUM. Hence, the algorithms of Elberfeld et al. [2010] and Kane [2010] yield algorithms that solve the search and optimization versions of SUBSETSUM exactly, in time polynomial in n and t , and space $O(\log n + \log t)$ (on the Turing machine model). This significantly improves the space bound of the classical exact solutions from space $O(t)$ to $O(\log n + \log t)$. Note that this also gives an improvement over the algorithm devised by Lokshтанov and Nederlof [2010] that achieves a $\tilde{O}(m^2)$ space bound and has running time pseudopolynomial in n .

1.2. Our Results and Techniques

We start by considering the space complexity of approximating SUBSETSUM in Section 2. We first show that the decision version of approximate SUBSETSUM has a simple deterministic one-pass streaming algorithm. The algorithm takes input ϵ , t , \tilde{a} and, using space $O(\frac{\log t}{\epsilon})$, decides whether some subset of numbers from the stream add up to a sum in $[(1 - \epsilon)t, (1 + \epsilon)t]$ (Theorem 2.1). The ONEPASSAPPROXDECISION algorithm uses logarithmic space as long as all numbers involved are polynomial in m , the number of integers in the input. In the unit-cost RAM model, it uses $O(\frac{1}{\epsilon})$ registers. In the strict turnstile streaming model, we give a simpler deterministic one-pass streaming algorithm that uses $O(\log t + \log m)$ space for the special case when $\epsilon \geq 1/3$ (Theorem 2.6).

Next we design algorithms for the search and optimization versions of approximate SUBSETSUM (Theorems 2.7 and 2.12). We achieve polynomial running time (depending on the approximation factor but independent of the value of t) and obtain small space bounds at the same time. We give space-efficient FPTAS for the optimization version of SUBSETSUM. FPTAS were known for SUBSETSUM, but our algorithms use significantly less space at the expense of only slightly more time. Whereas the previous best FPTAS for SUBSETSUM run in time $O(\min(m/\epsilon, m + 1/\epsilon^2 \log \frac{1}{\epsilon}))$ and space $O(m + 1/\epsilon)$ on unit-cost RAM [Kellerer et al. 2003], our algorithms run in $O(\frac{1}{\epsilon} m^2)$ time and $O(\frac{1}{\epsilon})$ space on unit-cost RAM, where $1 + \epsilon$ is the approximation factor. This implies constant space FPTAS on unit-cost RAMs when ϵ is a constant, where the running time as well as the space used on unit-cost RAMs does not depend on the sizes of the numbers involved at all. On the Turing machine model, we cannot completely eliminate the dependence of the space bounds on the sizes of the numbers, as we have to charge space for writing down even a single number. But we do not use any extra space beyond this “trivial” dependence: on Turing machines, our space bounds are of the form $O(\frac{1}{\epsilon} \log t + \log m)$ —that is, the space used corresponds to storing just a constant number of integers for any constant $\epsilon > 0$. This should be contrasted with previous FPTAS, where the space used was linear in m (the number of integers in the input) even on unit-cost RAM.

An additional interesting feature of our search algorithm is that it outputs the colexicographically first acceptable subset with respect to the given input ordering, and if an acceptable subset exists within a short prefix of the input, it may run in sublinear time. It can be viewed as a streaming algorithm, with the number of passes depending on the size of the subset found.

A key idea in these algorithms is the observation that the ONEPASSAPPROXDECISION (from Theorem 2.1) can be modified to yield interesting space bounds even if the numbers involved are not limited to be polynomial in m . We achieve our results on SUBSETSUM by repeatedly running ONEPASSAPPROXDECISION, with an extra trick for the optimization version to avoid all dependence on t (the value of the target sum) in the running time, to ensure fully polynomial running time.

In Section 3, we highlight a connection between approximation algorithms for PARTITION and finding balanced graph separators. Based on this connection and the space-efficient approximation algorithms for SUBSETSUM, we give space-efficient algorithms to construct balanced graph separators (Theorem 3.3). For graphs with constant size balanced separators, our algorithm finds a minimum size separator in polynomial time and logarithmic space. Moreover, we find a separator of the smallest size without knowing the size of the smallest separator in advance. We define separators and give more details on how our bounds relate to previous results in Section 3.

In Section 4, we consider two variants of SUBSETSUM where similar algorithms can be designed. We extend our results on SUBSETSUM to KNAPSACK problems with equal weights (Theorems 4.1). This allows us to find the smallest possible acceptable subset for approximating the SUBSETSUM problem in a space-efficient way (Theorem 4.2). We also obtain space-efficient approximation algorithms for the unbounded KNAPSACK problem with equal weights, which includes the change-making problem as a special case (Theorems 4.3 and 4.4). Section 4 provides definitions.

We consider the space complexity of exactly solving USS in the streaming model in Section 5. In light of the logspace algorithms [Elberfeld et al. 2010; Kane 2010] for USS, a natural question to ask is, is it crucial to have access to the entire input at any time to be able to solve USS in logspace? In other words, how hard is USS with regard to space when the inputs are read in a stream? The input numbers arrive in a stream, and we want to design a small space algorithm that makes one pass, or maybe a few passes, over the input stream and decides the instance. We use lower bounds from communication complexity to show that any randomized s -pass streaming algorithm for USS that makes error bounded by $1/3$ must use $\Omega(\frac{\sqrt{N}}{s})$ space (Theorem 5.2). In addition, we give a simple deterministic two-pass streaming algorithm for USS that uses space $O(\sqrt{N} \log N)$.

In Section 6, we leave the uniform computation model and look at SUBSETSUM from the perspective of monotone circuit complexity. As mentioned earlier, USS is known to be complete for the extremely small circuit class TC^0 . Note that SUBSETSUM is naturally monotone in the following sense: if the number of occurrences of a number i in the stream is increased, a YES instance remains a YES instance. Can we therefore construct small-depth monotone circuits for it? To model this monotonicity, we consider the following encoding of USS: for each positive integer B , the input consists of the frequency of each number in the stream in unary. In other words, an instance consists of B blocks of B bits each, where the i -th block w_i has as many ones as the number of occurrences m_i of the number i in the stream. Thus, the input records the multiplicity of each number in $[B]$ (without loss of generality, no multiplicity exceeds B). Call this problem the multiplicity-USS problem, mUSS. We show, by a monotone reduction to reachability, that this problem has monotone circuits of polynomial size and $O(\log^2 B)$ depth (Theorem 6.2). The circuit that we construct can also be used to solve the approximate version of USS.

We know that if we are allowed USS as gate primitives (or as an oracle in an oracle circuit), then we can capture all of TC^0 within polynomial-size and constant depth, as USS is TC^0 complete. In a related vein, we ask the following: how powerful are ε -approximators when used as gate primitives in circuits? We observe that ε -approximators for mUSS (we call them ApproxUSS gates) are at least as powerful as threshold gates (Lemma 6.5). Using a technique introduced by Nisan [1994], we also show that any circuit computing the disjointness function using ε -mUSS gates requires $\Omega(n/\log n)$ gates (Lemma 6.6). However, we have not been able to compare ApproxUSS gates explicitly with linear threshold gates. We conclude in Section 7 with some open problems.

2. STREAMING AND SPACE-EFFICIENT APPROXIMATION ALGORITHMS FOR SUBSET SUM

2.1. Decision Algorithms

To begin with, we consider the streaming model. We wish to approximate SUBSETSUM using a small number of passes on the input and using space polylogarithmic in the length of the input. For any ϵ and t and input stream $\tilde{a} = a_1, \dots, a_n$ where each $a_i \in [t]$, we say that set $S \subseteq [n]$ is an ϵ -approximator of t in \tilde{a} if $(\sum_{i \in S} a_i) \in [t(1 - \epsilon), t(1 + \epsilon)]$. Given ϵ, t, \tilde{a} , we want to decide whether there is an ϵ -approximator of t in \tilde{a} . We prove the following theorem.

THEOREM 2.1. *There is a deterministic one-pass streaming algorithm, ONEPASSAPPROXDECISION, that on an input stream ϵ, t, \tilde{a} with $0 < \epsilon < 1$ outputs 1 if and only if there exists an ϵ -approximator for t in the stream \tilde{a} . The space used by the algorithm is*

- (1) $O(\frac{1}{\epsilon})$ in the unit-cost RAM model, and
- (2) $O(\frac{\log t}{\epsilon})$ in the Turing machine model.

PROOF. Consider the following algorithm. Before seeing why the algorithm is correct, we first consider the space analysis. Note that at the beginning of each iteration, T has a set of disjoint intervals and each interval has size at least $2t\epsilon$. There can be at most $t/(2t\epsilon)$ disjoint intervals from 1 to t , so at any given time, $|T| \leq \frac{1}{\epsilon}$. Since T' has two intervals for each interval of T , $|T'|$ is also $O(\frac{1}{\epsilon})$. Thus, the algorithm needs to store the endpoints of at most $O(\frac{1}{\epsilon})$ intervals. In the unit-cost RAM model, each endpoint requires one register, so the space used is $O(|T|) \in O(\frac{1}{\epsilon})$. In the Turing machine model, since $\epsilon < 1$, all endpoints are in the range $\{0, 1, \dots, 2t\}$, and so the space required to store any one endpoint is $O(\log t)$. Thus, the total space used is $O(\frac{\log t}{\epsilon})$.

ALGORITHM: ONEPASSAPPROXDECISION

Maintain a set of intervals T .

Initialize: $T \leftarrow \{[t(1 - \epsilon), t(1 + \epsilon)]\}$.

while End of stream not reached **do**

$a \leftarrow$ Next number in stream.

if \exists interval $[\alpha, \beta] \in T$ such that $a \in [\alpha, \beta]$ **then**

Output YES and stop.

else

$T' \leftarrow \{[\alpha, \beta], [\alpha - a, \beta - a] \mid [\alpha, \beta] \in T\}$;

$T \leftarrow T'$.

Merge overlapping intervals in T to get a set of pairwise disjoint intervals.

(If $[a, b], [c, d] \in T$ and $a \leq c \leq b \leq d$, remove $[a, b], [c, d]$ and add $[a, d]$.)

end if

end while

We now show that the ONEPASSAPPROXDECISION algorithm is correct—that is, it outputs YES if and only if there exists a subset of the input numbers that has sum in $[t(1 - \epsilon), t(1 + \epsilon)]$. The intuition behind the correctness is the following. We maintain the set of intervals T such that if any number in the union of the intervals in T is seen as input, then there indeed exists a subset that generates t . This is true in the beginning by the way we initialize T . When a number a is read, a copy of each interval in T is shifted down by a to create a new interval. Therefore, if a number in any of these new intervals is seen, then it can be combined with a to give a number in one of the older intervals. (The original intervals are also retained, so we can also choose to

not use a in creating a subset.) In addition, this property is maintained by updating T with every number seen. Note that no interval in T gets deleted. Intervals in T only get merged into other intervals to become larger intervals, and this does not affect the invariant property.

We now describe the proof more formally. For a set of intervals T , define

$$R(T) = \{a \mid \exists[\alpha, \beta] \in T : a \in [\alpha, \beta]\}.$$

In other words, $R(T)$ is the union of all intervals in T . Let $\ell = t(1 - \epsilon)$ and $r = t(1 + \epsilon)$. Initially, $R(T) = \{a \mid \ell \leq a \leq r\}$.

\Rightarrow : Assume that `ONEPASSAPPROXDECISION` outputs YES. Let T_k denote the collection of intervals after reading k numbers from the stream. `ONEPASSAPPROXDECISION` accepts at a stage k when it reads a number $a_k \in R(T_{k-1})$. We show by induction on k , that if $a \in R(T_k)$, then there is a subset of $\{a_1, \dots, a_k\} \cup \{a\}$ with sum in $[\ell, r]$. This establishes that the YES answers are correct.

In the beginning, T_0 is initialized to $\{[\ell, r]\}$. Thus, $a \in R(T_0) \Rightarrow a \in [\ell, r]$.

Now assume that the property holds after reading $k - 1$ numbers. In other words, if $a \in R(T_{k-1})$, then there is a subset of $\{a_1, \dots, a_{k-1}\} \cup \{a\}$ with sum in $[\ell, r]$.

