


Weakly Approximating Knapsack in Subquadratic Time

Lin Chen ✉ 🏠 

Zhejiang University, Hangzhou, China

Jiayi Lian ✉ 🏠 

Zhejiang University, Hangzhou, China

Yuchen Mao ✉ 🏠 

Zhejiang University, Hangzhou, China

Guochuan Zhang ✉ 🏠 

Zhejiang University, Hangzhou, China

Abstract

We consider the classic Knapsack problem. Let t and OPT be the capacity and the optimal value, respectively. If one seeks a solution with total profit at least $\text{OPT}/(1+\varepsilon)$ and total weight at most t , then Knapsack can be solved in $\tilde{O}(n + (\frac{1}{\varepsilon})^2)$ time [Chen, Lian, Mao, and Zhang '24][Mao '24]. This running time is the best possible (up to a logarithmic factor), assuming that $(\min, +)$ -convolution cannot be solved in truly subquadratic time [Künemann, Paturi, and Schneider '17][Cygan, Mucha, Węgrzycki, and Włodarczyk '19]. The same upper and lower bounds hold if one seeks a solution with total profit at least OPT and total weight at most $(1+\varepsilon)t$. Therefore, it is natural to ask the following question.

If one seeks a solution with total profit at least $\text{OPT}/(1+\varepsilon)$ and total weight at most $(1+\varepsilon)t$, can Knapsack be solved in $\tilde{O}(n + (\frac{1}{\varepsilon})^{2-\delta})$ time for some constant $\delta > 0$?

We answer this open question affirmatively by proposing an $\tilde{O}(n + (\frac{1}{\varepsilon})^{7/4})$ -time algorithm.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Knapsack, FPTAS

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.51

Category Track A: Algorithms, Complexity and Games

Related Version Full Version: <https://arxiv.org/abs/2504.15001>

Funding Yuchen Mao: National Natural Science Foundation of China [Project No. 62402436].

Guochuan Zhang: National Natural Science Foundation of China [Project No. 12271477].

1 Introduction

In the Knapsack problem, one is given a knapsack of capacity t and a set of n items. Each item i has a weight w_i and a profit p_i . The goal is to select a subset of items such that their total weight does not exceed t , and their total profit is maximized.

Knapsack is one of Karp's 21 NP-complete problems [26], which motivates the study of approximation algorithms. In particular, it is well known that Knapsack admits fully polynomial-time approximation schemes (FPTASes), which can return a feasible solution with a total profit of at least $\text{OPT}/(1+\varepsilon)$ in $\text{poly}(n, \frac{1}{\varepsilon})$ time. (We use OPT to denote the maximum total profit one can achieve using capacity t .) There has been a long line of research on designing faster FPTASes for Knapsack [21, 33, 29, 37, 10, 24, 18, 15, 34]. Very recently, Chen, Lian, Mao and Zhang [15] and Mao [34] independently obtained



© Lin Chen, Jiayi Lian, Yuchen Mao, and Guochuan Zhang;
licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis

Article No. 51; pp. 51:1–51:18



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



$\tilde{O}(n + (\frac{1}{\varepsilon})^2)$ -time FPTASes¹. On the negative side, there is no $O((n + \frac{1}{\varepsilon})^{2-\delta})$ -time FPTAS for any constant $\delta > 0$ assuming $(\min, +)$ -convolution cannot be solved in subquadratic time [17, 32]. Hence, further improvement on FPTASs for Knapsack seems hopeless unless the $(\min, +)$ -convolution conjecture is overturned. Due to the symmetry of weight and profit, the same upper and lower bounds also hold if we seek a solution of total profit at least OPT , but relax the capacity to $(1 + \varepsilon)t$. In this paper, we consider *weak approximation schemes*, which relax both the total profit and the capacity. More precisely, a weak approximation scheme seeks a subset of items with total weight at most $(1 + \varepsilon)t$ and total profit at least $\text{OPT}/(1 + \varepsilon)$.

Weak approximation algorithms are commonly studied in the context of resource augmentation (see e.g., [20, 22]). One often resorts to resource augmentation when there is a barrier to further improvement on standard (i.e., non-weak) approximation algorithms. The two most relevant examples are Subset Sum and Unbounded Knapsack. Both problems admit an $\tilde{O}(n + (\frac{1}{\varepsilon})^2)$ -time FPTAS [27, 9, 23], while an $O((n + \frac{1}{\varepsilon})^{2-\delta})$ -time FPTAS has been ruled out assuming the $(\min, +)$ -convolution conjecture [32, 17, 35, 9]. Their weak approximation, however, can be done in truly subquadratic time. For Subset Sum, Mucha, Węgrzycki and Włodarczyk [35] showed a $O(n + (\frac{1}{\varepsilon})^{5/3})$ -time weak approximation scheme, and very recently, Chen et al. [13] further improved the running time to $\tilde{O}(n + \frac{1}{\varepsilon})$ -time. For Unbounded Knapsack, Bringmann and Cassis [6] established an $\tilde{O}(n + (\frac{1}{\varepsilon})^{3/2})$ -time weak approximation scheme. Given that the standard approximation of Knapsack has been in a similar situation, it is natural to ask whether the weak approximation of Knapsack can be done in truly subquadratic time.

It is worth mentioning that weak approximation schemes for Subset Sum and Unbounded Knapsack are not straightforward extensions of the existing standard (non-weak) approximation schemes. The first truly subquadratic-time weak approximation scheme for Subset Sum [35] as well as its recent improvement [13] both rely on a novel application of additive combinatorics, and the first truly subquadratic-time weak approximation scheme for Unbounded Knapsack [6] relies on a breakthrough on bounded monotone $(\min, +)$ -convolution [11, 16]. Unfortunately, techniques developed in these prior papers are not sufficient to obtain a truly subquadratic-time weak approximation scheme for Knapsack.

Our contribution

We propose the first truly subquadratic-time weak approximation scheme for Knapsack.

► **Theorem 1.** *Knapsack admits an $\tilde{O}(n + (\frac{1}{\varepsilon})^{7/4})$ -time weak approximation scheme, which is randomized and succeeds with high probability.²*

Our algorithm can be easily generalized to Bounded Knapsack, in which, every item i can be selected up to u_i times. When $u_i = O(1)$ for all i , Bounded Knapsack is equivalent to Knapsack (up to a constant factor), and therefore can be handled by our algorithm. For the general case, there is a standard approach of reducing Bounded Knapsack to instances where $u_i = O(1)$ (see e.g., [35, Lemma 4.1][31, Lemma 2.2]). Note that such a reduction will blow up item weights and profits by a factor $O(u_i)$, but this does not make much difference in (weak) approximation schemes due to scaling.

¹ In the paper, we use an $\tilde{O}(\cdot)$ notation to hide polylogarithmic factors in n and $\frac{1}{\varepsilon}$.

² Throughout the paper, “with high probability” stands for “with probability at least $1 - (n + \frac{1}{\varepsilon})^{-\Omega(1)}$ ”.

1.1 Technical Overview

We give a brief overview of the techniques that are used by our algorithm.

Bounded Monotone (min, +)-Convolution

(min, +)-convolution (more precisely, its counterpart, (max, +)-convolution) has been used in almost all the recent progress on Knapsack [10, 24, 14, 25, 5, 15, 34]. It allows one to cut the input instance into sub-instances, deal with each sub-instance independently, and then merge the results using (min, +)-convolution. In general, it takes $O(n^2)$ time to compute the (min, +)-convolution of two length- n sequences³, and it is conjectured that there is no algorithm with time $O(n^{2-\delta})$ for any constant $\delta > 0$. Nevertheless, some special cases of (min, +)-convolution can be computed in subquadratic time (See [1, 2, 36, 11, 16] for example), and among them is bounded monotone (min, +)-convolution where the two sequences are monotone and their entries are integers bounded by $O(n)$. In a celebrated work, Chan and Lewenstein [11] showed that bounded monotone (min, +)-convolution can be computed in $O(n^{1.859})$ time. This was later improved to $\tilde{O}(n^{1.5})$ by Chi, Duan, Xie, and Zhang [16], and further generalized to rectangular monotone (min, +)-convolution by Bringmann, Dürr and Polak [8].

Bounded monotone (min, +)-convolution is closely related to Unbounded Knapsack. Indeed, Bringmann and Cassis presented a $\tilde{O}(n + (\frac{1}{\epsilon})^{3/2})$ -time weak approximation scheme [6] that builds upon bounded monotone (min, +)-convolution, and they also showed an equivalence result between the two problems in the sense that a truly subquadratic-time weak approximation scheme for Unbounded Knapsack implies a truly subquadratic algorithm for bounded monotone (min, +)-convolution, and vice versa.

For our problem, we follow the general framework mentioned above: we cut the input instance into sub-instances (called the reduced problem), deal with each sub-instance independently, and then merge the results using bounded monotone (min, +)-convolution. Note that following this framework, for each sub-instance, we need to compute the near-optimal objective value for all (approximate) knapsack capacities.

One might anticipate that a subquadratic algorithm for bounded monotone (min, +)-convolution leads to a subquadratic weak approximation scheme for Knapsack straightforwardly, however, this is not the case. Bounded monotone (min, +)-convolution is only used to merge results; the main challenge lies in dealing with each group, where we need the following two techniques.