If $a_k \in R(T_{k-1})$, the algorithm terminates here and there is nothing more to prove. Otherwise, $a_k \notin R(T_{k-1})$, and the algorithm goes on to construct T_k . The update sets $R(T_k)$ to contain all of $R(T_{k-1})$ and all numbers b such that $a_k + b \in R(T_{k-1})$. Now consider an $a \in R(T_k)$. If it also holds that $a \in R(T_{k-1})$, then we can pretend that a_k was not read at all and, using induction, pull out a subset of $\{a_1, \dots, a_{k-1}\} \cup \{a\}$ with sum in $[\ell, r]$. If $a \notin R(T_{k-1})$, then $a_k + a \in R(T_{k-1})$. By induction, we have a subset of $\{a_1, \dots, a_{k-1}\} \cup \{a_k + a\}$ with sum in $[\ell, r]$. Hence, either way we have a subset of $\{a_1, \dots, a_{k-1}, a_k\} \cup \{a\}$ with sum in $[\ell, r]$, as desired.

\Leftarrow : Let S , $|S| = k$, be the first subset of numbers in the input stream that has sum in $[\ell, r]$. In other words,

- $S = \{a_{i_1}, a_{i_2}, \dots, a_{i_k}\}$ for some $i_1 < i_2 < \dots < i_k$,
- $\sum_{j=1}^k a_{i_j} = t - t\lambda$ for some $|\lambda| < |\epsilon|$, and
- there is no such subset in a_1, \dots, a_{i_k-1} .

We will show that `ONEPASSAPPROXDECISION` outputs YES on reading a_{i_k} .

To simplify notation, let s_j denote a_{i_j} .

Observe that if a number a enters $R(T)$ at any stage, then it remains in $R(T)$ until the end of the algorithm. This is because an interval is deleted only when an interval containing it is added.

Now we observe the way that T gets updated. After reading s_1 , $R(T)$ will contain the intervals $\{[\ell, r], [\ell - s_1, r - s_1]\}$. (It may contain more numbers as well, but that is irrelevant.) After reading s_2 , $R(T)$ will contain $\{[\ell, r], [\ell - s_1, r - s_1], [\ell - s_2, r - s_2], [\ell - s_1 - s_2, r - s_1 - s_2]\}$. Proceeding in this way, and using the preceding observation that $R(T)$ never shrinks, after reading s_1, s_2, \dots, s_{k-1} , $R(T)$ will contain $[\ell - (s_1 + \dots + s_{k-1}), r - (s_1 + s_2 + \dots + s_{k-1})]$. But this interval is the following:

$$\begin{aligned} & [t(1 - \epsilon) - (s_1 + \dots + s_{k-1}), t(1 + \epsilon) - (s_1 + \dots + s_{k-1})] \\ &= [t(1 - \epsilon) - (t - t\lambda - s_k), t(1 + \epsilon) - (t - t\lambda - s_k)] \\ &= [s_k + t\lambda - t\epsilon, s_k + t\lambda + t\epsilon] \\ &= [s_k - t(\epsilon - \lambda), s_k + t(\epsilon + \lambda)]. \end{aligned}$$

Since $\epsilon > 0$ and $|\lambda| < |\epsilon|$, $s_k \in [s_k + t(\lambda - \epsilon), s_k + t(\lambda + \epsilon)]$. Hence, the algorithm will output YES when s_k is read. \square

We state here some observations concerning the `ONEPASSAPPROXDECISION` algorithm that will be useful in later sections. These observations can be proved by induction on r .

Definition 2.2. Let $S = \{a_1, \dots, a_m\}$ and $0 < \epsilon < 1$. A subset $S' \subseteq S = \{a_1, \dots, a_m\}$ is said to be a *good set* with respect to t and ϵ if

$$(1 - \epsilon)t \leq \sum_{a_i \in S'} a_i \leq (1 + \epsilon)t.$$

In other words, for some $T \subseteq [m]$ that is a ϵ -approximator of t in S , as defined in the beginning of this section, S' is the corresponding set of values $\{a_i \mid i \in T\}$.

Let $S = \{a_1, \dots, a_m\}$ of m positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$. Suppose that the `ONEPASSAPPROXDECISION` algorithm answers YES on input ϵ, t, S , and let a_r be the last element read from the stream before the algorithm stops.

OBSERVATION 2.3. *There exists a good set with respect to t and ϵ that contains a_r .*

OBSERVATION 2.4. *If $(1 - \epsilon)t - a_r > 0$, then there is a subset H' of $H = \{a_1, \dots, a_{r-1}\}$ such that the sum of the elements in H' satisfies*

$$(1 - \epsilon)t - a_r \leq \sum_{a_i \in H'} a_i \leq (1 + \epsilon)t - a_r.$$

OBSERVATION 2.5. *Let $P = \{a_1, \dots, a_p\} \subseteq S = \{a_1, \dots, a_m\}$ be the first p elements of S read when reading the set S as a stream. If P contains a good subset (with respect to a given t' and ϵ'), then the algorithm stops with YES (when running it on S with t' and ϵ'), and the last element read is a member of P .*

The strategy in proving Theorem 2.1 is to maintain a collection of intervals such that a single element in any of these intervals guarantees a YES answer. For $\epsilon \geq 1/3$, a simpler strategy is to maintain the sum of small enough elements. This yields a one-pass deterministic algorithm even in the strict turnstile model, which we describe next.

First, we briefly recall the definition of the turnstile model. Here the input is given as a stream $(\alpha_1, \alpha_2, \dots, \alpha_n)$, where each α_i is a pair $\langle x, c_i \rangle$, for $x \in \mathbb{Z}$ and $c_i \in \{1, -1\}$. If $c_i = 1$, then α_i is said to *insert* x ; otherwise, it is said to *delete* x . The model is called a *strict turnstile* if for every prefix of the input stream and for every element in the input, the number of copies inserted is at least as many as the number of copies deleted. Formally, let $I_x = \{i \mid \alpha_i = \langle x, c_i \rangle\}$. Then $\forall i \in I_x : \sum_{j \leq i, j \in I_x} c_j \geq 0$. The turnstile model arises naturally in applications, such as handling updates to a *large* database and tracking IP traffic on a network link. Moreover, if the entries/counts are guaranteed to stay nonnegative, the model naturally arising is in fact the strict turnstile model. For more details, see Muthukrishnan [2005] and Gilbert et al. [2001].

THEOREM 2.6. *Let $S = \{a_1, \dots, a_m\}$ be a finite set of m positive integers. There is a deterministic one-pass streaming algorithm in the strict turnstile model that, given $t \in \mathbb{Z}^+$, $\epsilon \geq 1/3$, and a stream $(\alpha_1, \alpha_2, \dots, \alpha_n)$, where each α_i is a pair $\langle a, c_i \rangle$, for $a \in S$ and $c_i \in \{1, -1\}$, decides if there is a subset $S' \subseteq S$ and integers $0 \leq b_a \leq \sum_{i \in I_a} c_i$ for $a \in S'$ such that*

$$(1 - \epsilon)t \leq \sum_{a \in S'} b_a a \leq (1 + \epsilon)t.$$

The algorithm runs in linear time and $O(\log t + \log m)$ space on Turing machines. It runs in $O(m)$ time and constant space on unit-cost RAMs.

ALGORITHM: TURNSTILEAPPROXDECISION

```

Count ← 0; Sum ← 0.
while End of stream not reached do
   $\langle a, c_i \rangle \leftarrow$  Next number in stream.
  if  $a \leq (1 + \epsilon)t$  then
    if  $a < (1 - \epsilon)t$  then
      Sum ← Sum +  $a \times c_i$ 
    else
      Count ← Count +  $c_i$ 
    end if
  end if
end while
if Count > 0 or Sum  $\geq (1 - \epsilon)t$  then
  return YES.
else
  return NO.
end if

```

PROOF. The algorithm is as follows.

To see that the algorithm is correct, first note that $Count > 0$ means that at least one element a such that $(1 - \epsilon)t \leq a \leq (1 + \epsilon)t$ survived—that is, $\sum_{i \in I_a} c_i > 0$ after reading the whole stream. Any such element forms an acceptable sum in itself.

Next suppose that $Sum \geq (1 - \epsilon)t$ in the end. If we also have $Sum \leq (1 + \epsilon)t$, then the sum of all surviving elements (except those that exceed $(1 + \epsilon)t$) is an acceptable sum. If $Sum > (1 + \epsilon)t$, then we must have a good surviving subset as well, as the elements contributing to Sum are each less than $(1 - \epsilon)t \leq 2\epsilon t = (1 + \epsilon)t - (1 - \epsilon)t$. Thus, some subcollection of them must form an acceptable sum.

Finally, we argue that if both $Count \leq 0$ and $Sum < (1 - \epsilon)t$, then no acceptable sum can be formed by the surviving elements. $Count \leq 0$ means that none of the surviving elements is in the correct range by itself. $Sum < (1 - \epsilon)t$ means that the sum of all surviving elements (except those that exceed $(1 + \epsilon)t$) is smaller than $(1 - \epsilon)t$, but then no subset of the surviving elements can be in the correct range. \square

2.2. Search Algorithms

Our approximation algorithms for the search and optimization versions of SUBSETSUM are based on repeatedly running the ONEPASSAPPROXDECISION algorithm. Unlike in the straightforward reduction to the decision version described in Section 1, we do not keep in working memory the elements found so far, and each time we run the algorithm on the same set S .

First, consider the search for an approximate solution.

THEOREM 2.7. *There is a deterministic algorithm APPROXSEARCH that, given a finite set $S = \{a_1, \dots, a_m\}$ of m positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$, outputs a subset $S' \subseteq S$ such that*

$$(1 - \epsilon)t \leq \sum_{a_i \in S'} a_i \leq (1 + \epsilon)t$$

if such a subset exists, and otherwise it outputs NO. The algorithm runs in $O(\frac{1}{\epsilon}mn) = O(\frac{1}{\epsilon}n^2)$ time, and $O(\frac{1}{\epsilon} \log t + \log m)$ space on Turing machines, where $n = \sum_{i=1}^m \log a_i + \log t$, and in $O(\frac{1}{\epsilon}m^2)$ time, and $O(\frac{1}{\epsilon})$ space on unit-cost RAMs.

Furthermore, if an appropriate subset exists among the first r elements of the set S , the algorithm reads only the r -element prefix of the input as a stream in at most r passes

and terminates within $O(\frac{1}{\epsilon}r^2)$ steps on the unit-cost RAM model, and $O(\frac{1}{\epsilon}rn_r)$ steps on Turing machines, where $n_r = \sum_{i=1}^r \log a_i + \log t$.

Before we prove the theorem, we make some remarks concerning the result itself.

Remark 2.8. We note that we could also obtain the same worst-case bounds using algorithm `ONEPASSAPPROXDECISION` as a blackbox, with a careful reduction. However, our current proof yields the additional property that on certain inputs the algorithm may run in sublinear time and may not have to read the whole input. This is because the algorithm outputs the co-lexicographically first subset that gives an appropriate sum with respect to the given input ordering. If there is such a subset among the first few elements of the input, the algorithm will run in sublinear time and does not have to read the whole input. In fact, one can argue that the algorithm reads only the shortest possible prefix of the input.

Remark 2.9. Note that the algorithm runs in polynomial time regardless of t , and it uses constant space on unit-cost RAMs for any constant $\epsilon > 0$.

Remark 2.10. We state our space bounds on Turing machines in terms of $\log m$. Of course, this could be replaced by $\log n$ as well since $m \leq n$. However, in general, $\log m$ may be significantly smaller than $\log n$ if the input includes very large values for some a_i -s. As noted earlier, we can simply ignore such values if they are larger than $(1 + \epsilon)t$, and we can detect if a given value is too large using $O(\log t)$ workspace.

PROOF OF THEOREM 2.7. The idea is to repeatedly run the algorithm `ONEPASSAPPROXDECISION` on the same set S , thus keeping the workspace used by the algorithm small. We only change the values of the target value t' and the approximation factor γ for the various runs. In the following algorithm, we treat `ONEPASSAPPROXDECISION` as a function call that returns either NO or a number r , where r is the last element read from the stream before deciding YES.

ALGORITHM: APPROXSEARCH

```

 $t' \leftarrow t; \gamma \leftarrow \epsilon; r \leftarrow \text{ONEPASSAPPROXDECISION}(S, t', \gamma).$ 
if  $r = \text{NO}$  then
  Output NO and stop.
else
  Write  $a_r$  on output tape.
  while  $r > 1$  do
    if  $(1 - \gamma)t' \leq a_r \leq (1 + \gamma)t'$  then
      Stop.
    else
       $\gamma \leftarrow \frac{\gamma t'}{t' - a_r}; t' \leftarrow t' - a_r; r \leftarrow \text{ONEPASSAPPROXDECISION}(S, t', \gamma).$ 
      Write  $a_r$  on output tape.
    end if
  end while
  Stop.
end if

```

First, we run algorithm `ONEPASSAPPROXDECISION` on S , $\gamma = \epsilon$, and $t' = t$. If the algorithm answers NO, we stop and output NO.

Suppose that the algorithm answers YES on this input, and let a_r be the last element read from the stream before the algorithm stops. By Observation 2.3, we know that a_r can be part of a correct output (a good set), so we write it on the output tape even though we may not have finished finding a good set yet.

If $r = 1$, or if $r > 1$ and also $(1 - \gamma)t' \leq a_r \leq (1 + \gamma)t'$, then the good set construction is complete and we stop.

Otherwise, note that $(1 - \gamma)t' - a_r > 0$ must hold. We now update the target t' to $t'' = t' - a_r$, the target sum for the rest of the good set. We also update the approximation factor to γ' so that the intervals $[(1 - \gamma)t' - a_r, (1 + \gamma)t' - a_r]$ and $[(1 - \gamma')t'', (1 + \gamma')t'']$ are the same. By Observation 2.4, we know that the `ONEPASSAPPROXDECISION` algorithm with these parameters must terminate with YES, and if $a_{r'}$ is the last element read, then by Observations 2.4 and 2.5, we know that $r' < r$.

We iterate the preceding process until $r = 1$ or the element a_r completes the good set construction. Since r strictly decreases in each iteration, the process must terminate in at most m iterations. When the procedure terminates, a good set is listed on the output tape.

Next we estimate the space used. In each iteration, we store at most $O(\frac{1}{\epsilon})$ intervals, and it takes constant space on unit-cost RAMs and $O(\log t)$ space on Turing machines to store one interval (by storing its two endpoints). We reuse space for subsequent iterations. The time bound follows since each iteration scans the input once, and processing each newly read integer may take up to $O(\frac{1}{\epsilon})$ steps to update the intervals. \square

In the special case of $\epsilon = \frac{1}{3}$, the same idea that was used in Theorem 2.6 can be used to get an algorithm that is even simpler than what would follow by substituting $\epsilon = \frac{1}{3}$ in Theorem 2.7. Not only is the algorithm simpler, it is a two-pass streaming algorithm.

THEOREM 2.11. *There is a deterministic two-pass streaming algorithm that, given a finite set $S = \{a_1, \dots, a_m\}$ of m positive integers and $t \in \mathbb{Z}^+$, outputs a subset $S' \subseteq S$ such that*

$$\frac{2}{3}t \leq \sum_{a_i \in S'} a_i \leq \frac{4}{3}t$$

if such subset exists, and otherwise it outputs NO. The algorithm runs in linear time $O(n)$ and $O(\log t + \log m)$ space, where $n = \sum_{i=1}^m \log a_i + \log t$, on Turing machines, and in $O(m)$ time and constant space on unit-cost RAMs.

PROOF. During the algorithm, we will only keep in working memory the value of a current sum s and some counters.

During the first pass, we start by initializing s to 0. For each $i = 1, \dots, n$, we look at the value a_i on the input tape. There are three cases:

- (1) If $a_i > \frac{4}{3}t$, we move to the next element.
- (2) If $\frac{2}{3}t \leq a_i \leq \frac{4}{3}t$, we output a_i and stop.
- (3) If $a_i < \frac{2}{3}t$, we let $s = s + a_i$. Then, if $s < \frac{2}{3}t$, we move to the next element. If we run out of elements and still have $s < \frac{2}{3}t$, we stop and output NO. If $\frac{2}{3}t \leq s \leq \frac{4}{3}t$, we move to the second pass (see the following). Note that $s > \frac{4}{3}t$ is not possible, as the previous sum was less than $\frac{2}{3}t$ and we added a value less than $\frac{2}{3}t$.

The only reason we need a second pass is because during the first pass we did not know yet if we have to answer NO or output a set, and the output tape is write only.

In the second pass, we proceed exactly as in the first pass; however, in addition, every time we add an element a_i to the sum, we also write it on the output tape. The algorithm stops when we reach for the first time a value $s \geq \frac{2}{3}t$. At this point, we also have the members of a good subset listed on the output tape.

We argue that the algorithm finds a good subset if one exists as follows. Elements larger than $\frac{4}{3}t$ cannot participate in a good subset, and if a single element a_i is in

the correct range, we find it during the first scan. If no such element exists, then all members of a good subset must be less than $\frac{2}{3}t$. If the first scan terminates with s still less than $\frac{2}{3}t$, it means that the sum of all elements less than $\frac{2}{3}t$ is a value also less than $\frac{2}{3}t$. But then the sum of any subset of these elements is also less than $\frac{2}{3}t$, and no good subset exists. \square

We can extend this argument to other values of ϵ and keep the running time linear in n for small values of $\frac{1}{\epsilon}$. In particular, we still get a two-pass algorithm for $\epsilon = \frac{1}{5}$ with $O(\log t + \log m)$ space. But we cannot keep the algorithm linear time and logarithmic space for arbitrary constant ϵ . As $\frac{1}{\epsilon}$ gets larger, our quadratic time algorithm of Theorem 2.7 gives better running time, while still using logarithmic space on Turing machines and constant space on unit-cost RAM, for arbitrary constant ϵ .

2.3. FPTAS for the Optimization Version of SUBSETSUM

We now describe an FPTAS for the optimization question.

THEOREM 2.12. *There is a deterministic approximation scheme for SUBSETSUM optimization that for any $\epsilon > 0$ achieves approximation factor $1 + \epsilon$ and runs in time $O(\frac{1}{\epsilon}n^2)$ and space $O(\frac{1}{\epsilon} \log t + \log m)$ on Turing machines, where $n = \sum_{i=1}^m \log a_i + \log t$, and in $O(\frac{1}{\epsilon}m(m + \log m + \log \frac{1}{\epsilon}))$ time and $O(\frac{1}{\epsilon})$ space on unit-cost RAMs.*

PROOF. At the outset, we compute the sum of all elements a_i that do not exceed t . If all elements exceed t (or if the set S is empty), we stop and output NO. If the sum of the elements not exceeding t is nonzero and it is smaller than t , then this is the optimal sum and there is nothing else to do. Thus, in what follows, we may assume that the sum of all elements a_i that do not exceed t is nonzero, and t is bounded above by this sum. On the other hand, the optimal sum t^* is at least as large as the largest a_i not exceeding t . This implies $t^* \geq t/m$. We will use this fact later when we analyze the running time.

Given an instance of SUBSETSUM optimization with S , $t \in \mathbb{Z}^+$, and $\epsilon > 0$, we run the following algorithm after the preprocessing described earlier.

ALGORITHM: FPTASOPTIMIZATION

```

Initialize:  $i \leftarrow 0$ ;  $c_0 \leftarrow \frac{t}{2}$ ;  $\xi_0 \leftarrow \frac{1}{2}$ ;  $\gamma_0 \leftarrow 1$ .
/* Maintain the invariants  $\gamma_i c_i = \xi_i t$ ;  $c_i - \xi_i t \leq t^* \leq c_i + \xi_i t \leq t$ . */
while  $\gamma_i > \frac{\epsilon}{2+\epsilon}$  do
   $i \leftarrow i + 1$ ;
   $\xi_i \leftarrow \frac{1}{2}\xi_{i-1}$ ;
   $t_i \leftarrow c_{i-1} + \xi_i t$ ;
   $r_i \leftarrow \frac{\xi_i t}{t_i}$ ;
  if ONEPASSAPPROXDECISION( $S, t_i, r_i$ ) = YES then
     $c_i \leftarrow t_i$ ;  $\gamma_i \leftarrow r_i$ ;
  else
     $c_i \leftarrow c_{i-1} - \xi_i t$ ;  $\gamma_i \leftarrow \frac{\xi_i t}{c_i}$ ;
  end if
end while
 $\ell \leftarrow i$ ;  $t' \leftarrow c_\ell$ ;  $\gamma \leftarrow \gamma_\ell$ .
Run APPROXSEARCH
return  $T$ .

```

The while loop implements a binary search procedure for an approximation to the optimal sum t^* by repeatedly calling the ONEPASSAPPROXDECISION algorithm with

varying parameters. In the while loop at stage i , initially the optimal is known to lie in the interval of width $2\xi_{i-1}t$ centered around c_{i-1} , and we want to check whether it is in the upper half of this interval. Since the width halves at each stage, this is precisely the interval $[c_{i-1} \pm \xi_i t]$, and the approximation parameter r_i is chosen so that it is also the same as the interval $[t_i(1 \pm r_i)]$. With these parameters, we call `ONEPASSAPPROX-DECISION`. If the answer is YES, we can set c_i, γ_i to t_i, r_i , respectively. Otherwise, t^* lies in the lower half of the previous interval, namely $[c_{i-1} - \xi_i t \pm \xi_i t]$, so c_i and γ_i are chosen accordingly. Thus, at the beginning of each stage, we maintain the invariants that $c_i \gamma_i = \xi_i t$ and $c_i - \xi_i t \leq t^* \leq c_i + \xi_i t \leq t$. Thus, $c_i = \hat{t} - \gamma_i t_i$, where \hat{t} is the upper end of the current target-search interval. In other words, sums larger than \hat{t} have already been eliminated from consideration. Throughout, we have $\hat{t} \leq t$.

Let ℓ denote the value of the index i when we exit the while loop - that is, $\gamma_\ell \leq \frac{\epsilon}{2+\epsilon}$ —for the first time. We will show later that $O(\log m + \log \frac{1}{\epsilon})$ iterations will suffice. Therefore, our running time will not depend on t . (Note that with a more straightforward reasoning, we would get up to $\log t$ iterations, which we want to avoid.) Setting $t' = c_\ell$ and $\gamma = \gamma_\ell$, the invariants tell us that

$$(1 - \gamma)t' \leq t^* \leq (1 + \gamma)t' \leq t,$$

where t^* denotes the value of the optimal sum.

At this point, we know that there exists a good subset S' of S with respect to t' and γ —that is,

$$(1 - \gamma)t' \leq \sum_{a_i \in S'} a_i \leq (1 + \gamma)t'.$$

Next we run our `APPROXSEARCH` algorithm from Theorem 2.7 that outputs such a subset. Note that we have chosen the parameters so that $(1 - \delta)t^* \leq (1 - \gamma)t'$. By definition, we cannot find a sum larger than the optimal sum. Thus, we obtain an approximation scheme with approximation factor $\frac{1}{1-\delta} = 1 + \epsilon$.

It remains to estimate ℓ . Recall that by our initial remarks, $t^* \geq t/m$. Therefore, $c_\ell = t' \geq \frac{1}{1+\gamma} \frac{t}{m}$, and $\frac{t}{c_\ell} \leq (1 + \gamma)m \leq 2m$. Recall that $\gamma_\ell = \xi_\ell \frac{t}{c_\ell}$, and we stop as soon as $\gamma_\ell \leq \frac{\epsilon}{2+\epsilon}$ —that is, we stop as soon as $\xi_\ell \leq \frac{\epsilon}{2+\epsilon} \frac{c_\ell}{t}$. Hence, we certainly stop by the stage when $\xi_\ell \leq \frac{\epsilon}{2+\epsilon} \frac{1}{2m}$. Thus, we stop after at most $\ell = O(\log m + \log \frac{1}{\epsilon})$ iterations. \square

3. AN APPLICATION: FINDING BALANCED SEPARATORS

Balanced separators are used in VLSI circuit design [Ullman 1984], algorithms on planar graphs [Lipton and Tarjan 1979, 1980; Miller 1986], graphs with excluded minors [Alon et al. 1990], and in general in algorithmic applications on graphs that involve a divide and conquer strategy.