Knapsack with Items of Similar Efficiencies

When all items have the same efficiency, Knapsack becomes Subset Sum, and it can be weakly approximated in $\tilde{O}(n + \frac{1}{\epsilon})$ time [13]. We extend the result to the case where all the items have their efficiencies (that is, profit-to-weight ratio) in the range $[\rho, \rho + \Delta]$ for constant ρ , and show that in such a case, Knapsack can be weakly approximated in $\tilde{O}(n + \frac{1}{\epsilon} + (\frac{1}{\epsilon})^2 \Delta)$ time. This result may be of independent interest. The following is a rough idea. By scaling, we can assume that there are only $O(\frac{\Delta}{\epsilon})$ distinct item efficiencies. Then we divide the items into $O(\frac{\Delta}{\epsilon})$ groups so that the items in the same group have the same efficiency, and therefore,

³ A slightly faster $O(n^2/2^{\Omega(\sqrt{\log n})})$ -time algorithm is known by Bremner et.al.'s reduction to (min, +)-matrix multiplication [3] and William's algorithm for the latter problem [38] (which was derandomized by Chan and Williams [12]).

each group can be treated as a Subset Sum problem. For each group, we solve it by invoking Chen et al.'s $\tilde{O}(n + \frac{1}{\epsilon})$ -time weak approximation scheme for Subset Sum [13]. After solving all groups, we merge their results using 2D-FFT and divide-and-conquer, which takes $\tilde{O}((\frac{1}{\epsilon})^2 \Delta)$ time.

We note that the general idea of using Subset Sum to solve Knapsack when items have similar efficiencies was used by Jin [24], and the idea of merging Subset Sum results using 2D-FFT and divide-and-conquer was used by Deng, Jin and Mao [18].

Proximity Bounds

Proximity is another tool that is crucial to the recent progress of Knapsack [36, 14, 25, 5, 15]. Proximity bounds capture the intuition that the optimal solution should contain a lot of high-efficiency items and only a few low-efficiency items. (The efficiency of an item is defined to be its profit-to-weight ratio.) It was pioneered by Eisenbrand and Weismantel [19]. Their proximity bound works for a more general problem of integer programming and is based on Steinitz's lemma. The proof can be greatly simplified if one only considers Knapsack [36]. This bound was subsequently improved using tools from additive combinatorics [14, 5, 25]. These bounds depend on the maximum item weight w_{\max} or the maximum item profits p_{\max} , and leads to $\tilde{O}(n + w_{\max}^2)$ -time and $\tilde{O}(n + p_{\max}^2)$ -time algorithm for Knapsack [25, 5]. (We remark that all the proximity bounds of this type assume that the item weights and profits are integers.) These results, however, seem not helpful in breaking the quadratic barrier: even if we can scale the w_{\max} and p_{\max} to integers bounded by $O(\frac{1}{\epsilon})$, it still takes $\tilde{O}(n + (\frac{1}{\epsilon})^2)$ time to solve the problem.

We shall use another type of proximity bounds, which depends on the item efficiencies. Such a proximity bound was used to design FPTASes for Knapsack [24, 18] before. This proximity states that the optimal solution selects almost all the high-efficiency items, and exchanges only a few of them for the low-efficiency items. Moreover, the number of exchanged items is inversely proportional to the efficiency gap between high- and low-efficiency items. Fewer exchanged items allow a larger additive error per exchanged item, and therefore, a faster approximation.

Our Main Technical Contribution

While each of the above-mentioned techniques appeared in previous papers before, it seems that none of them alone can break the quadratic barrier for weakly approximating Knapsack. Our main contribution is to modify and then combine these techniques to obtain a truly subquadratic-time algorithm.

Note that if the item efficiencies differ significantly, then the efficiency-based proximity bound becomes sharp; If item efficiencies stay similar, then the Subset Sum based approach becomes efficient. To combine both ideas, the main difficulty is that we need to approximate the optimal objective value for all (approximate) capacities, and the efficiency-based proximity bound may be sharp for certain capacities, and weak for other capacities. Towards this, the main idea is to use different rounding precision for different items. We roughly explain it below. For simplicity, when we say we approximately solve a group of items, we mean to approximately compute the optimal objective value for all knapsack capacities when taking items from this group. We partition all items into multiple groups of fixed size. If the items in a group have similar efficiencies, then we approximately solve this group with high precision using the Subset Sum based approach; Otherwise, we approximately solve this group with low precision. Low precision means a larger error per item, but if proximity guarantees that

only a few items from this group will contribute to the optimal solution (or only a few items will not contribute), then the total error may still be bounded. So the challenge is to prove that a sharp proximity holds for such a group regardless of the knapsack capacity.

We remark that a simpler implementation of the above idea leads to an algorithm with running time $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$. By grouping the items in a more sophisticated way, the running time can be improved to $\tilde{O}(n + (\frac{1}{\varepsilon})^{7/4})$.

1.2 Further Related Work

FPTASes with Running Time in Multiplicative Form

So far, we have discussed FPTASes for Knapsack whose running times are $n + f(\frac{1}{\varepsilon})$. There is also a line of research focusing on FPTASes with a running time of the form $n^{O(1)} \cdot f(\frac{1}{\varepsilon})$ (see, e.g., [33, 28, 10]). The best known running time of this form is $\tilde{O}(\frac{n^{3/2}}{\varepsilon})$ due to Chan [10]. It remains open whether there exists an FPTAS of running time $O(\frac{n}{\varepsilon})$. It is not even clear whether Knapsack admits a weak approximation scheme in $O(\frac{n}{\varepsilon})$ time.

Pseudo-polynomial Time Algorithm

Recent years have witnessed important progress in the study of pseudo-polynomial time algorithms for Knapsack and Unbounded Knapsack [6, 7, 14, 5, 25]. In particular, the best-known algorithm for Knapsack that is parametrized by the largest weight w_{\max} (or the largest profit p_{\max}) runs in $\tilde{O}(n + w_{\max}^2)$ -time (or $\tilde{O}(n + p_{\max}^2)$ -time) [5, 25], and is the best possible assuming the (min, +)-convolution conjecture [17, 32]. It is, however, a major open problem whether the quadratic barrier can be broken by combining both parameters, that is, whether an $\tilde{O}(n + (w_{\max} + p_{\max})^{2-\delta})$ -time algorithm exists for Knapsack. Notably, for Unbounded Knapsack, one can find an algorithm that runs in $\tilde{O}(n + (w_{\max} + p_{\max})^{3/2})$ time [6].

1.3 Paper Outline

In Section 2, we introduce all the necessary notation and terminology. In Section 3, we reduce the problem so that it suffices to consider instances with nice properties. In Section 4, we present all the ingredients that are needed by our algorithm. In Section 5, we present an $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$ -time algorithm to illustrate our main idea. In Section 6, we show how to improve the running time to $\tilde{O}(n + (\frac{1}{\varepsilon})^{7/4})$. All the missing proofs can be found in the full version of the paper.

2 Preliminaries

We assume that $\varepsilon \leq 1$, since otherwise we can approximate using a linear-time 2-approximation algorithm for Knapsack [30, Section 2.5]. We denote the set of real numbers between a and b as $[a, b]$. All logarithms are based 2.

To ease the presentation of our algorithm, we interpret Knapsack from the perspective of tuples. An item can be regarded as a tuple (w, p) , where w and p represent the weight and profit, respectively. Throughout the paper, we do not distinguish an item and a tuple, and always refer to the first entry as weight and the second entry as profit. The efficiency of a tuple (w, p) is defined to be its profit-to-weight ratio $\frac{p}{w}$. The efficiency of $(0, 0)$ is defined to be 0.

Let I be a multiset of tuples. We use $w(I)$ to denote the total weight of tuples in I . That is, $w(I) = \sum_{(w,p) \in I} w$. Similarly, we define $p(I) = \sum_{(w,p) \in I} p$. Define the set of all (2-dimensional) subset sums of I as follows:

$$\mathcal{S}(I) = \{(w(I'), p(I')) : I' \subseteq I\}.$$

Every tuple $(w, p) \in \mathcal{S}(I)$ corresponds to a Knapsack solution with total weight w and total profit p . We say a tuple (w, p) dominates another tuple (w', p') if $w \leq w'$, $p \geq p'$, and at least one of the two inequalities holds strictly. Let $\mathcal{S}^+(I)$ be the set obtained from $\mathcal{S}(I)$ by removing all dominated tuples. Every tuple in $\mathcal{S}^+(I)$ corresponds to a Pareto optimal solution for Knapsack, and vice versa. Let t be the given capacity. Knapsack can be formulated as follows:

$$\max\{p : (w, p) \in \mathcal{S}^+(I) \text{ and } w \leq t\}.$$

We consider the more general problem of (weakly) approximating $\mathcal{S}^+(I)$.

► **Definition 2.** Let S and S' be two sets of tuples.

- (i) S' approximates S with factor $1 + \varepsilon$ if for any $(w, p) \in S$, there exists $(w', p') \in S'$ such that $w' \leq (1 + \varepsilon)w$ and $p' \geq p/(1 + \varepsilon)$;
- (ii) S' approximates S with additive error (δ_w, δ_p) if for any $(w, p) \in S$, there exists $(w', p') \in S'$ such that $w' \leq w + \delta_w$ and $p' \geq p - \delta_p$;
- (iii) S dominates S' if for any $(w', p') \in S'$, there exists $(w, p) \in S$ such that $w \leq w'$ and $p \geq p'$.

When S' approximates a set S that contains only one tuple (w, p) , we also say that S' approximates the tuple (w, p) .