Informally, a *balanced separator* of a graph G is a set of nodes whose removal yields two disjoint subgraphs of G that are comparable in size. More formally, balanced separators are defined as follows.

Definition 3.1 [Lipton and Tarjan 1979]. Let $G = (V, E)$ be an undirected graph. A set of nodes $S \subseteq V$ is a *balanced separator* of G if the removal of S disconnects G into two subgraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ that satisfy $|V_i| \geq \frac{1}{3}|V|$ for $i = 1, 2$.

We say that S is a *balanced h -separator* of G if S is a balanced separator of G and S consists of h vertices. We say that G is *h -separable* if it has a balanced h -separator.

Note that the definition can be easily extended to directed graphs. Additionally, instead of requiring that the size of both components is between $\frac{1}{3}|V|$ and $\frac{2}{3}|V|$, we can consider components with size closer to $\frac{1}{2}|V|$.

Definition 3.2. Let $G = (V, E)$ be an undirected graph. A set of nodes $S \subseteq V$ is a ϵ -balanced separator of G if the removal of S disconnects G into two subgraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ that satisfy $|V_i| \geq (1 - \epsilon) \frac{|V|}{2}$ for $i = 1, 2$.

The two subgraphs G_1 and G_2 could be disconnected within themselves. This implies that any graph $G = (V, E)$ has a ϵ -balanced separator of size $|V|$ for any constant $1 > \epsilon > 0$.

One special case of the SUBSETSUM problem is the PARTITION problem, which was also listed among Karp's first 21 NP-complete problems [Karp 1972]. The problem is defined as follows. Given a finite set S of positive integers, decide if there is a subset $S' \subseteq S$ such that $\sum_{a_i \in S'} a_i = \sum_{b_i \in S \setminus S'} b_i$. For the search version, output such a subset if it exists. In the SUBSETSUM problem, if we set $t = \frac{A}{2}$, where $A = \sum_{a_i \in S} a_i$, then we get the PARTITION problem. As for the SUBSETSUM problem, the classical solution for PARTITION is to use dynamic programming [Dantzig 1957; Bellman 1957], which runs in time $O(mt)$ and space $O(t)$.

We observe that space-efficient approximation algorithms for the decision version of the PARTITION problem yield space-efficient algorithms to find balanced separators.

THEOREM 3.3. *There is a deterministic algorithm that, given an undirected graph $G = (V, E)$ and $0 < \epsilon < 1$ as input, outputs a smallest-size ϵ -balanced separator. If the smallest ϵ -balanced separator has size h , the algorithm runs in time $O(\frac{1}{\epsilon} |V|^{(h+O(1))})$ and uses space $O((h + \frac{1}{\epsilon}) \log |V|)$.*

We offer some comments before we formally prove the theorem. Note that the value h is not provided as input to the algorithm, and the algorithm outputs a separator of the smallest possible size. The algorithm runs in polynomial time and logarithmic space for graphs with constant size balanced separators.

As far as we know, this is the first space-efficient algorithm that finds balanced separators in the sense of Definition 3.1 for general graphs. In the case of planar graphs, Imai et al. [2013] gave a polynomial time $O(n^{1/2+\epsilon})$ space algorithm to find balanced separators. Elberfeld et al. [2010] considered a different notion of separators (also commonly referred to as balanced separators in the literature), which we call *component-wise* balanced separators, to distinguish them from the definition of Lipton and Tarjan [1979] that we use.

Definition 3.4 [Elberfeld et al. 2010]. Let $G = (V, E)$ be an undirected graph. A set of nodes $S \subseteq V$ is a *component-wise* balanced separator of G if the removal of S disconnects G into connected subgraphs $G_i = (V_i, E_i)$ that satisfy $|V_i| \leq \frac{1}{2} |V|$ for each i .

Elberfeld et al. [2010] observed the following simple space-efficient algorithm for finding component-wise balanced separators. Enumerate all possible subsets S of V first by size and then by lexicographic order for each size, starting from subsets of size 1. For each subset S , use Reingold's algorithm [Reingold 2008], which solves undirected connectivity in logspace to compute the size of each component in the graph obtained by removing S from G . Output the first subset such that the size condition is satisfied by each component. It is clear that this algorithm can compute the component-wise balanced separator with minimum size and uses $O(h \log |V|)$ space if the smallest component-wise separator has size h .

In particular, this algorithm runs in logspace for graphs that have constant size component-wise separators. This was used in Elberfeld et al. [2010] to find the tree decomposition of graphs with constant tree width in logarithmic space.

However, for finding balanced separators in the sense of Definition 3.1, we also need to check whether the components obtained after removing a given subset of vertices can be grouped into two subsets of vertices such that each has size at least $(1 - \epsilon) \frac{|V|}{2}$. We observe that this corresponds to solving an instance of the ϵ -approximate PARTITION problem (a special case of ϵ -approximate SUBSETSUM), where the set $S = \{a_i\}$ consists of the sizes of the components $a_i = |V_i|$.

Moreover, since the values a_i are sizes of subsets of vertices of the graph, we are guaranteed that the value $A = \sum_i a_i$ is polynomial in $n = |V|$. Thus, our space-efficient algorithms for approximating the SUBSETSUM and PARTITION problems yield space-efficient algorithms for finding balanced separators.

Since for finding balanced separators we only need to check if a subset of components with the appropriate sum of sizes exists, the ONEPASSAPPROXDECISION algorithm, which solves the decision version of approximate SUBSETSUM, can be used.

Putting all of this together, we obtain the proof of Theorem 3.3.

PROOF OF THEOREM 3.3. We define three logspace bounded Turing machines to perform the following tasks. Note that each machine uses the original graph G as input together with some additional inputs, including the list of vertices in a fixed set W . Other than the space used for storing their respective inputs, each machine uses at most $O(\log |V|)$ workspace, where V is the set of vertices of V . Moreover, the length of the output of each machine is $O(\log |V|)$.

For an undirected graph $G = (V, E)$, a set $W \subseteq V$, and a node $v \in V \setminus W$, we denote by G_W the graph obtained from G by removing the vertices in W and all edges adjacent to vertices in W . We denote by $G_W(v)$ the connected component of G_W that contains v .

The first Turing machine M_1 on input G , W , and v computes the number of nodes connected to v in G_W —that is, M_1 outputs $|G_W(v)|$. This can be done by using Reingold's logspace algorithm for undirected connectivity [Reingold 2008] to check, one-by-one, for each node $u \in V \setminus W \setminus \{v\}$ if u is connected to v in G_W . We let Reingold's algorithm run on the original graph G (we do not store the representation of G_W), and we simply ignore edges of G adjacent to vertices in W when running the algorithm. We reuse space for each repetition.

M_2 on input G and W outputs the number of connected components of G_W . Note that there is always at least one component, so we start by initializing a counter to $count = 1$. We assume that the first vertex belongs to the first component. Next, for $i = 2, \dots, |V \setminus W|$, we check (using Reingold's algorithm) if $v_i \in V \setminus W$ is connected to any of the vertices $v_j \in V \setminus W$ for $1 \leq j < i$. If v_i is not connected to any of the previously considered vertices, we increase the counter and move on to considering the next vertex. If v_i is connected to some previously considered vertex, then we move on to the next vertex without increasing the value of $count$. When all vertices in $V \setminus W$ have been considered, the value of $count$ equals the number of connected components of G_W .

Let G_W^ℓ denote the ℓ -th connected component of G_W —that is, $G_W^1 = G_W(v_1)$ and $G_W^\ell = G_W(v_i)$, where i is the smallest integer such that v_i is not contained in any of the components $G_W^1, \dots, G_W^{\ell-1}$.

M_3 on input G , W , and ℓ computes the size of the ℓ -th connected component $|G_W^\ell|$ as follows. Initialize a counter to $count = 1$. Next, similarly to M_2 , for $i = 2, \dots, |V \setminus W|$, we check (using Reingold's algorithm) if $v_i \in V \setminus W$ is connected to any of the vertices $v_j \in V \setminus W$ for $1 \leq j < i$. If v_i is connected to some previously considered vertex, then we move on to the next vertex without increasing the value of $count$. If v_i is not connected to any of the previously considered vertices, we increase the counter by one, and then if $count < \ell$, we move on to considering the next vertex. If $count = \ell$, we let $v = v_i$ and use M_1 on input G , W , and v to compute $|G_W^\ell| = |G_W(v)|$.

We are now ready to describe the algorithm. We consider integers $1 \leq k \leq |V|$ in increasing order, and for each k we consider subsets $W \subseteq V$ of size k in lexicographic order. For fixed k and W , we use $k \log |V|$ workspace to keep the list of vertices in W . We use M_2 on G and W to compute the number of connected components of G_W . We denote this number by m . For $\ell = 1, \dots, m$, let a_ℓ denote $|G_W^\ell|$, the size of the ℓ -th connected component.

We now need to check whether the connected components can be grouped into two subsets of vertices such that each has size at least $(1 - \epsilon) \frac{|V|}{2}$. In other words, we need to solve an instance of ϵ -approximate SUBSETSUM on input $S = \{a_\ell\}$ and $t = \frac{|V|}{2}$. However, we do not have enough work space to explicitly keep the list of integers in S .

We run the ONEPASSAPPROXDECISION algorithm (the decision algorithm for solving ϵ -approximate SUBSETSUM) on $S = \{a_\ell\}$ and $t = \frac{|V|}{2}$ by using M_3 to compute the value a_ℓ when the algorithm needs it. Note that since the values a_i are sizes of subsets of vertices of the graph, we are guaranteed that the value $A = \sum_i a_i$ is polynomial in $|V|$. Thus, the space used for solving approximate SUBSETSUM is always bounded by $O(\log |V|)$. Storing the sets W takes $O(h \log |V|)$ space, where h is the size of the smallest ϵ -balanced separator. \square

Note that our running time is mostly dominated by the time we spend enumerating the $\binom{|V|}{k}$ possible subsets of size $k \leq h$. All other computations take polynomial time. For graphs with constant size balanced separators, our algorithm finds a minimum size separator in polynomial time and logarithmic space.

Even the simpler problem of finding minimum size component-wise balanced separators is NP-hard for general graphs [Hyafil and Rivest 1973], and as far as we know, our running time is comparable to the running time of previously known algorithms to find minimum size balanced separators. However, our algorithm is more space efficient than previous solutions.

4. EXTENSIONS TO OTHER VARIANTS OF KNAPSACK

We also consider variants of the more general KNAPSACK problem. Similarly to SUBSETSUM, one can consider decision, search, and optimization versions. The *decision version* of the 0-1 KNAPSACK problem can be defined as follows. Given a capacity $c \in \mathbb{Z}^+$, a target value $t \in \mathbb{Z}^+$, and a set of size m , where each element has value a_i and weight w_i , output YES if and only if there is a subset $S \subseteq [m]$ such that $\sum_{i \in S} a_i = t$ and $\sum_{i \in S} w_i \leq c$. In the search version, we have to output the subset S if it exists. The *optimization version* is defined as follows. Given a capacity $c \in \mathbb{Z}^+$ and a set of size m , where each element has value a_i and weight w_i , find a subset $S \subseteq [1, m]$ such that $\sum_{i \in S} w_i \leq c$ and $\sum_{i \in S} a_i$ is as large as possible. Note that the optimal value is not known in advance. An equivalent definition of the optimization version is to maximize $\sum_{i=1}^m a_i x_i$ with the constraints $\sum_{i=1}^m w_i x_i \leq c$ and $x_i \in \{0, 1\}$ for $1 \leq i \leq m$. Note that we get the SUBSETSUM problem as a special case of 0-1 KNAPSACK by letting $w_i = a_i$ and $c = t$. A common generalization of the 0-1 KNAPSACK problem is the *integer KNAPSACK problem*, where every x_i is a nonnegative integer. This is also known as the *unbounded KNAPSACK problem* since the number of times we can pick each item is unbounded.