Let (I, t) be a Knapsack instance. Let OPT be the optimal value. Our goal is to compute a set S' such that

- S' approximates $\mathcal{S}^+(I) \cap ([0, t] \times [0, \text{OPT}])$ with additive error $(\varepsilon t, \varepsilon \cdot \text{OPT})$ and
- S' is dominated by $\mathcal{S}(I)$.

The former ensures that S' provides a good (weak) approximation of the optimal solution, and the latter ensures that every tuple in S' is not better than a true solution. Our definition is essentially the same as that by Bringmann and Cassis [6], albeit their definition is for sequences.

Let $I_1 \cup I_2$ be a partition of I . It is easy to observe the following.

- $\mathcal{S}(I) = \mathcal{S}(I_1) + \mathcal{S}(I_2)$, where the sumset $X + Y$ is defined by the following formula.

$$X + Y = \{(w_1 + w_2, p_1 + p_2) : (w_1, p_1) \in X \text{ and } (w_2, p_2) \in Y\};$$

- $\mathcal{S}^+(I) = \mathcal{S}^+(I_1) \oplus \mathcal{S}^+(I_2)$, where \oplus represents the $(\max, +)$ -convolution defined by the following formula.

$$X \oplus Y = \{(w, p) \in X + Y : (w, p) \text{ is not dominated in } X + Y\}.$$

Using Chi et al.'s $\tilde{O}(n^{3/2})$ -time algorithm for bounded monotone $(\min, +)$ -convolution [16] and Bringmann and Cassis's algorithm for approximating the $(\max, +)$ -convolution [6, Lemma 28 in the full version], we can approximate the $(\max, +)$ -convolution of two sets in $\tilde{O}((\frac{1}{\varepsilon})^{3/2})$ time. Using the divide-and-conquer approach from [10], we can generalize it to multiple sets. The proof of the following lemma can be found in the appendix of the full version.

► **Lemma 3.** *Let $S_1, S_2, \dots, S_m \subseteq (\{0\} \cup [A, B]) \times (\{0\} \cup [C, D])$ be sets of tuples with total size ℓ . In $\tilde{O}(\ell + (\frac{1}{\varepsilon})^{3/2}m)$ time⁴ and with high probability, we can compute a set \tilde{S} of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $S_1 \oplus S_2 \oplus \dots \oplus S_m$ with factor $1 + O(\varepsilon)$. Moreover, \tilde{S} is dominated by $S_1 + S_2 + \dots + S_m$.*

With Lemma 3, we can approximate $\mathcal{S}^+(I) \cap ([0, t] \times [0, OPT])$ as follows. Partition I into several groups, approximate each group independently, and merge their results using Lemma 3. As long as there are not too many groups, the merging time will be subquadratic. The following lemmas explain how the multiplicative factor and additive error accumulate during the computation. Their proofs can be found in the full version.

► **Lemma 4.** *Let S , S_1 , and S_2 be sets of tuples.*

- (i) *If S_1 approximates S with factor $1 + \varepsilon_1$ and S_2 approximates S_1 with factor $1 + \varepsilon_2$, then S_2 approximates S with factor $(1 + \varepsilon_1)(1 + \varepsilon_2)$.*
- (ii) *If S_1 approximates S with additive error $(\delta_{w1}, \delta_{p1})$ and S_2 approximates S_1 with factor $(\delta_{w2}, \delta_{p2})$, then S_2 approximates S with additive error $(\delta_{w1} + \delta_{w2}, \delta_{p1} + \delta_{p2})$.*

► **Lemma 5.** *Let S_1, S_2, S'_1, S'_2 be sets of tuples.*

- (i) *If S'_1 and S'_2 approximates S_1 and S_2 respectively with factor $1 + \varepsilon$, then $S'_1 \oplus S'_2$ and $S'_1 + S'_2$ approximate $S_1 \oplus S_2$ and $S_1 + S_2$ respectively with factor $1 + \varepsilon$.*
- (ii) *If S'_1 and S'_2 approximates S_1 and S_2 with additive error $(\delta_{w1}, \delta_{p1})$ and $(\delta_{w2}, \delta_{p2})$, respectively, then $S'_1 \oplus S'_2$ and $S'_1 + S'_2$ approximate $S_1 \oplus S_2$ and $S_1 + S_2$ respectively with additive error $(\delta_{w1} + \delta_{w2}, \delta_{p1} + \delta_{p2})$.*

In the rest of the paper, we may switch between the multiplicative factor and the additive error, although our goal is to bound the latter. It is helpful to keep the following observation in mind.

► **Observation 6.** *If S' approximates S with factor $1 + \varepsilon$, then S' approximates each tuple $(w, p) \in S$ with additive error $(\varepsilon w, \varepsilon p)$.*

3 Reducing the Problem

With Lemma 3, we can reduce the problem of weakly approximating Knapsack to the following problem.

The reduced problem $RP(\alpha)$: given a set I of n items from $[1, 2] \times [1, 2]$ and a real number $\alpha \in [\frac{1}{2}, \frac{1}{4\varepsilon}]$, compute a set S of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $\mathcal{S}^+(I) \cap ([0, \frac{1}{\alpha\varepsilon}] \times [0, \frac{1}{\alpha\varepsilon}])$ with additive error $(O(\frac{1}{\alpha}), O(\frac{1}{\alpha}))$. Moreover, S should be dominated by $\mathcal{S}(I)$.

► **Lemma 7.** *An $\tilde{O}(n + T(\frac{1}{\varepsilon}))$ -time algorithm for the reduced problem $RP(\alpha)$ implies an $\tilde{O}(n + (\frac{1}{\varepsilon})^{3/2} + T(\frac{1}{\varepsilon}))$ -time weak approximation scheme for Knapsack.*

The approach to reducing the problem is similar to that in [15]. The details of the proof can be found in the full version.

Throughout the rest of the paper, we write $\mathcal{S}^+(I) \cap ([0, \frac{1}{\alpha\varepsilon}] \times [0, \frac{1}{\alpha\varepsilon}])$ as $\mathcal{S}^+(I; \frac{1}{\alpha\varepsilon})$ for short, and when the weight and profit both have an additive error of δ , we simply say that additive error is δ .

⁴ Here, the $\tilde{O}(\cdot)$ notation hides polylogarithmic factor not only in n and $\frac{1}{\varepsilon}$, but also in other parameters such as $\frac{B}{A}$ and $\frac{D}{C}$. These parameters will eventually be set to be polynomial in n and $\frac{1}{\varepsilon}$.

We also remark that when presenting algorithms for the reduced problem $\text{RP}(\alpha)$, we will omit the requirement that the resulting set S should be dominated by $\mathcal{S}(I)$ since it is guaranteed by the algorithms in a quite straightforward way. More specifically, in those algorithms (including Lemma 3), rounding is the only place that may incur approximation errors, and whenever we round a tuple (w, p) , we always round w up and round p down. As a consequence, a set after rounding is always dominated by the original set.

4 Algorithmic Ingredients

We present a few results that will be used by our algorithm as subroutines.

4.1 An Algorithm for Small Solution Size

Consider the reduced problem $\text{RP}(\alpha)$. Since items are from $[1, 2] \times [1, 2]$, it is easy to see that $|I'| \leq \frac{1}{\alpha\varepsilon}$ for any $I' \subseteq I$ with $(w(I'), p(I')) \in \mathcal{S}^+(I; \frac{1}{\alpha\varepsilon})$. When α is large, $|I'|$ is small, and such a case can be tackled using Bringmann's two-layer color-coding technique [4]. Roughly speaking, we can divide I into $k = \tilde{O}(\frac{1}{\alpha\varepsilon})$ subsets I_1, \dots, I_k so that with high probability, $|I_j \cap I'| \leq 1$ for every j . Then it suffices to (approximately) compute $I_1 \oplus \dots \oplus I_k$, and the resulting set will provide a good approximation for the tuple $(w(I'), p(I'))$. Since this technique has become quite standard nowadays, we defer the proof to the full version.

► **Lemma 8.** *Let I be a set of items from $[1, 2] \times [1, 2]$. In $\tilde{O}(n + (\frac{1}{\varepsilon})^{5/2} \frac{1}{\alpha})$ time and with high probability, we can compute a set S of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $\mathcal{S}^+(I; \frac{1}{\alpha\varepsilon})$ with factor $1 + O(\varepsilon)$.*

4.2 An Efficiency-Based Proximity Bound

A proximity bound basically states that the optimal solution of Knapsack must contain lots of high-efficiency items and very few low-efficiency items. Let $I = \{(w_1, p_1), \dots, (w_n, p_n)\}$ be a set of items. Assume that the items are labeled in a non-increasing order of efficiencies. That is, $\frac{p_1}{w_1} \geq \dots \geq \frac{p_n}{w_n}$. (The labeling may not be unique. We fix an arbitrary one in that case.) Given a capacity w , let b be the minimum integer such that $w_1 + \dots + w_b > w$. We say that the item (w_b, p_b) is the breaking item with respect to w . (When no such b exists, we simply let $b = n + 1$ and define the breaking item to be $(+\infty, 0)$.)

The following proximity bound is essentially the same as the one used in [24, 18]. For completeness, we provide a proof in the full version.

► **Lemma 9.** *Let I be a set of items from $[1, 2] \times [1, 2]$. Let $(w^*, p^*) \in \mathcal{S}^+(I)$. Let $I^* \subseteq I$ be a subset of items with $w(I^*) = w^*$ and $p(I^*) = p^*$. Let ρ_b be the efficiency of the breaking item with respect to w^* . For any $I' \subseteq I$,*

(i) *if every item in I' has efficiency at most $\rho_b - \Delta$ for some $\Delta > 0$, then*

$$w(I' \cap I^*) \leq \frac{2}{\Delta};$$

(ii) *if every item in I' has efficiency at least $\rho_b + \Delta$ for some $\Delta > 0$, then*

$$w(I' \setminus I^*) \leq \frac{2}{\Delta}.$$

4.3 An Algorithm for Items with Similar Efficiencies

When all the items have the same efficiency, Knapsack becomes Subset Sum, and there is an $\tilde{O}(n + \frac{1}{\varepsilon})$ -time weak approximation scheme for Subset Sum [13]. The following lemma extends this result to the case where the item efficiencies differ by only a small amount.

► **Lemma 10.** *Let I be a set of n items from $[1, 2] \times [0, \infty]$ whose efficiencies are in the range $[\rho, \rho + \Delta]$. In $\tilde{O}(n + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$ time and with high probability, we can compute a set of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $\mathcal{S}^+(I)$ with factor $1 + O(\varepsilon)$.*

We first consider the $(\max, +)$ -convolution of two sets of tuples.

► **Lemma 11.** *Let S_1, S_2 be two sets of tuples from $(\{0\} \cup [A, B]) \times [0, +\infty]$ with total size ℓ . Suppose that the efficiencies of the tuples in $S_1 \cup S_2$ are in the range $\{0\} \cup [\rho, \rho + \Delta]$. In $\tilde{O}(\ell + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$ time⁵, we can compute a set S of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $S_1 \oplus S_2$ with factor $1 + O(\varepsilon)$. Moreover, S is dominated by $S_1 + S_2$.*

Proof. It suffices to consider the case where $\rho = 1$. If $\rho \neq 1$, we can replace every tuple (w, p) with $(w, \frac{p}{\rho})$, and replace Δ with $\frac{\Delta}{\rho}$.

Let $a \in [A, B]$. We first describe how to approximate $(S_1 \oplus S_2) \cap ([0, a] \times [0, +\infty])$ with additive error $O(\varepsilon a)$.

Let S'_1 and S'_2 be the sets obtained from S_1 and S_2 , respectively, by replacing every tuple (w, p) with $(w, p - w)$. The range of item efficiencies of S'_1 and S'_2 becomes $[0, \Delta]$. We assume that every tuple $(w, p) \in S'_1 \cup S'_2$ has $w \leq a$ since the tuples with $w > a$ can be safely ignored. For every tuple (w, p) , we round w up to and round p down to the nearest multiple of εa . Denote the resulting sets by S''_1 and S''_2 . By scaling, we can assume that every tuple in $S''_1 \cup S''_2$ is a tuple of integers from $[0, \frac{1}{\varepsilon}] \times [0, \frac{\Delta}{\varepsilon}]$. Then we can compute $S = S''_1 + S''_2$ using 2D-FFT in $\tilde{O}(\frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \Delta)$ time. It is easy to see that S is of size $\tilde{O}((\frac{1}{\varepsilon})^2 \Delta)$, and is dominated by $S'_1 + S'_2$ (since we round w up and round p down). Moreover, for each $(w, p) \in S'_1 + S'_2$, there is a tuple $(w', p') \in S$ such that $w \leq w' \leq w + 2\varepsilon a$ and $p - 2\varepsilon a \leq p' \leq p$, which implies that $p + w - 4\varepsilon a \leq p' + w' - 2\varepsilon a \leq p + w$. Now consider the set S' obtained from S by replacing each tuple (w, p) with $(w, p + w - 2\varepsilon a)$. One can verify that S' is dominated by $S_1 + S_2$ and approximates $S_1 + S_2$ with additive error $O(\varepsilon a)$. Since $S_1 \oplus S_2 \subseteq S_1 + S_2$, we have that S' also approximates $S_1 \oplus S_2$ with additive error $O(\varepsilon a)$. Then we remove all the dominated tuples from S' . This can be done in $O(|S'| \log |S'|) = \tilde{O}((\frac{1}{\varepsilon})^2 \Delta)$ time. After this, the size of S' becomes $O(\frac{1}{\varepsilon})$ because for each integer $w \in [0, \frac{1}{\varepsilon}]$, there can be at most one tuple with weight w in S' .

To finish the proof, we repeat the above procedure for every $a \in [A, B]$ that is a power of 2, and take the union of the $\log \frac{B}{A}$ resulting sets. Note that for a tuple $(w, p) \in S_1 + S_2$ with $w \in [a/2, a]$, we have that $p \geq \rho w = w \geq a/2$. Therefore, the additive error $O(\varepsilon a)$ implies an approximation factor of $1 + O(\varepsilon)$. The running time increases by a factor of $\log \frac{B}{A}$. ◀

Now we are ready to prove Lemma 10.

Proof of Lemma 10. If $\Delta/\rho < \varepsilon$, we can round down item profits so that the efficiencies of all items are ρ . This increases the approximation factor by at most $1 + \varepsilon$. Now the problem degenerates to Subset Sum, and therefore, we can compute a set S that approximates $\mathcal{S}^+(I)$ in $\tilde{O}(n + \frac{1}{\varepsilon})$ time by the weak approximation scheme⁶.

⁵ Here, the $\tilde{O}(\cdot)$ notation hides polylogarithmic factor not only in n and $\frac{1}{\varepsilon}$, but also in $\frac{B}{A}$. These parameters will eventually be set to be polynomial in n and $\frac{1}{\varepsilon}$.

⁶ The weak approximation scheme in [13] actually returns a set \tilde{S} that approximates $\mathcal{S}^+(I_j)$ with additive error εt . This can be easily modified to return a set that approximates $\mathcal{S}^+(I_j)$ with factor $1 + \varepsilon$ by repeating it for $t = 1, 2, 4, 8, \dots$, and taking the union of the resulting sets.

Now we assume $\Delta/\rho \geq \varepsilon$. By rounding down item profits, we assume that the efficiencies of items are powers of $1 + \varepsilon$. This increases the approximation factor by at most $1 + \varepsilon$. So there are at most $m = \log_{1+\varepsilon} \frac{\Delta}{\rho} = O(\frac{\Delta}{\varepsilon\rho})$ distinct efficiencies. We divide the items into m groups I_1, \dots, I_m by their efficiencies. For each group I_j , since the items have the same efficiency, and again, we can compute a set S_j that approximates $\mathcal{S}^+(I_j)$ in $\tilde{O}(|I_j| + \frac{1}{\varepsilon})$ time by the weak approximation scheme for Subset Sum [13]. The total time cost for computing all sets S_j is $\tilde{O}(n + \frac{m}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho}) = \tilde{O}(n + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$.

Then we shall merge S_1, \dots, S_m using a divide-and-conquer approach. The tuples in S_j have the same efficiency (except that the tuple $(0,0)$ has efficiency 0), which is the efficiency of the items in I_j . Let ρ_j be this efficiency. Without loss of generality, assume that $\rho_1 < \dots < \rho_m$. We recursively (and approximately) compute $S_1 \oplus \dots \oplus S_{m/2}$ and $S_{m/2+1} \oplus \dots \oplus S_m$, and then merge the resulting two sets via Lemma 11. The recursion tree has $\log m$ levels. One can verify that the total merging time in each level is $\tilde{O}(\frac{m}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\rho_m - \rho_1}{\rho_1})$, which is bounded by $\tilde{O}((\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$ since $\rho \leq \rho_j \leq \rho + \Delta$. The approximation factor is $(1 + \varepsilon)^{\log m} = 1 + O(\varepsilon \log m)$. We can adjust ε by a factor of $\log m$, which increases the running time by a polylogarithmic factor. \blacktriangleleft

The Lemma 10 immediately implies the following.

► **Corollary 12.** *Let I be a set of n items from $[1, 2] \times [0, +\infty]$ whose efficiencies are in the range $[\rho, \rho + \Delta]$. In $\tilde{O}(n + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$ time and with high probability, we can compute a set of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates each tuple $(w, p) \in \mathcal{S}^+(I)$ with additive error $(\varepsilon w, \varepsilon p)$.*

In the above corollary, the additive error is proportional to the value of the entry of the tuple: a large entry means a large additive error. We can also achieve the contrary and approximate in a way so that a large entry means a small additive error. Basically, this can be done by (approximately) solving a problem symmetric to Knapsack: instead of determining which items should be selected, we determine which items should not be selected.

► **Corollary 13.** *Let I be a set of n items from $[1, 2] \times [0, +\infty]$ whose efficiencies are in the range $[\rho, \rho + \Delta]$. In $\tilde{O}(n + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$ time and with high probability, we can compute a set of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $\mathcal{S}^+(I)$ with additive error $(\varepsilon(w(I) - w), \varepsilon(p(I) - p))$.*

Proof. Define

$$\mathcal{S}^-(I) = \{(w, p) \in \mathcal{S}(I) : (w, p) \text{ does not dominate any other tuple in } \mathcal{S}(I)\}.$$

From the perspective of Knapsack, every tuple in $\mathcal{S}^+(I)$ corresponds to an optimal solution of Knapsack, while every tuple (w, p) in $\mathcal{S}^-(I)$ corresponds to an optimal solution of the symmetric problem that seeks for the minimum total profit that can be achieved using capacity at least w . It is easy to observe that

$$\mathcal{S}^+(I) = \{(w(I) - w, p(I) - p) : (w, p) \in \mathcal{S}^-(I)\}.$$

Therefore, to approximate $\mathcal{S}^+(I)$, it suffices to approximate $\mathcal{S}^-(I)$.