An interesting variant is the *change-making problem*. The decision version is defined as follows. Given positive integers a_i for $1 \leq i \leq m$, t , and c , output YES if and only if $\sum_{i=1}^m a_i x_i = t$ and $\sum_{i=1}^m x_i \leq c$, where $x_i \geq 0$ is an integer. This corresponds to paying a fixed amount t using a limited number of coins, where a_i are the coin denominations. Note that this is a special case of the unbounded KNAPSACK problem by letting $w_i = 1$. The change-making problem was first studied by Lueker [1975], who showed that the problem is NP-hard. It was later studied by Kozen and Zaks [1994], and Adamaszek

and Adamaszek [2010]. Additionally, see Martello and Toth [1990] and Kellerer et al. [2004] for comprehensive coverage of the problem.

As we noted in Remark 2.8, the algorithm in Theorem 2.7 finds the colexicographically first subset with acceptable weight. This yields the property that the algorithm reads the shortest possible prefix of the input: it has to read at least as far as the last element of the subset with the earliest last element with respect to the given ordering of the input.

However, this does not mean that the algorithm finds a good subset with the smallest possible number of elements. For example, there may be a good subset of size k among the first r input elements, but the last $k' < k$ elements might also form a good subset, which will not be found, as the algorithm will not read further than the first r elements of the input sequence. This may be reasonable if faster running time is more important.

We can modify our algorithm to always find the smallest possible good subset, but for this we always have to read the whole input, so this version will never run in sublinear time.

The problem of finding a good subset with fewest elements is in fact another variant of the more general KNAPSACK problem: the search version of the 0-1 KNAPSACK problem with weights $w_i = 1$ and capacity $c = k$ corresponds to finding a good subset of size at most k .

4.1. 0-1 Knapsack with Equal Weights

To keep the notation simple, we present our results for $w_i = 1$, but they directly generalize to any KNAPSACK instance where all weights are equal.

THEOREM 4.1. *There is a deterministic one-pass streaming algorithm that, given an integer $m \geq k > 0$, a finite set $S = \{a_1, \dots, a_m\}$ of m positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$, decides if there is a subset $S' \subseteq S$ of size at most k such that*

$$(1 - \epsilon)t \leq \sum_{a_i \in S'} a_i \leq (1 + \epsilon)t.$$

The algorithm runs in $O(\frac{1}{\epsilon}kn)$ time and $O(\frac{1}{\epsilon}k \log t + \log m)$ space on Turing machines and in $O(\frac{1}{\epsilon}km)$ time and $O(\frac{1}{\epsilon}k)$ space on unit-cost RAMs.

Moreover, if a good subset of size at most k exists among the first r elements of the set S , the algorithm terminates within $O(\frac{1}{\epsilon}kr)$ steps on the unit-cost RAM model, and $O(\frac{1}{\epsilon}kn_r)$ steps on Turing machines, where $n_r = \sum_{i=1}^r \log a_i + \log t$.

PROOF (SKETCH). Similarly to the ONEPASSAPPROXDECISION algorithm, we maintain a collection of intervals. In addition, for each interval, we assign an index, indicating the length of partial sums associated with the interval. At the beginning of the algorithm, there is only one interval and its index is 0. As we proceed, whenever we form a new interval, its index will be the index of the interval from which the new interval is formed plus 1. We also merge overlapping intervals like before, but now we only merge intervals that have the same index. This way, every partial sum associated with a given interval has the same size, and the index of a given interval corresponds to the number of elements participating in the possible partial sums associated with the interval. More precisely, if the index of an interval is j , and we find that some $a_i \in S$ belongs to the interval, then we know that there is a good sum using a_i and some other j elements a_ℓ with $\ell < i$.

In addition, whenever creating a new interval would result in an index greater than or equal to k , we do not create the interval and continue with the algorithm. This is

because creating an interval with index at least k would only be useful for detecting sums of size greater than k .

This way, the space bound is increased by a factor of k , and we maintain up to $\frac{1}{\epsilon}k$ intervals. The algorithm stops with YES if it finds that some a_r belongs to one of the existing intervals. In this case, a_r is the last element read from the input, and the rest of the input is never read. The algorithm stops with NO if it read all of the input and none of the elements belonged to an existing interval. The bounds on running time follow since we read each input element at most once, and for each input we perform at most $O(\frac{1}{\epsilon}k)$ comparisons to check if it belongs to one of the intervals. \square

We can use this algorithm to find the smallest good subset without knowing its size in advance in a space-efficient way.

THEOREM 4.2. *There is a deterministic algorithm that, given a finite set $S = \{a_1, \dots, a_m\}$ of m positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$, outputs the smallest subset $S' \subseteq S$ such that*

$$(1 - \epsilon)t \leq \sum_{a_i \in S'} a_i \leq (1 + \epsilon)t$$

if such subset exists, and otherwise it outputs NO. If the size of the smallest good subset is k , the algorithm runs in $O(\frac{1}{\epsilon}k^2n)$ time and $O(\frac{1}{\epsilon}k \log t + \log m)$ space on Turing machines and in $O(\frac{1}{\epsilon}k^2m)$ time and $O(\frac{1}{\epsilon}k)$ space on unit-cost RAMs.

PROOF (SKETCH). We run the algorithm from Theorem 4.1 repeatedly for increasing k' starting from $k' = 1$. If a good set of size k' exists, the algorithm will stop with YES. If the smallest good set has size k , we will get a YES answer during the k -th repetition of running the algorithm.

After this, another $k - 1$ repetition allows us to output the good set, similarly to the proof of Theorem 2.7, by running the algorithm of Theorem 4.1 repeatedly, this time for decreasing k' starting from $k' = k - 1$. \square

4.2. Change-Making Problem: Unbounded Knapsack with Equal Weights

In the SUBSETSUM problem, as well as in the 0-1 KNAPSACK problem, each element a_i can be used at most once to contribute to the sum. The unbounded KNAPSACK problem allows the use of the same element a_i several times toward the sum.

If such repetitions are allowed, then we can save a factor of k in the space bounds, where k is the number of elements (counted with multiplicities) contributing toward the sum. The running time remains the same, but there is a catch: in the case of 0-1 KNAPSACK, the size k of a good subset must be always at most m , as each element can be used at most once. Hence, our algorithms for 0-1 KNAPSACK with equal weights are always polynomial time. But if repetitions are allowed, then in general k may be much larger than m , so our algorithms are not polynomial time in general: they are polynomial in k and n (and in k and m on unit-cost RAMs). On the other hand, since the space bound does not depend on k , they run in constant space on unit-cost RAMs for any constant approximation factor.

This version of unbounded KNAPSACK with weights $w_i = 1$ is sometimes called the *change-making problem* (see Section 1). In the terminology of the change-making problem, k (the size of the subset counted with multiplicities) corresponds to the number of coins used. To keep the notation simple, we present our results for $w_i = 1$ —that is, for the change-making problem—but they directly generalize to any unbounded KNAPSACK instance where all weights are equal.

THEOREM 4.3. *Given an integer $k > 0$, there is a deterministic k -pass streaming algorithm that, given a finite set $S = \{a_1, \dots, a_m\}$ of m positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$, decides if there is a subset $S' \subseteq S$ such that*

$$(1 - \epsilon)t \leq \sum_{a_i \in S'} a_i x_i \leq (1 + \epsilon)t,$$

where each $x_i \geq 0$ is an integer and $\sum_{a_i \in S'} x_i \leq k$. The algorithm runs in $O(\frac{1}{\epsilon}kn)$ time and $O(\frac{1}{\epsilon} \log t + \log m)$ space on Turing machines, and in $O(\frac{1}{\epsilon}km)$ time and $O(\frac{1}{\epsilon})$ space on unit-cost RAMs.

PROOF (SKETCH). The proof is similar to the proof of Theorem 4.1, but now we only need to maintain $O(\frac{1}{\epsilon})$ intervals at a time. First we run the algorithm from Theorem 4.1 once with $k = 1$. Unless a good subset of size 1 exists, this will read the whole input without finding an element that belongs to the initial interval $[l, r]$. We keep the intervals created by this run in working memory; note that each of these intervals has index 1, except the initial interval $[l, r]$, which has index 0. At this point, we discard the interval with index 0, as we will not need it any longer. Next we proceed similarly to running the algorithm of Theorem 4.1 with $k = 2$, but now we already have a collection of intervals with index 1 stored in memory. Every time we read a new element a_i , we check if it belongs to one of the intervals with index 1. If YES, we stop with YES. If not, we create a new interval, but only using the intervals with index 1, and thus we only create intervals with index 2 during this pass. If we do not succeed in finding an element that belongs to one of the intervals during this pass, we eliminate all intervals with index 1 but keep all intervals with index 2.

During the j -th pass, we will check if the elements read belong to one of the intervals with index $j - 1$ that we kept from the previous pass. During the j -th pass, we only create new intervals with index j . At the end of the j -th pass, we discard the intervals with index $j - 1$ and keep all intervals with index j .

Thus, at any time during the process, we keep at most $2\frac{1}{\epsilon}$ intervals. Each pass of the input will take $O(\frac{1}{\epsilon}n)$ time, and we make at most k passes. \square

The reason this approach with the smaller space bound could not be used in the proof of Theorem 4.1 is that this way, we cannot prevent including the same element a_i more than once into the sum. The algorithm in Theorem 4.1 is a one-pass algorithm, so we never even consider the same element more than once. But in the algorithm in the proof of Theorem 4.3, we keep restarting reading the input, and we have no way of knowing if the currently read element a_i has already contributed toward one of the partial sums associated with a given interval.

We can use this algorithm to find the smallest possible k (i.e., the smallest number of coins in the change-making problem) without knowing its value in advance in a space-efficient way.

THEOREM 4.4. *There is a deterministic algorithm that, given a finite set $S = \{a_1, \dots, a_m\}$ of m positive integers, $t \in \mathbb{Z}^+$, and $0 < \epsilon < 1$, outputs the smallest integer k and a sequence b_1, \dots, b_k such that $b_i \in S$ for $i = 1, \dots, k$ and*

$$(1 - \epsilon)t \leq \sum_{i=1}^k b_i \leq (1 + \epsilon)t.$$

The algorithm runs in $O(\frac{1}{\epsilon}k^2n)$ time and $O(\frac{1}{\epsilon} \log t + \log m)$ space on Turing machines, and in $O(\frac{1}{\epsilon}k^2m)$ time and $O(\frac{1}{\epsilon})$ space on unit-cost RAMs.

PROOF (SKETCH). Similarly to the proof of Theorem 4.2, we run the algorithm from Theorem 4.3 repeatedly for increasing k' starting from $k' = 1$. If a good set of size k' exists, the algorithm will stop with YES. If the smallest good set has size k , we will get a YES answer during the k -th repetition of running the algorithm.

After this, another $k - 1$ repetition allows us to output the good set, similarly to the proof of Theorem 2.7, by running the algorithm of Theorem 4.3 again repeatedly this time for decreasing k' starting from $k' = k - 1$, \square

5. EXACT UNARY SUBSET SUM IN THE STREAMING MODEL

We first show lower bounds for one-pass and many-pass streaming algorithms that decide USS. We do so by lower bounding a communication game corresponding to SUBSETSUM. A communication game corresponding to SUBSETSUM is as follows. Both Alice and Bob are given an integer B . Further, each of them has a set of numbers, and they have to determine if there is a subset of numbers among the union of their sets that adds to B . The goal is to minimize the number of bits exchanged between Alice and Bob. Additionally, there may be constraints on how often the communication exchange changes direction (the number of rounds).