Geometrically, $\mathcal{S}^+(I)$ represents the upper envelope of $\mathcal{S}(I)$, while $\mathcal{S}^-(I)$ represents the lower envelope. But we can make $\mathcal{S}^-(I)$ be the upper envelope by exchanging the x -coordinate and y -coordinate. Then we can approximate $\mathcal{S}^-(I)$ using Corollary 12.

More precisely, let $I' = \{(p, w) : (w, p) \in I\}$. The efficiencies of the items in I' are in the range $[\rho', \rho' + \Delta']$, where $\rho' = \frac{1}{\rho + \Delta}$ and $\rho' + \Delta' = \frac{1}{\rho}$. One can verify that $\frac{\Delta'}{\rho'} = \frac{\Delta}{\rho}$. Then we can compute a set S of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates each tuple $(p, w) \in \mathcal{S}^+(I')$ with additive error $(\varepsilon p, \varepsilon w)$ via Corollary 12 in $\tilde{O}(n + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{\Delta}{\rho})$ time. Let

$$S' = \{(w(I) - w, p(I) - p) : (p, w) \in S\}.$$

One can verify that S' approximates each tuple $(w, p) \in \mathcal{S}^+(I)$ with additive error $(\varepsilon(w(I) - w), \varepsilon(p(I) - p))$. \blacktriangleleft

5 An Algorithm with Exponent 11/6

To illustrate our main idea, we first present a simpler algorithm that runs in $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$ time. When $\alpha \geq (\frac{1}{\varepsilon})^{2/3}$, the algorithm in Lemma 8 already runs in $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$ time. It remains to consider the case where $\frac{1}{2} \leq \alpha \leq (\frac{1}{\varepsilon})^{2/3}$. In the rest of this section, we assume that $I = \{(w_1, p_1), \dots, (w_n, p_n)\}$ and that $\frac{p_1}{w_1} \geq \dots \geq \frac{p_n}{w_n}$.

Let τ be a parameter that will be specified later. We first partition I into groups as follows. For the items $(w_1, p_1), \dots, (w_n, p_n)$ in I , we bundle every τ items as a group until we obtain $\lceil \frac{1}{\alpha\varepsilon\tau} \rceil + 1$ groups. For the remaining items, say $\{(w_i, p_i), \dots, (w_n, p_n)\}$, we further divide them into two groups $\{(w_i, p_i), \dots, (w_{i'}, p_{i'})\}$ and $\{(w_{i'+1}, p_{i'+1}), \dots, (w_n, p_n)\}$, where i' is the maximum index such that $\frac{p_i}{w_i} - \frac{p_{i'}}{w_{i'}} \leq \frac{1}{\tau}$. At the end, we obtain $m = \lceil \frac{1}{\alpha\varepsilon\tau} \rceil + 3$ groups I_1, I_2, \dots, I_m . For simplicity, we assume that none of these groups is empty. This assumption is without loss of generality, since empty groups only make things easier.

For each group I_j , we define Δ_j to be the maximum difference between the efficiencies of the items in I_j . That is,

$$\Delta_j = \max \left\{ \frac{p}{w} - \frac{p'}{w'} : (w, p), (w', p') \in I_j \right\}.$$

By the way we construct the groups, we have that $\sum_j \Delta_j \leq \frac{3}{2}$.

We say a group I_j is good if $\Delta_j \leq \frac{1}{\tau}$, and bad otherwise.

Note that the efficiency of any item in I is within the range $[\frac{1}{2}, 2]$. For each good group I_j , its item efficiencies are in the range $[\rho, \rho + \frac{1}{\tau}]$ for some constant $\rho \in [\frac{1}{2}, 2]$, so we can compute a set S_j of $\tilde{O}(\frac{1}{\varepsilon})$ that approximates each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $(\varepsilon w, \varepsilon p)$ via Corollary 12 in $O(|I_j| + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{1}{\tau})$ time. Since the items are from $[1, 2] \times [1, 2]$, for any tuple in $\mathcal{S}^+(I_j)$, its weight and profit are of the same magnitude, so are the corresponding additive errors. We can say that S_j approximates each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $O(\varepsilon w)$.

For bad groups, we shall approximate them less accurately: we shall use $\frac{1}{\alpha\tau}$ instead of ε as the accuracy parameter. (Later we will show that we can still obtain a good bound on the total additive error using the proximity bound in Lemma 9.) We compute a set S'_j of size $\tilde{O}(\alpha\tau)$ that approximates each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $(\frac{w}{\alpha\tau}, \frac{p}{\alpha\tau})$ via Corollary 12 in $O(|I_j| + \alpha\tau + \alpha^2\tau^2\Delta_j)$ time. We also compute a set S''_j of size $\tilde{O}(\alpha\tau)$ that approximates each $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $(\frac{w(I)-w}{\alpha\tau}, \frac{p(I)-p}{\alpha\tau})$ via Corollary 13 in $O(|I_j| + \alpha\tau + \alpha^2\tau^2\Delta_j)$ time. Again, since the weight and the profit are of the same magnitude, we can say that S'_j approximates each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $O(\frac{1}{\alpha\tau} \cdot w)$, and that S''_j approximates each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $O(\frac{1}{\alpha\tau} \cdot (w(I) - w))$. Now consider the set $S_j = S'_j \cup S''_j$. One can verify that S_j approximates each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $\tilde{O}(\frac{1}{\alpha\tau} \cdot \min\{w, (w(I) - w)\})$.

The last step is to compute a set S of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $S_1 \oplus \dots \oplus S_m$ with factor $1 + O(\varepsilon)$ via Lemma 3. Note that each set S_j is of size $\tilde{O}(\frac{1}{\varepsilon})$ or $\tilde{O}(\alpha\tau)$. This step takes time $\tilde{O}(\frac{m}{\varepsilon} + m\alpha\tau + (\frac{1}{\varepsilon})^{3/2}m)$.

Bounding Total Time Cost

Partitioning the item into I_1, \dots, I_m can be done by sorting and scanning, and therefore, it takes only $O(n \log n + m)$ time. Each good group costs $O(|I_j| + \frac{1}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{1}{\tau})$ time. Since there are at most m good groups, the good groups cost total time

$$\tilde{O}\left(n + \frac{m}{\varepsilon} + (\frac{1}{\varepsilon})^2 \frac{m}{\tau}\right).$$

51:12 Weakly Approximating Knapsack in Subquadratic Time

Each bad group I_j costs $O(|I_j| + \alpha\tau + \alpha^2\tau^2\Delta_j)$ time, so the total time cost for bad groups is

$$\tilde{O}(n + m\alpha\tau + \alpha^2\tau^2 \sum_j \Delta_j) \leq \tilde{O}(n + m\alpha\tau + \alpha^2\tau^2).$$

The inequality is due to that $\sum_j \Delta_j \leq \frac{3}{2}$. Merging all the sets S_j via Lemma 3 costs $\tilde{O}(\frac{m}{\varepsilon} + m\alpha\tau + (\frac{1}{\varepsilon})^{3/2}m)$ time. Taking the sum of all these time costs, we have that the running time of the algorithm is

$$\tilde{O}\left(n + \frac{m}{\varepsilon} + m\alpha\tau + \left(\frac{1}{\varepsilon}\right)^2 \frac{m}{\tau} + \alpha^2\tau^2 + \left(\frac{1}{\varepsilon}\right)^{3/2}m\right).$$

Recall that $m = \lceil \frac{1}{\alpha\varepsilon\tau} \rceil$ and that $\frac{1}{2} \leq \alpha \leq (\frac{1}{\varepsilon})^{2/3}$. Set $\tau = \lceil (\frac{1}{\varepsilon})^{\frac{11}{12}} \frac{1}{\alpha} \rceil$. One can verify that the running time is $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$.

Bounding Additive Error

To prove the correctness of the algorithm, it suffices to show that the set S returned by the algorithm approximates every tuple in $\mathcal{S}^+(I; \frac{1}{\alpha\varepsilon})$ with additive error $O(\frac{1}{\alpha})$.

Let (w^*, p^*) be an arbitrary tuple in $\mathcal{S}^+(I; \frac{1}{\alpha\varepsilon})$. Let (w_b, p_b) be the breaking item with respect to w^* . Let $\rho_b = \frac{p_b}{w_b}$ be the efficiency of the breaking item. We first show that the proximity bound can be applied to all but four bad groups.

► **Lemma 14.** *At most four bad groups may contain an item whose efficiency is within the range $[\rho_b - \frac{1}{\tau}, \rho_b + \frac{1}{\tau}]$. Moreover, each of these groups is of size τ .*

Proof. We first show that there are at most four bad groups each containing an item whose efficiency is within the range $[\rho_b - \frac{1}{\tau}, \rho_b + \frac{1}{\tau}]$. Suppose, for the sake of contradiction, that there are five bad groups each containing an item whose efficiency is in the range $[\rho_b - \frac{1}{\tau}, \rho_b + \frac{1}{\tau}]$. Let i_1, i_2, i_3, i_4, i_5 be these five items. Without loss of generality, assume that $\frac{p_{i_1}}{w_{i_1}} \geq \dots \geq \frac{p_{i_5}}{w_{i_5}}$. Since they belong to five bad groups, and within each bad group, the maximum and the minimum item efficiency differ by at least $\frac{1}{\tau}$, we have that $\frac{p_{i_1}}{w_{i_1}} - \frac{p_{i_5}}{w_{i_5}} \geq \frac{3}{\tau}$. But this is impossible since the efficiencies of i_1 and i_5 in the range $[\rho_b - \frac{1}{\tau}, \rho_b + \frac{1}{\tau}]$ whose length is $\frac{2}{\tau}$.

Note that by our construction, I_{m-1} is a good group, so it cannot be one of these four bad groups. Since the first $m-2$ groups are of size τ , to finish the proof, it suffices to show that I_m cannot have an item whose efficiency is within the range $[\rho_b - \frac{1}{\tau}, \rho_b + \frac{1}{\tau}]$. Since $(w^*, p^*) \in \mathcal{S}^+(I; \frac{1}{\alpha\varepsilon})$, we have that $w^* \leq \frac{1}{\alpha\varepsilon}$. Since every item is from $[1, 2] \times [1, 2]$, we have $b \leq \frac{1}{\alpha\varepsilon} + 1$. Recall that $m = \lceil \frac{1}{\alpha\varepsilon\tau} \rceil + 3$ and that the first $m-2$ groups each has τ items. So it must be that $(w_b, p_b) \in I_j$ for some $j \leq m-2$. Let i' be the first item (the item with the highest efficiency) in I_{m-1} . We have $\frac{p_{i'}}{w_{i'}} \leq \frac{p_b}{w_b}$. Let i'' be the first item in I_m . By the construction of I_{m-1} , we have that $\frac{p_{i''}}{w_{i''}} < \frac{p_{i'}}{w_{i'}} - \frac{1}{\tau}$. Therefore, $\frac{p_{i''}}{w_{i''}} < \frac{p_b}{w_b} - \frac{1}{\tau}$. This implies that every item in I_m has its efficiency strictly less than $\frac{p_b}{w_b} - \frac{1}{\tau}$. ◀

Now we are ready to bound the total additive error.

► **Lemma 15.** $S_1 \oplus S_2 \oplus \dots \oplus S_m$ approximates (w^*, p^*) with additive error $O(\frac{1}{\alpha})$.

Proof. Since $I = I_1 \cup \dots \cup I_m$ and $(w^*, p^*) \in \mathcal{S}^+(I)$, we have that there is a tuple $(w_j^*, p_j^*) \in \mathcal{S}^+(I_j)$ for each j such that

$$(w^*, p^*) = (w_1^*, p_1^*) + \dots + (w_m^*, p_m^*).$$

Let J_{good} be the indices of good groups. Let J_{bad} be the set of indices of bad groups. Recall that for each $j \in J_{good}$, S_j approximates (w_j^*, p_j^*) with additive error $O(\varepsilon w_j^*)$, and that for each $j \in J_{bad}$, the set S_j approximates (w_j^*, p_j^*) with additive error $O(\frac{1}{\alpha\tau} \min\{w_j^*, w(I_j) - w_j^*\})$. Therefore, $S_1 \oplus S_2 \oplus \dots \oplus S_m$ approximates (w^*, p^*) with additive error

$$O(\varepsilon \cdot \sum_{j \in J_{good}} w_j^* + \frac{1}{\alpha\tau} \cdot \sum_{j \in J_{bad}} \min\{w_j^*, w(I_j) - w_j^*\}).$$

It is easy to see that

$$\sum_{j \in J_{good}} w_j^* \leq w^* \leq \frac{1}{\alpha\varepsilon}. \quad (1)$$

It remains to bound $\sum_{j \in J_{bad}} \min\{w_j^*, w(I_j) - w_j^*\}$. We further divide J_{bad} into $J_{bad}^H \cup J_{bad}^M \cup J_{bad}^L$ so that

- For $j \in J_{bad}^H$, all items in I_j have efficiency at least $\rho_b + \frac{1}{\tau}$
- For $j \in J_{bad}^M$, some items in I_j has efficiency in the range $[\rho_b - \frac{1}{\tau}, \rho_b + \frac{1}{\tau}]$.
- For $j \in J_{bad}^L$, all items in I_j have efficiency at most $\rho_b - \frac{1}{\tau}$.

The proximity Lemma 9 implies that

$$\sum_{j \in J_{bad}^H} (w(I_j) - w_j^*) \leq 2\tau \quad \text{and} \quad \sum_{j \in J_{bad}^L} w_j^* \leq 2\tau. \quad (2)$$

Lemma 14 implies that $|J_{bad}^M| \leq 4$ and $I_j \leq \tau$ for each $j \in J_{bad}^M$. Recall that the item weights are at most 2. We have

$$\sum_{j \in J_{bad}^M} w_j^* \leq 4 \cdot \tau \cdot 2 = 8\tau. \quad (3)$$

In view of (2) and (3), we have

$$\sum_{j \in J_{bad}} \min\{w_j^*, w(I_j) - w_j^*\} \leq \sum_{j \in J_{bad}^H} (w(I_j) - w_j^*) + \sum_{j \in J_{bad}^L} w_j^* + \sum_{j \in J_{bad}^M} w_j^* \leq 12\tau. \quad (4)$$

In view of (1) and (4), we have $S_1 \oplus S_2 \oplus \dots \oplus S_m$ approximates (w^*, p^*) with additive error

$$O(\varepsilon \cdot \sum_{j \in J_{good}} w_j^* + \frac{1}{\alpha\tau} \cdot \sum_{j \in J_{bad}} \min\{w_j^*, w(I_j) - w_j^*\}) \leq O(\varepsilon \cdot \frac{1}{\alpha\varepsilon} + \frac{1}{\alpha\tau} \cdot 12\tau) = O(\frac{1}{\alpha}). \quad \blacktriangleleft$$

The above lemma immediately implies the following since S approximates $S_0 \oplus S_1 \oplus \dots \oplus S_m$ with factor $1 + O(\varepsilon)$.

► **Corollary 16.** S approximates (w^*, p^*) with additive error $O(\frac{1}{\alpha})$.

We summarize this section by the following lemma.

► **Lemma 17.** *There is an $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$ -time algorithm for the reduced problem, which is randomized and succeeds with high probability.*

6 An Improved Algorithm with Exponent 7/4

Due to the space limit, we present only the most crucial part of the algorithm and the analysis. The missing part can be found in the full version of the paper.

In our $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$ -time algorithm, only the bad groups benefit from the proximity bound, which allows us to approximate with larger factor $1 + \frac{1}{\alpha\tau}$, while for the good groups, we simply approximate with factor $1 + \varepsilon$. To further improve the running time, we shall partition the items in a way that all the groups can benefit from the proximity bound.

We assume that $\frac{1}{2} \leq \alpha \leq (\frac{1}{\varepsilon})^{3/4}$ since the algorithm in Lemma 8 already runs in $\tilde{O}(n + (\frac{1}{\varepsilon})^{7/4})$ time for all $\alpha \geq (\frac{1}{\varepsilon})^{3/4}$. We assume that $I = \{(w_1, p_1), \dots, (w_n, p_n)\}$ and that $\frac{p_1}{w_1} \geq \dots \geq \frac{p_n}{w_n}$.

We first partition I into the head group I_{head} and the tail group I_{tail} where I_{head} contains the first $\lceil \frac{1}{\alpha\varepsilon} \rceil + 1$ items and I_{tail} contains the rest of the items. Below we only give an algorithm for I_{head} , and analyze its additive error. The remaining analysis and the algorithm for I_{tail} can be found in the full version.

6.1 Approximating Head Group

Let $n' = \lceil \frac{1}{\alpha\varepsilon} \rceil + 1$. We shall show that in $\tilde{O}((\frac{1}{\varepsilon})^{7/4})$ time, we can compute a set S of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximate $\mathcal{S}^+(I_{\text{head}})$ with additive error $O(\frac{1}{\alpha})$.

Let τ be a parameter that will be specified later. Roughly speaking, we will further partition I_{head} into good groups and bad groups. The bad groups are the same as before: a bad group is a group of τ items whose efficiencies differ by at least $\frac{1}{\tau}$. For good groups, we will make them large than before: when we obtain a good group I' of size τ , we will keep adding items to I' until the difference between the maximum and the minimum item efficiencies exceeds $\frac{1}{|I'|}$ with the next item. A more precise description of the partition process is given below.

We partition I_{head} into I_1, \dots, I_m as follows. Initially, $j = 1$ and $k = 1$, and the remaining items are $\{(w_k, p_k), \dots, (w_{n'}, p_{n'})\}$. Let k' be the minimum integer such that

$$(k' - k) \cdot \left(\frac{p_k}{w_k} - \frac{p_{k'}}{w_{k'}} \right) > 1.$$

If such k' does not exist, we set $I_j = \{(w_k, p_k), \dots, (w_{n'}, p_{n'})\}$, set $\Delta_j = \Delta'_j = \frac{p_k}{w_k} - \frac{p_{n'}}{w_{n'}}$, set $m := j$, and finish. Assume that k' exists.