A standard lower-bound technique (see Alon et al. [1999]) shows that an s -pass space $O(S)$ streaming algorithm yields a protocol with communication complexity $O(sS)$ and $2s - 1$ rounds. Thus, a communication complexity lower bound yields a streaming space lower bound. We use this technique to show that any one-pass streaming algorithm for USS needs $\Omega(\sqrt{N})$ space.

LEMMA 5.1. *Any deterministic or randomized one-pass streaming algorithm for USS uses space $\Omega(\sqrt{N})$.*

PROOF. We reduce the INDEX problem to SUBSETSUM. The INDEX_n function is defined as follows. Alice has $x \in \{0, 1\}^n$, and Bob has an index $k \in [n]$. The goal is to find x_k . Alice can send one message to Bob, after which Bob should announce what he believes is the value of x_k . It is known that the one-way randomized communication complexity of INDEX_n is $\Theta(n)$ (see Bar-Yossef et al. [2002] or Kremer et al. [1995]).

The reduction from INDEX_n to SUBSETSUM is as follows. The target sum is $B = 2n$. Alice creates a set $S = \{2n - i \mid x_i = 1\}$. Bob creates the set $T = \{k\}$. Since each number in S is at least n , any subset of S that has two or more numbers would have a sum strictly greater than $2n$. Hence, any subset of $S \cup T$ that has a sum of $2n$ can have at most one number from S . Now it is easy to see that if $x_k = 1$, the subset $\{(2n - k), k\}$ has sum $2n$. However, if $x_k = 0$, there is no subset of $S \cup T$ that has sum $2n$. Thus, a protocol that correctly decides the SUBSETSUM game instance where $B = 2n$, Alice has S and Bob has T , with communication cost c also correctly decides $\text{INDEX}_n(x, k)$ with communication cost c .

Assume that there is a space $f(N)$ one-pass streaming algorithm for USS, where f is some nondecreasing function. Then there is a cost $f(N)$ protocol for the SUBSETSUM game, and hence by the preceding reduction a cost $f(\ell)$ protocol for INDEX_n , where ℓ is the length of the resulting USS instance. Since $\ell \leq n^2$, $f(\ell) \leq f(n^2)$. By the lower bound for INDEX, $f(n^2) \in \Omega(n)$, and so $f(N) \in \Omega(\sqrt{N})$. \square

A generalization of the preceding proof gives a space lower bound for streaming USS that depends on the number of passes.

THEOREM 5.2. *Any deterministic or randomized s -pass streaming algorithm for USS uses space $\Omega(\sqrt{N}/s)$.*

PROOF. We give a reduction from $\overline{\text{DISJ}}_n$ to SUBSETSUM. The disjointness problem DISJ_n is defined as follows. For $x, y, \in \{0, 1\}^n$, $\text{DISJ}_n(x, y) = \bigwedge_{i=1}^n \neg(x_i \wedge y_i)$. (In other words, if x and y are characteristic vectors of sets $X, Y \subseteq [n]$, then $\text{DISJ}_n(x, y) = 1$ if and only if $X \cap Y = \emptyset$.) Its complement $\overline{\text{DISJ}}_n$ is the intersection problem. Alice and Bob are given $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$, respectively. The goal is to determine if there exists an $i \in [n]$ such that $x_i = y_i = 1$. It is known [Kalyanasundaram and Schnitger 1992; Razborov 1992; Bar-Yossef et al. 2004] that any randomized protocol for $\overline{\text{DISJ}}_n$, with any number of rounds, must exchange $\Omega(n)$ bits to bound error probability by $1/3$.

The reduction from $\overline{\text{DISJ}}_n$ to SUBSETSUM is as follows. We set the target to $B = 12n - 1$. Alice constructs the set $S = \{8n - 2i \mid x_i = 1\}$. Bob constructs the set $T = \{4n + 2i - 1 \mid y_i = 1\}$. Notice that all numbers in S are greater than $B/2$ and that all numbers in T lie in the interval $(B/3, B/2)$. Further note that each number in S is even and that each number in T is odd. We claim that $\overline{\text{DISJ}}_n = 1$ exactly when $S \cup T$ has a subset adding to B . To see why, first observe the following:

- (1) Using numbers only from S cannot give a sum of B since B itself does not appear in S , and the sum of any two numbers from S exceeds B .
- (2) Using numbers only from T cannot give a sum of B since
 - (a) B does not appear in T ;
 - (b) Any two numbers in T add to an even number greater than $B/2$, but B is odd; and
 - (c) adding three or more numbers from T gives a sum greater than B .

Thus, we see that if any subset of $S \cup T$ adds to B , then it must contain exactly one number from S and one from T . In other words, it must be of the form $\{8n - 2i, 4n + 2j - 1\}$. To add up to $12n - 1$, it must be the case that $i = j$. Hence, such a subset exists if and only if there exists an $i \in [n]$ such that $x_i = y_i = 1$.

Now, as in Lemma 5.1, assume that there is an s -pass streaming algorithm for USS using space $f(N)$, where $f(\cdot)$ is some nondecreasing function. Then there is a cost $(2s - 1)f(N)$ protocol for the SUBSETSUM game with $2s - 1$ rounds, and hence by the preceding reduction, a cost $(2s - 1)f(n^2)$ protocol for $\overline{\text{DISJ}}_n$. (The reduction from $\overline{\text{DISJ}}_n$ creates USS instances of length $N \in O(n^2)$.) By the lower bound for $\overline{\text{DISJ}}_n$, $(2s - 1)f(n^2) \in \Omega(n)$, and so $f(n) \in \Omega(\sqrt{N}/s)$. \square

Next we show that the preceding lower bounds are close to tight by giving a deterministic two-pass algorithm using $O(\sqrt{N} \log N)$ space.

THEOREM 5.3. *There is a deterministic two-pass streaming algorithm for USS that uses space $O(\sqrt{N} \log N)$ and runs in polynomial time.*

PROOF. In the first pass, we simply read the input to learn its length—that is, the value of N , as well as the value B . We write these down using $O(\log N)$ space. Next we initialize an array C of size \sqrt{N} to all zeros. We will use the entries of this array as counters, with values at most N , so we reserve $\sqrt{N} \log N$ space for the entries of the array. In addition, we write a list of numbers using space not occupied by the array.

In the second pass, whenever we see a number $a_i \leq \sqrt{N}$, we increment the entry $C[a_i]$ of the array by one. If we see a number $a_i > \sqrt{N}$, we append its value (in binary) to the list after the array.

After finishing the second pass, we have all of the input sequence stored locally in an implicit form using space $O(\sqrt{N} \log N)$. Next we run the algorithm of Kane [2010] on this locally stored input. This will take at most $O(\log N)$ additional space and time polynomial in N . \square

We also observe that there are two simple one-pass streaming algorithms for `SUBSET-SUM` that can be used for `USS` and `SUBSETSUM` when the input is represented in binary. We assume that B is the first number in the stream. The first algorithm uses space $O(m \log B)$ storing those input numbers that do not exceed B . The other algorithm uses space $O(B \log m)$ storing how many times each integer in $[B]$ appears in the input. Then both algorithms run Kane's algorithm [Kane 2010] on the locally stored implicit input representation. However, note that these bounds are incomparable with \sqrt{N} in general.

6. CIRCUIT COMPLEXITY OF MONOTONE APPROXIMATE USS

6.1. Problem Formulation

In this section, we consider the monotone circuit complexity of `USS`. Without the monotone restrictions, it is known that `USS` is complete for the circuit class TC^0 ([Elberfeld et al. 2012]). However, in a very natural sense, `SUBSETSUM` is a monotone problem, and so we can consider monotone circuits for it. The encoding of the input becomes crucial for achieving monotonicity. We choose the following encoding.

For each positive integer B , the input consists of the frequency of each number in the stream in unary. An instance $w \in \{0, 1\}^{B^2}$ consists of B blocks of B bits each. For each $k \in [B]$, if k occurs in the stream m_k times, then the k -th block w_k has exactly m_k ones—that is, $\sum_{j=1}^B w_{kj} = m_k$. Thus, the input records the multiplicity of each number in $[B]$ (we assume that no multiplicity exceeds B).

Define the transitive relation \leq : for $u = (u_1, u_2, \dots, u_B)$, $v = (v_1, v_2, \dots, v_B)$ with $u_k, v_k \in \{0, 1\}^B$, $u \leq v$ if and only if $\forall k \in [B]$, $\sum_{j=1}^B u_{kj} \leq \sum_{j=1}^B v_{kj}$.

We define the multiplicity-`USS` problem, denoted as `mUSS`, and its approximation variant ϵ -`mUSS`, as follows:

$$\begin{aligned} \text{mUSS}(w, B) = 1 &\iff \exists y = (y_1, y_2, \dots, y_B) : \\ &y_k \in \{0, 1\}^B \forall k \in [B], y \leq w, \text{ and} \\ &B = \left(\sum_{k=1}^B k \left(\sum_{j=1}^B y_{kj} \right) \right) \\ \epsilon\text{-mUSS}(w, B) = 1 &\iff \exists y = (y_1, y_2, \dots, y_B) : \\ &y_k \in \{0, 1\}^B \forall k \in [B], y \leq w, \text{ and} \\ &B(1 - \epsilon) \leq \left(\sum_{k=1}^B k \left(\sum_{j=1}^B y_{kj} \right) \right) \leq B(1 + \epsilon). \end{aligned}$$

We call such a y a *witness* for (w, B) . The vector y represents a subset of the multiset represented by w such that the elements in y sum to B (or to a number within ϵ of B , respectively).

For example, for $B = 4$, the stream 1 3 2 2 1 4 3 can be encoded by any of the following strings (and by many more): 1100 1100 1100 1000, 1010 0101 0011 0010. Some witnesses for this instance are 1100 1000 0000 0000 (use two ones and a two), 0100 0000 0001 0000 (use a one and a three), 0000 0000 000 1000 (use a four).

FACT 6.1. *mUSS is a monotone function. For instance, for each positive integer B , and for each $u = (u_1, u_2, \dots, u_B)$, if $\text{mUSS}(u, B) = 1$, and if $v = (v_1, v_2, \dots, v_B)$ is obtained from u by changing some zeros to ones, then $\text{mUSS}(v, B) = 1$.*

Similarly, for each ϵ and B , ϵ -mUSS is a monotone function.

6.2. Monotone Depth Upper Bound

It has been known for more than three decades [Monien and Sudborough 1980] that USS is in nondeterministic logspace; hence, USS reduces to the problem Reach defined as follows:

Given: A layered directed acyclic graph G , two designated nodes s, t .
Output: One if there is a path from s to t in G , and zero otherwise.

It is well known that Reach has monotone circuits of depth $O(\log^2 n)$, where n is the number of vertices in the input instance. (This follows from the construction of Savitch [1970]. For example, see Arora and Barak [2009].) We show that with the encoding described earlier, exact and approximate versions of mUSS reduce to Reach via a very weak form of reduction known as monotone projections.

$f : A \rightarrow B$ is called a *monotone projection* if $\forall x \in A, y \in B$ such that $f(x) = y$, it holds that $\forall i, \exists j$ such that $y_i = x_j$. In other words, a monotone projection produces an output string by either rearranging bits of the input or copying bits of the input or both.

Showing a monotone projection from mUSS to Reach gives us small-depth monotone circuits for mUSS.

THEOREM 6.2. *For every positive integer B , mUSS(\cdot, B) and ϵ -mUSS(\cdot, B) have monotone circuits of depth $O(\log^2 B)$.*

PROOF. We prove this by reducing an instance of mUSS into an instance of Reach via a monotone projection.

The following is the high-level idea. For every integer B , and given $w \in \{0, 1\}^{B^2} = (w_1, w_2, \dots, w_B)$, we create a graph with $B^2 + 1$ layers. The zero-th layer consists of the source vertex, and the other B^2 layers have $(B + 1)$ vertices each. We further partition the set of B^2 layers into B blocks of B consecutive layers each.

Let $v_{j,k}^i$ denote the i -th vertex in the layer j in the block k . Intuitively, each layer corresponds to a bit position in the input string. We add edges to ensure that a vertex $v_{k,j}^i$ is reachable from the source vertex if and only if the stream corresponding to the first $k - 1$ blocks of w and j bits from the k -th block has a subset that adds to i .