- If $k' - k \geq \tau$, we set $I_j = \{(w_k, p_k), \dots, (w_{k'-1}, p_{k'-1})\}$, set $\Delta_j = \frac{p_k}{w_k} - \frac{p_{k'-1}}{w_{k'-1}}$, set $\Delta'_j = \frac{p_k}{w_k} - \frac{p_{k'}}{w_{k'}}$, and proceed with $j := j + 1$ and $k := k'$. In this case, we say I_j is a good group.
- Otherwise, we let $I_j = \{(w_k, p_k), \dots, (w_{k+\tau-1}, p_{k+\tau-1})\}$, set $\Delta_j = \frac{p_k}{w_k} - \frac{p_{k+\tau-1}}{w_{k+\tau-1}}$, set $\Delta'_j = \frac{p_k}{w_k} - \frac{p_{k+\tau}}{w_{k+\tau}}$, and proceed with $j := j + 1$ and $k := k + \tau$. In this case, we say that I_j is a bad group.

(We remark that Δ_j and Δ'_j are not required by the algorithm. They are maintained only for the purpose of analysis. Δ_j is the actual difference between the maximum and minimum item efficiencies in I_j . For technical reasons, we also need Δ'_j . Basically, compared to Δ_j , the gap Δ'_j also includes the efficiency gap between the minimum efficiency in I_j and the maximum efficiency in I_{j+1} . We will use Δ_j to bound the time cost of I_j , and use Δ'_j to create an efficiency gap for other groups.)

We have the following observations.

► **Observation 18.** *The groups I_1, \dots, I_m satisfy the following properties.*

- (i) $m \leq \frac{1}{\alpha\varepsilon\tau} + 1$
- (ii) $\sum_{j=1}^m |I_j| = |I_{\text{head}}| = \frac{1}{\alpha\varepsilon} + 1$ and $\sum_{j=1}^m \Delta_j \leq \sum_{j=1}^m \Delta'_j \leq \frac{3}{2}$.
- (iii) For every group I_j , the efficiencies of the items in I_j differ by at most Δ_j .
- (iv) For every good group I_j , we have $|I_j|\Delta_j \leq 1$ and $|I_j|\Delta'_j \geq \frac{1}{2}$.
- (v) For every bad group I_j , we have $|I_j|\Delta'_j \geq |I_j|\Delta_j > 1$.
- (vi) For the last group I_m , we have $|I_m|\Delta_m \leq 1$.

For each group, we shall approximate it in exactly the same way as we did for the bad groups in the $\tilde{O}(n + (\frac{1}{\varepsilon})^{11/6})$ -time algorithm, except that we shall use $\frac{1}{\alpha|I_j|}$ as the accuracy parameter. More specifically, we use Corollary 12 and Corollary 13 to compute a set S_j of size $\tilde{O}(\alpha|I_j|)$ that approximate each tuple $(w, p) \in \mathcal{S}^+(I_j)$ with additive error $\frac{1}{\alpha|I_j|} \cdot \min\{w, w(I) - w\}$. The time cost for I_j is $\tilde{O}(|I_j| + \alpha|I_j| + \alpha^2|I_j|^2\Delta_j)$.

Then we compute a set S_{head} of size $\tilde{O}(\frac{1}{\varepsilon})$ that approximates $S_1 \oplus \dots \oplus S_m$ with factor $1 + O(\varepsilon)$ via Lemma 3.

Bound Additive Error

We shall show that S_{head} approximates every tuple $(w^*, p^*) \in \mathcal{S}^+(I_{\text{head}})$ with additive error $\tilde{O}(\frac{1}{\alpha})$.

Let (w^*, p^*) be an arbitrary tuple in $\mathcal{S}^+(I_{\text{head}})$. Let (w_b, p_b) be the breaking item with respect to w^* . Let $\rho_b = \frac{p_b}{w_b}$ be the efficiency of the breaking item.

Let $I_{j'}$ be the group containing the breaking item (w_b, p_b) . To apply the proximity bound, we will show that the groups can be divided into $O(\log \frac{1}{\varepsilon})$ collections so that for each collection $\{I_{j_k+1}, I_{j_k+2}, \dots, I_{j_{k+1}-1}\}$, the efficiency gap between ρ_b and the groups in this collection is at least the maximum Δ'_j of these groups. To do this, we need the following auxiliary lemma. Its proof is deferred to the full version.

► **Lemma 19.** *Let $\Delta_1, \dots, \Delta_n$ be a sequence of positive real numbers. Let Δ_{\min} and Δ_{\max} be the minimum and maximum numbers in the sequence. There exists $h = O(\log \frac{\Delta_{\max}}{\Delta_{\min}})$ indices $1 = j_1 < j_2 < \dots < j_h = n$ such that for any $k \in \{1, \dots, h-1\}$, we have that*

$$\max\{\Delta_j : j_k < j < j_{k+1}\} \leq \sum\{\Delta_j : j_{k+1} \leq j \leq n\}, \quad (5)$$

where the maximum of an empty set is defined to be 0.

► **Lemma 20.** $S_1 \oplus \dots \oplus S_m$ approximates (w^*, p^*) with additive error $\tilde{O}(\frac{1}{\alpha})$.

Proof. Since $I_{\text{head}} = I_1 \cup \dots \cup I_m$ and $(w^*, p^*) \in \mathcal{S}^+(I_{\text{head}})$, we have that there is a tuple $(w_j^*, p_j^*) \in \mathcal{S}^+(I_j)$ for each j such that $(w^*, p^*) = (w_1^*, p_1^*) + \dots + (w_m^*, p_m^*)$. Recall that S_j approximates (w_j^*, p_j^*) with the additive error $\delta_j = \frac{1}{\alpha|I_j|} \cdot \min\{w_j^*, w(I_j) - w_j^*\}$. It suffices to bound $\sum_j \delta_j$. Let $I_{j'}$ be the group containing the breaking item (w_b, p_b) . Since the weights of the items are in $[1, 2]$. It is easy to see that $\delta_j \leq \frac{1}{\alpha|I_j|} w_j^* \leq \frac{1}{\alpha}$ for any j . Therefore, $\delta_{j'} + \delta_m \leq \frac{2}{\alpha}$. It remains to consider $\sum_{j=1}^{j'-1} \delta_j$ and $\sum_{j=j'}^{m-1} \delta_j$. We shall bound them using Lemma 19 and the proximity bound (Lemma 9).

Consider $\Delta'_1, \Delta'_2, \dots, \Delta'_{j'-1}$. We shall apply Lemma 19 to them. Let $\Delta'_{\max} = \max_{j=1}^{j'-1} \Delta'_j$ and let $\Delta'_{\min} = \min_{j=1}^{j'-1} \Delta'_j$. Note that $\Delta'_{\max} \leq 3/2$ since the items are from $[1, 2] \times [1, 2]$, and that $\Delta'_{\min} \geq 1/(2 \sum_j |I_j|) \geq \frac{\alpha\varepsilon}{2} \geq \frac{\varepsilon}{4}$ due to Observation 18(ii), (iv), and (v). Therefore, $\log(\frac{\Delta'_{\max}}{\Delta'_{\min}}) = O(\log \frac{1}{\varepsilon})$. By Lemma 19, there exists $h = O(\log \frac{1}{\varepsilon})$ indices $1 = j_1 < j_2 < \dots < j_h = j' - 1$ such that for any $k \in \{1, \dots, h-1\}$, we have that

$$\max\{\Delta'_j : j_k < j < j_{k+1}\} \leq \sum\{\Delta'_j : j_{k+1} \leq j \leq j' - 1\}. \quad (6)$$

Fix some k . Consider the groups I_j with $j_k < j < j_{k+1}$. Let $\Delta' = \max\{\Delta'_j : j_k < j < j_{k+1}\}$. The inequality (6) implies that the items in these groups I_j have efficiencies at least $\rho_b + \Delta'$. By Lemma 9, we have

$$\sum_{j_k < j < j_{k+1}} (w(I_j) - w_j^*) \leq \frac{2}{\Delta'}. \quad (7)$$

Also note that for each I_j with $j_k < j < j_{k+1}$, we have $\frac{1}{|I_j|} \leq 2\Delta'_j \leq 2\Delta'$ (due to Observation 18(iv) and (v)). Therefore,

$$\sum_{j_k < j < j_{k+1}} \delta_j \leq \sum_{j_k < j < j_{k+1}} \frac{w(I_j) - w_j^*}{\alpha |I_j|} \leq \sum_{j_k < j < j_{k+1}} \frac{2\Delta'(w(I_j) - w_j^*)}{\alpha} \leq \frac{1}{\alpha}.$$

The last inequality is due to (7). Recall that $\delta_j \leq \frac{1}{\alpha}$ for any j . We have the following.

$$\sum_{j=1}^{j'-1} \delta_j = \sum_{k=1}^h \delta_{j_k} + \sum_{k=1}^{h-1} \sum_{j_k < j < j_{k+1}} \delta_j = O\left(\frac{h}{\alpha}\right) = \tilde{O}\left(\frac{1}{\alpha}\right).$$

In a symmetric way, we can show that $\sum_{j=j'+1}^{m-1} \delta_j \leq \tilde{O}\left(\frac{1}{\alpha}\right)$. ◀

The above lemma immediately implies the following since S_{head} approximates $S_1 \oplus \dots \oplus S_m$ with factor $1 + O(\varepsilon)$.

► **Corollary 21.** S_{head} approximates (w^*, p^*) with additive error $\tilde{O}\left(\frac{1}{\alpha}\right)$.

The remaining analysis and the algorithm for I_{tail} can be found in the full version. We summarize this with the following lemma.

► **Lemma 22.** *There is an $\tilde{O}(n + (\frac{1}{\varepsilon})^{7/4})$ -time algorithm for the reduced problem, which is randomized and succeeds with high probability.*

References

- 1 Alok Aggarwal, Maria M. Klawe, Shlomo Moran, Peter Shor, and Robert Wilber. Geometric applications of a matrix-searching algorithm. *Algorithmica*, 2(1):195–208, November 1987. doi:10.1007/BF01840359.
- 2 Kyriakos Axiotis and Christos Tzamos. Capacitated Dynamic Programming: Faster Knapsack and Graph Algorithms. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, pages 19:1–19:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.19.
- 3 David Bremner, Timothy M Chan, Erik D Demaine, Jeff Erickson, Ferran Hurtado, John Iacono, Stefan Langerman, Mihai Pătraşcu, and Perouz Taslakian. Necklaces, convolutions, and $x+y$. *Algorithmica*, 69(2):294–314, 2014. doi:10.1007/S00453-012-9734-3.
- 4 Karl Bringmann. A Near-Linear Pseudopolynomial Time Algorithm for Subset Sum. In *Proceedings of the 2017 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*, pages 1073–1084. Society for Industrial and Applied Mathematics, 2017. doi:10.1137/1.9781611974782.69.
- 5 Karl Bringmann. Knapsack with small items in near-quadratic time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 259–270. Association for Computing Machinery, 2024. doi:10.1145/3618260.3649719.
- 6 Karl Bringmann and Alejandro Cassis. Faster Knapsack Algorithms via Bounded Monotone Min-Plus-Convolution. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, pages 31:1–31:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ICALP.2022.31.
- 7 Karl Bringmann and Alejandro Cassis. Faster 0-1-Knapsack via Near-Convex Min-Plus-Convolution. In *31st Annual European Symposium on Algorithms (ESA 2023)*, pages 24:1–24:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ESA.2023.24.
- 8 Karl Bringmann, Anita Dür, and Adam Polak. Even faster knapsack via rectangular monotone min-plus convolution and balancing. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, September 2-4, 2024, Royal Holloway, London, United Kingdom*, volume 308 of *LIPIcs*, pages

- 33:1–33:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ESA.2024.33.
- 9 Karl Bringmann and Vasileios Nakos. A Fine-Grained Perspective on Approximating Subset Sum and Partition. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA 2021)*, pages 1797–1815. Society for Industrial and Applied Mathematics, 2021. doi:10.1137/1.9781611976465.108.
 - 10 Timothy M. Chan. Approximation Schemes for 0-1 Knapsack. In *1st Symposium on Simplicity in Algorithms (SOSA 2018)*, pages 5:1–5:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/OASIcs.SOSA.2018.5.
 - 11 Timothy M. Chan and Moshe Lewenstein. Clustered Integer 3SUM via Additive Combinatorics. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing (STOC 2015)*, pages 31–40. Association for Computing Machinery, 2015. doi:10.1145/2746539.2746568.
 - 12 Timothy M. Chan and Ryan Williams. Deterministic apsp, orthogonal vectors, and more: Quickly derandomizing razborov-smolensky. In *Proceedings of the 2016 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016)*, pages 1246–1255, 2016. doi:10.1137/1.9781611974331.ch87.
 - 13 Lin Chen, Jiayi Lian, Yuchen Mao, and Guochuan Zhang. Approximating partition in near-linear time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 307–318. Association for Computing Machinery, 2024. doi:10.1145/3618260.3649727.
 - 14 Lin Chen, Jiayi Lian, Yuchen Mao, and Guochuan Zhang. Faster algorithms for bounded knapsack and bounded subset sum via fine-grained proximity results. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2024)*, pages 4828–4848. Society for Industrial and Applied Mathematics, 2024. doi:10.1137/1.9781611977912.171.
 - 15 Lin Chen, Jiayi Lian, Yuchen Mao, and Guochuan Zhang. A nearly quadratic-time fptas for knapsack. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 283–294. Association for Computing Machinery, 2024. doi:10.1145/3618260.3649730.
 - 16 Shucheng Chi, Ran Duan, Tianle Xie, and Tianyi Zhang. Faster min-plus product for monotone instances. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*, pages 1529–1542. Association for Computing Machinery, 2022. doi:10.1145/3519935.3520057.
 - 17 Marek Cygan, Marcin Mucha, Karol Węgrzycki, and Michał Włodarczyk. On problems equivalent to $(\min, +)$ -convolution. *ACM Transactions on Algorithms*, 15(1):1–25, January 2019. doi:10.1145/3293465.
 - 18 Mingyang Deng, Ce Jin, and Xiao Mao. Approximating Knapsack and Partition via Dense Subset Sums. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2023)*, pages 2961–2979. Society for Industrial and Applied Mathematics, 2023. doi:10.1137/1.9781611977554.ch113.
 - 19 Friedrich Eisenbrand and Robert Weismantel. Proximity Results and Faster Algorithms for Integer Programming Using the Steinitz Lemma. *ACM Transactions on Algorithms*, 16(1):5:1–5:14, November 2019. doi:10.1145/3340322.
 - 20 Aleksei V. Fishkin, Olga Gerber, Klaus Jansen, and Roberto Solis-Oba. On packing squares with resource augmentation: maximizing the profit. In *Proceedings of the 2005 Australasian Symposium on Theory of Computing, CATS '05*, pages 61–67, AUS, 2005. Australian Computer Society, Inc. URL: <https://dl.acm.org/doi/10.5555/1082260.1082268>.
 - 21 Oscar H. Ibarra and Chul E. Kim. Fast Approximation Algorithms for the Knapsack and Sum of Subset Problems. *Journal of the ACM*, 22(4):463–468, October 1975. doi:10.1145/321906.321909.
 - 22 Kazuo Iwama and Guochuan Zhang. Online knapsack with resource augmentation. *Information Processing Letters*, 110(22):1016–1020, 2010. doi:10.1016/j.ipl.2010.08.013.

- 23 Klaus Jansen and Stefan E. J. Kraft. A faster FPTAS for the Unbounded Knapsack Problem. *European Journal of Combinatorics*, 68:148–174, February 2018. doi:10.1016/j.ejc.2017.07.016.
- 24 Ce Jin. An improved FPTAS for 0-1 knapsack. In *Proceedings of 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, pages 76:1–76:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.76.
- 25 Ce Jin. 0-1 Knapsack in Nearly Quadratic Time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 271–282. Association for Computing Machinery, 2024. doi:10.1145/3618260.3649618.
- 26 Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer US, Boston, MA, 1972. doi:10.1007/978-1-4684-2001-2_9.
- 27 Hans Kellerer, Renata Mansini, Ulrich Pferschy, and Maria Grazia Speranza. An efficient fully polynomial approximation scheme for the Subset-Sum Problem. *Journal of Computer and System Sciences*, 66(2):349–370, March 2003. doi:10.1016/S0022-0000(03)00006-0.
- 28 Hans Kellerer and Ulrich Pferschy. A New Fully Polynomial Time Approximation Scheme for the Knapsack Problem. *Journal of Combinatorial Optimization*, 3(1):59–71, July 1999. doi:10.1023/A:1009813105532.
- 29 Hans Kellerer and Ulrich Pferschy. Improved Dynamic Programming in Connection with an FPTAS for the Knapsack Problem. *Journal of Combinatorial Optimization*, 8(1):5–11, March 2004. doi:10.1023/B:JOCO.0000021934.29833.6b.
- 30 Hans Kellerer, Ulrich Pferschy, and Pisinger David. *Knapsack Problems*. Springer, Berlin Heidelberg, 2004. doi:10.1007/978-3-540-24777-7.
- 31 Konstantinos Koiliaris and Chao Xu. Faster Pseudopolynomial Time Algorithms for Subset Sum. *ACM Transactions on Algorithms*, 15(3):1–20, July 2019. doi:10.1145/3329863.
- 32 Marvin Künnemann, Ramamohan Paturi, and Stefan Schneider. On the Fine-Grained Complexity of One-Dimensional Dynamic Programming. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, pages 21:1–21:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.ICALP.2017.21.
- 33 Eugene L. Lawler. Fast Approximation Algorithms for Knapsack Problems. *Mathematics of Operations Research*, 4(4):339–356, November 1979. doi:10.1287/MOOR.4.4.339.
- 34 Xiao Mao. $(1 - \epsilon)$ -approximation of knapsack in nearly quadratic time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC 2024)*, pages 295–306. Association for Computing Machinery, 2024. doi:10.1145/3618260.3649677.
- 35 Marcin Mucha, Karol Węgrzycki, and Michał Włodarczyk. A Subquadratic Approximation Scheme for Partition. In *Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*, pages 70–88. Society for Industrial and Applied Mathematics, 2019. doi:10.1137/1.9781611975482.5.
- 36 Adam Polak, Lars Rohwedder, and Karol Węgrzycki. Knapsack and Subset Sum with Small Items. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, pages 106:1–106:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICALP.2021.106.
- 37 Donguk Rhee. *Faster fully polynomial approximation schemes for knapsack problems*. PhD thesis, Massachusetts Institute of Technology, 2015. URL: <https://dspace.mit.edu/handle/1721.1/98564>.
- 38 Ryan Williams. Faster all-pairs shortest paths via circuit complexity. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing (STOC 2014)*, pages 664–673. Association for Computing Machinery, 2014. doi:10.1145/2591796.2591811.