If after reading l bits of w there is a subset that adds to i , then this subset continues to exist even after reading more bits. To capture this phenomenon, we add horizontal edges from every vertex v in layer l to the copy of v in layer $l + 1$.

If the bit $w_{kj} = 1$, then using this copy of k , for each existing subset sum s , the subset sum $s + k$ can also be created. To capture this, we include slanted edges from each $v_{j,k}^i$ to $v_{j+1,k}^{i+k}$.

Thus, there is a path from the source to vertex $v_{B,B}^i$ exactly when there is a subset that sums to i . By connecting $v_{B,B}^i$ for appropriate i to a new target node t , we reduce mUSS or ϵ -mUSS to Reach.

Formally, the graph for reducing mUSS to Reach is defined as follows. Let $v_{j,k}^i$ denote the i -th vertex in the layer j in the block k , where $0 \leq i \leq B$, and $j, k \in [B]$. Let $V = \{v_{j,k}^i \mid 0 \leq i \leq B, 1 \leq j, k \leq B\} \cup \{v_{0,0}^0\}$. We now describe the edge set for the graph. We first describe the horizontal edges. Let $E_1 = \{v_{j,k}^i \rightarrow v_{j+1,k}^i \mid 0 \leq i \leq B, 1 \leq j \leq B - 1, 1 \leq k \leq B\}$ be the edges that connect the i -th vertex in the layer j of the block k to the i -th vertex in the layer $j + 1$ of the same block. Let $E_2 = \{v_{B,k}^i \rightarrow v_{1,k+1}^i \mid 0 \leq i \leq B, 1 \leq k \leq B - 1, \}$ be the interblock edges that connect the i -th vertex in the layer B of the block k to the i -th vertex in the first layer of the block $k + 1$. Let E_1'' and E_2'' denote

sets of intrablock and interblock edges obtained from the input instance w of mUSS, respectively (these are the slant edges):

$$E_1^w = \{v_{j,k}^i \rightarrow v_{j+1,k}^{i+k} \mid w_{jk} = 1, 0 \leq i \leq B-k, 1 \leq j \leq B-1, 1 \leq k \leq B\}$$

$$E_2^w = \{v_{B,k}^i \rightarrow v_{1,k+1}^{i+k} \mid w_{Bk} = 1, 0 \leq i \leq B-k, 1 \leq k \leq B-1\}$$

The Reach instance that we create is the following:

$$(G = (V, E = E_1 \cup E_2 \cup E_1^w \cup E_2^w \cup \{v_{0,0}^0 \rightarrow v_{1,1}^0\}), s = v_{0,0}^0, t = v_{B,B}^B).$$

CLAIM 6.3. *After reading the j -th bit of w_k , if $s_1, s_2, \dots, s_\alpha$ are all possible subset sums that have sum at most B , then the set of vertices reachable from $v_{0,0}^0$ in the j -th layer of the block k is $\{v_{j,k}^{s_1}, v_{j,k}^{s_2}, \dots, v_{j,k}^{s_\alpha}\}$.*

PROOF. We prove this by induction on j, k . The base case is $j = 1, k = 1$ (i.e., the first bit of w_1 is read). Either $w_{1,1} = 0$ or $w_{1,1} = 1$. In the first case, the possible subset sum is zero. And from our construction, $v_{0,0}^0 \rightarrow v_{1,1}^0$ is an edge in G . In the latter case, the possible subset sums are zero and one. As $w_{1,1} = 1$, there is an edge $v_{0,0}^0 \rightarrow v_{1,1}^1$ in E_1^w . And $v_{0,0}^0 \rightarrow v_{1,1}^0$ is also an edge in E . Hence, the base case.

After reading the j -th bit in w_k , let the possible subset sums be $s_1, s_2, \dots, s_\alpha$. By induction hypothesis, the set of vertices reachable from $v_{0,0}^0$ in the j -th layer of the block k is $\{v_{j,k}^{s_1}, v_{j,k}^{s_2}, \dots, v_{j,k}^{s_\alpha}\}$. The inductive step now proceeds as follows.

Case 1 ($j < B$, i.e., next bit in the same block). Either $w_{j+1,k} = 0$ or 1. If $w_{j+1,k} = 0$, then the possible subset sums after reading $j+1$ -th bit remain unchanged (i.e., $s_1, s_2, \dots, s_\alpha$). By our construction, $\{v_{j,k}^{s_1} \rightarrow v_{j+1,k}^{s_1}, v_{j,k}^{s_2} \rightarrow v_{j+1,k}^{s_2}, \dots, v_{j,k}^{s_\alpha} \rightarrow v_{j+1,k}^{s_\alpha}\}$ is a subset of E_1 . By the induction hypothesis, $v_{j,k}^{s_1}, v_{j,k}^{s_2}, \dots, v_{j,k}^{s_\alpha}$ are reachable from $v_{0,0}^0$. Therefore, $v_{j+1,k}^{s_1}, v_{j+1,k}^{s_2}, \dots, v_{j+1,k}^{s_\alpha}$ are reachable. No other vertex in layer $j+1$ is reachable (as E_1^w has no edges in this layer).

If $w_{j+1,k} = 1$, then the possible subset sums after reading $j+1$ -th bit are $s_1, s_2, \dots, s_\alpha, s_1+k, s_2+k, \dots, s_\alpha+k$. (If $s_l+k > B$ for some l , then s_l+k need not be considered.) By our construction, $\{v_{j,k}^{s_1} \rightarrow v_{j+1,k}^{s_1+k}, v_{j,k}^{s_2} \rightarrow v_{j+1,k}^{s_2+k}, \dots, v_{j,k}^{s_\alpha} \rightarrow v_{j+1,k}^{s_\alpha+k}\}$ is a subset of E_1^w . And $\{v_{j,k}^{s_1} \rightarrow v_{j+1,k}^{s_1}, v_{j,k}^{s_2} \rightarrow v_{j+1,k}^{s_2}, \dots, v_{j,k}^{s_\alpha} \rightarrow v_{j+1,k}^{s_\alpha}\}$ is a subset of E_1 . Using the induction hypothesis and these edges, the inductive step follows.

Case 2 ($j = B$, i.e., next bit in block $k+1$). The proof of this case is similar to that of Case 1. The arguments goes through by using the induction hypothesis and the interblock edges from sets E_2, E_2^w . \square

For the given instance, $\text{mUSS}(w, B) = 1$ if and only if after reading the last bit of w_B there is a subset that adds to B if and only if $v_{B,B}^B$ is one of the reachable nodes from $v_{0,0}^0$ (using Claim 6.3). As the graph G has $B^2(B+1) + 1 = O(B^3)$ vertices, Reach can be computed by a $O(\log^2 B)$ depth monotone circuit. Since the reduction can be performed using projections, the theorem follows. \square

Remark 6.4. A more direct way of seeing this was pointed out to us by an anonymous referee and is as follows. For a set $\{0\} \subseteq U \subseteq \{0, 1, \dots, n\}$, let χ_U be its characteristic string of length $n+1$ (so the i -th bit of χ_U is one if and only if $i \in U$). Let $U + V = \{x + y \mid x \in U \wedge y \in V\}$ be the ‘‘addition’’ of two sets. We can compute χ_{U+V} from χ_U and χ_V using a monotone constant depth circuit:

$$\chi_{U+V}(i) = \bigvee_{j=0}^i [\chi_U(j) \wedge \chi_V(i-j)].$$

Now, to test whether S contains a subset summing up to some t , consider the sets $\{\{0, x\} \mid x \in S\}$ and “sum up” their characteristic strings in a tree-like fashion. Then the result has a one at position t if and only if S has a subset summing up to t . Furthermore, we can compute this using a monotone circuit of polynomial size and $O(\log^2 B)$ depth. (If unbounded fanin \vee gates are allowed, it has logarithmic depth.)

6.3. Monotone Approximate Subset Sum as a Circuit Primitive

We now examine the power of ϵ -approximators for mUSS when used as a primitive to compute other functions. Nisan [1994] showed that any circuit for DISJ_n using linear threshold gates requires $\Omega(n/\log n)$ gates. We introduce a new kind of gate, an ApproxUSS gate, that we show is at least as powerful as a threshold or majority gate, and show that any circuit that uses ApproxUSS gates to compute disjointness needs size $\Omega(n/\log n)$. However, we do not know whether linear threshold gates can simulate ApproxUSS gates with at most sublogarithmic blowup in the number of gates or vice versa.

We define approximate USS gates, denoted ApproxUSS, as gates that solve the ϵ -mUSS problem defined in Section 6. An $\text{ApproxUSS}_{\epsilon, B}$ gate takes a bit string x of length B^2 as input and outputs one exactly when $\epsilon\text{-mUSS}(x, B) = 1$.

Although it is trivial to see that majority can be computed with a single call to an oracle for mUSS, it is not immediately clear that oracle access to ϵ -mUSS when $\epsilon > 0$ is also sufficient. We show that this is indeed the case by showing that ApproxUSS gates are at least as powerful as standard threshold gates. Specifically, we show that an ApproxUSS gate can simulate majority with only a polynomial blowup in the number of wires.

LEMMA 6.5. *The MAJ_{2n+1} function can be computed by an $\text{ApproxUSS}_{\epsilon, B}$ gate with $B = O(n^3)$ and a suitable nonzero value for ϵ .*

PROOF. Given $w \in \{0, 1\}^{2n+1}$ as an input to MAJ, we come up with a vector $y \in \{0, 1\}^{B^2}$ such that $\text{MAJ}(w) = 1$ if and only if $\text{ApproxUSS}_{\epsilon, B}(y) = 1$. We will specify B and ϵ , and also another parameter N , shortly. The reduction from w to y is a simple projection given as follows.

The string y that we wish to compute represents a stream of numbers. We describe the projection in terms of the stream. If $w_i = 1$, then we add $\lfloor N/i \rfloor$ copies of i to the stream. So $y_{i,j} = 1$ exactly when $w_i = 1$ and $j \leq \lfloor N/i \rfloor$. Thus, if $w_i = 0$, adding all copies of i in y gives zero, whereas if $w_i = 1$, adding all copies of i in y gives $i \times \lfloor N/i \rfloor$. Note that for every i , $N - 2n \leq N - (i - 1) \leq i \times \lfloor N/i \rfloor \leq N$.

Let $\alpha = nN$, $\beta = (n + 1)(N - 2n)$, and $\gamma = (2n + 1)N$. Let M denote the maximum subset sum in the stream corresponding to y . (Add up *all* numbers in the stream.) If w is a NO instance of MAJ, then $M \leq \alpha$. If w is a YES instance of MAJ, then $\beta \leq M \leq \gamma$.

To turn this into a proper reduction, it suffices to ensure that

- $N \leq B$ so that it is possible to create $\lfloor N/i \rfloor$ copies of i for each i ,
- $\alpha < \beta$ so that we can distinguish between YES and NO instances, and
- the range $[\beta, \gamma]$ coincides with $[B - B\epsilon, B + B\epsilon]$ so that a YES instance of MAJ gets mapped to a YES instance of ApproxUSS.

We see that the preceding conditions are met by choosing $N = 2(n + 1)^2$, $B = (\beta + \gamma)/2$, and $\epsilon = \frac{\gamma - \beta}{2B} = \frac{\gamma - \beta}{\gamma + \beta}$. This gives $\beta = 2n^3 + o(n^3)$, $\gamma = 4n^3 + o(n^3)$, and hence $B = 3n^3 + o(n^3)$, and $\epsilon = \frac{2n^3 + o(n^3)}{6n^3 + o(n^3)} \sim 1/3$. \square

On the other hand, it is not known whether ApproxUSS, for $\epsilon \neq 0$, can be decided with a single oracle call to majority. It is conceivable that demanding a YES answer for

a wider range of inputs (the approximation) makes the problem harder. It is therefore interesting to examine the power of circuits using ApproxUSS gates. We pursue this thread in the following.

The communication problem ccApproxUSS corresponding to ApproxUSS can be described as follows. Let $S \subseteq [B^2]$. Both Alice and Bob know S , B , and ϵ . Alice knows the bits x_i for $i \in S$, and Bob knows the remaining bits of x . They must decide whether $\text{ApproxUSS}_{\epsilon, B}(x) = 1$ —that is, whether ϵ -mUSS(x , B) = 1.

In Theorem 2.1, we proved that for every ϵ , there is a one-pass streaming algorithm that ϵ -approximates USS using space $O((\log B)/\epsilon)$. The algorithm works for every possible ordering of the numbers in the input stream. This implies that there is a $O((\log B)/\epsilon)$ bit one-round protocol for ccApproxUSS for worst-case partitioning of the input (for every $S \subseteq [B^2]$). (The string x determines the multiset of numbers. For any partition of x , Alice and Bob can construct a stream of numbers forming this multiset, where Alice has the initial part of the stream and Bob has the latter part. Treat the indices in B^2 as pairs k, j , where k is the block number and j is the index within the block. Alice includes in her stream a copy of k for each bit $x_{kj} = 1$ in her part. Bob does the same.) Therefore, using an argument similar to that of Nisan [1994], we can prove the following lemma.

LEMMA 6.6. *Let C be a circuit that computes DISJ_n using s $\text{ApproxUSS}_{\epsilon, B}$ gates, where $\epsilon \in \Theta(1)$, and the value of B at each $\text{ApproxUSS}_{\epsilon, B}$ is bounded above by a polynomial in n . Then $s \in \Omega(n/\log n)$.*

PROOF. Let C be such a circuit, with s ApproxUSS gates. Let t denote the maximum value of B in any of the gates. We use C to obtain a protocol for DISJ_n as follows. Alice and Bob evaluate C bottom-up, reaching a gate only after all of its children have been evaluated. At each ApproxUSS gate, we know that $\log B \in O(\log t) \subseteq O(\log n)$. When an ApproxUSS gate has to be evaluated, an efficient protocol of $O((\log t)/\epsilon)$ bits for ccApproxUSS is invoked with the appropriate partition of the inputs of the gate. As there are s ApproxUSS gates, the entire protocol for computing DISJ_n uses $O(((\log t)/\epsilon) \times s)$ bits of communication. However, we know that any protocol for DISJ_n requires $\Omega(n)$ bits of communication ([Kalyanasundaram and Schnitger 1992; Razborov 1992; Bar-Yossef et al. 2004]). Hence, $s \log t = \Omega(\epsilon n)$. By assumption, $\epsilon \in \Theta(1)$ and $\log t = O(\log n)$. Hence, $s = \Omega(n/\log n)$. \square

7. OPEN PROBLEMS

We now discuss a few aspects of our results and some open questions:

- Our algorithms work for multiset versions, where an element may appear in the input stream more than once. If the multiplicities of all numbers are restricted to be between $\{0, 1\}$, then the problem does not become easier. In fact, our lower bound proof for USS in Section 5 (Theorem 5.2) generates such instances.
- Our algorithms for 0-1 KNAPSACK with equal weights are fully polynomial time, but the space bounds depend on k , the size of the smallest good subset. Thus, in general, the space bounds for this problem may be linear in m if k is large (e.g., constant fraction of m). As we noted earlier, our algorithms for unbounded KNAPSACK with equal weights eliminate the dependence on k in the space bounds, but the running time is not polynomial, as in this version of the problem k does not have to be bounded by m . It remains open if KNAPSACK can be approximated in polynomial time and $o(m)$ space simultaneously, even in the case of equal weights.
- We know that USS is in TC^0 [Elberfeld et al. 2012], and we have proved that mUSS is in monotone NC^2 (polynomial-size \log^2 depth circuits). It is known that there exists a monotone formula of polynomial size that cannot be computed by constant

depth polynomial-size monotone threshold circuits [Yao 1989]. However, the question of whether monotone TC^0 is contained in monotone NC^1 is open (i.e., see Servedio [2004]). Majority is known to be in monotone NC^1 [Valiant 1984]. Moreover, it is easy to observe that majority reduces to USS. In the light of these results, the question of whether mUSS is contained in monotone NC^1 is interesting.

ACKNOWLEDGMENTS

Some of the collaboration on this work was initiated and continued in the excellent setting of Dagstuhl Seminars 12421 and 14391 on Algebraic and Combinatorial Methods in Computational Complexity and Dagstuhl Seminar 14121 on Computational Complexity of Discrete Problems.

The authors thank the anonymous referees, whose comments helped to improve the presentation of this article.

REFERENCES

- Anna Adamaszek and Michal Adamaszek. 2010. Combinatorics of the change-making problem. *European Journal of Combinatorics*, 47–63.
- Noga Alon, Yossi Matias, and Mario Szegedy. 1999. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences* 58, 1, 137–147.
- Noga Alon, Paul D. Seymour, and Robin Thomas. 1990. A separator theorem for graphs with an excluded minor and its applications. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*. 293–299. DOI: <http://dx.doi.org/10.1145/100216.100254>
- Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press.
- Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar. 2002. Information theory methods in communication complexity. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity*. 93–102.
- Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar. 2004. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences* 68, 4, 702–732.
- Cristina Bazgan, Miklos Santha, and Zsolt Tuza. 2002. Efficient approximation algorithms for the SUBSET-SUMS EQUALITY problem. *Journal of Computer and System Sciences* 64, 2, 160–170. DOI: <http://dx.doi.org/10.1006/jcss.2001.1784>
- Richard Bellman. 1957. *Dynamic Programming*. Princeton University Press, Princeton, NJ.
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. 2009. *Introduction to Algorithms* (3rd ed.). MIT Press, Cambridge, MA.
- George Dantzig. 1957. Discrete-variable extremum problems. *Operations Research* 5, 266–277.
- Michael Elberfeld, Andreas Jakoby, and Till Tantau. 2010. Logspace versions of the theorems of Bodlaender and Courcelle. In *Proceedings of the 2010 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS'10)*. 143–152.
- Michael Elberfeld, Andreas Jakoby, and Till Tantau. 2012. Algorithmic meta theorems for circuit classes of constant and logarithmic depth. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS'12)*. 66–77.
- M. R. Garey and D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco, CA.
- George Gens and Eugene Levner. 1980. Complexity of approximation algorithms for combinatorial problems: A survey. *ACM SIGACT News* 12, 3, 52–65.
- George Gens and Eugene Levner. 1979. Computational complexity of approximation algorithms for combinatorial problems. In *Proceedings of the 8th Symposium of the Mathematical Foundations of Computer Science*. 292–300. DOI: http://dx.doi.org/10.1007/3-540-09526-8_26
- A. C. Gilbert, Y. Kotidis, S. Muthukrishnan, and M. Strauss. 2001. Surfing wavelets on streams: One-pass summaries for approximate aggregate queries. In *Proceedings of 27th International Conference on Very Large Data Bases (VLDB'01)*. 79–88. <http://www.vldb.org/conf/2001/P079.pdf>.
- Laurent Hyafil and Ronald L. Rivest. 1973. *Graph Partitioning and Constructing Optimal Decision Trees are Polynomial Complete Problems*. Technical Report Rapport de Recherche no. 33. IRIA: Laboratoire de Recherche en Informatique et Automatique.

- Oscar H. Ibarra and Chul E. Kim. 1975. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM* 22, 4, 463–468. DOI: <http://dx.doi.org/10.1145/321906.321909>
- T. Imai, K. Nakagawa, A. Pavan, N. V. Vinodchandran, and O. Watanabe. 2013. An $O(n^{1/2+\epsilon})$ space and polynomial-time algorithm for directed planar reachability. In *Proceedings of the IEEE Conference on Computational Complexity*. IEEE, Los Alamitos, CA, 277–286.
- Bala Kalyanasundaram and Georg Schnitger. 1992. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics* 5, 4, 545–557.
- Daniel M. Kane. 2010. Unary Subset-Sum is in logspace. arXiv:10102.1336.
- Richard M. Karp. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*. IBM Research Symposia Series. Springer, 85–103.
- Hans Kellerer, Renata Mansini, Ulrich Pferschy, and Maria Grazia Speranza. 2003. An efficient fully polynomial approximation scheme for the subset-sum problem. *Journal of Computer and System Sciences* 66, 2, 349–370. DOI: [http://dx.doi.org/10.1016/S0022-0000\(03\)00006-0](http://dx.doi.org/10.1016/S0022-0000(03)00006-0)
- Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. *Knapsack Problems*. Springer.
- Dexter Kozen and Shmuel Zaks. 1994. Optimal bounds for the change-making problem. *Theoretical Computer Science* 123, 2, 377–388. DOI: [http://dx.doi.org/10.1016/0304-3975\(94\)90134-1](http://dx.doi.org/10.1016/0304-3975(94)90134-1)
- Ilan Kremer, Noam Nisan, and Dana Ron. 1995. On randomized one-round communication complexity. *Computational Complexity* 8, 596–605.
- Nutan Limaye, Meena Mahajan, and Karteek Sreenivasaiah. 2012. The complexity of unary Subset Sum. In *Computing and Combinatorics*. Lecture Notes in Computer Science, Vol. 7434. Springer, 458–469.
- Richard J. Lipton and Robert Endre Tarjan. 1979. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics* 36, 177–189. DOI: <http://dx.doi.org/10.1137/0209046>
- Richard J. Lipton and Robert Endre Tarjan. 1980. Applications of a planar separator theorem. *SIAM Journal on Computing* 9, 3, 615–627. DOI: <http://dx.doi.org/10.1137/0209046>
- Daniel Lokshtanov and Jesper Nederlof. 2010. Saving space by algebraization. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*. 321–330. DOI: <http://dx.doi.org/10.1145/1806689.1806735>
- George Lueker. 1975. *Two NP-Complete Problems in Nonnegative Integer Programming*. Technical Report 178. Computer Science Laboratory, Princeton University, Princeton, NJ.
- Silvano Martello and Paolo Toth. 1990. *Knapsack Problems: Algorithms and Computer Implementations*. Wiley.
- Gary L. Miller. 1986. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences* 32, 3, 265–279. DOI: [http://dx.doi.org/10.1016/0022-0000\(86\)90030-9](http://dx.doi.org/10.1016/0022-0000(86)90030-9)
- B. Monien and I. H. Sudborough. 1980. The interface between language theory and complexity theory. In *Formal Language Theory*, R.V. Book (Ed.). Academic Press, 287–324.
- S. Muthukrishnan. 2005. Data streams: Algorithms and applications. *Foundations and Trends in Theoretical Computer Science* 1, 2, 117–236. DOI: <http://dx.doi.org/10.1561/0400000002>
- Noam Nisan. 1994. The communication complexity of threshold gates. In *Proceedings of Combinatorics, Paul Erdos Is Eighty*. 301–315.
- Alexander A. Razborov. 1992. On the distributional complexity of disjointness. *Theoretical Computer Science* 106, 2, 385–390.
- Omer Reingold. 2008. Undirected connectivity in log-space. *Journal of the ACM* 55, 4, Article No. 17.
- W. Savitch. 1970. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences* 4, 2, 177–192.
- Rocco A. Servedio. 2004. Monotone Boolean formulas can approximate monotone linear threshold functions. *Discrete Applied Mathematics* 142, 1–3, 181–187.
- Jeffrey D. Ullman. 1984. *Computational Aspects of VLSI*. Computer Science Press, Rockville, MD.
- L. G. Valiant. 1984. Short monotone formulae for the majority function. *Journal of Algorithms* 5, 3, 363–366.
- Vijay V. Vazirani. 2001. *Approximation Algorithms*. Springer.
- Gerhard J. Woeginger and Zhongliang Yu. 1992. On the equal-subset-sum problem. *Information Processing Letters* 42, 6, 299–302. DOI: [http://dx.doi.org/10.1016/0020-0190\(92\)90226-L](http://dx.doi.org/10.1016/0020-0190(92)90226-L)
- A. C. Yao. 1989. Circuits and local computation. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC'89)*. ACM, New York, NY, 186–196.

Received March 2015; revised February 2016; accepted February 2016