Approximation Algorithms for Interdiction Problem with Packing Constraints

Lin Chen[1]
Department of Computer Science, Texas Tech University, Lubbock, TX, USA

Xiaoyu Wu
School of Mathematical Sciences, Zhejiang University, Hangzhou, China

Guochuan Zhang
School of Computer Science, Zhejiang University, Hangzhou, China

Abstract

We study a bilevel optimization problem which is a zero-sum Stackelberg game. In this problem, there are two players, a leader and a follower, who pick items from a common set. Both the leader and the follower have their own (multi-dimensional) budgets, respectively. Each item is associated with a profit, which is the same to the leader and the follower, and will consume the leader's (follower's) budget if it is selected by the leader (follower). The leader and the follower will select items in a sequential way: First, the leader selects items within the leader's budget. Then the follower selects items from the remaining items within the follower's budget. The goal of the leader is to minimize the maximum profit that the follower can obtain. Let s_A and s_B be the dimension of the leader's and follower's budget, respectively. A special case of our problem is the bilevel knapsack problem studied by Caprara et al. [SIAM Journal on Optimization, 2014], where $s_A = s_B = 1$. We consider the general problem and obtain an $(s_B + \epsilon)$-approximation algorithm when s_A and s_B are both constant. In particular, if $s_B = 1$, our algorithm implies a PTAS for the bilevel knapsack problem, which is the first $\mathcal{O}(1)$-approximation algorithm. We also complement our result by showing that there does not exist any $(4/3 - \epsilon)$-approximation algorithm even if $s_A = 1$ and $s_B = 2$. We also consider a variant of our problem with resource augmentation when s_A and s_B are both part of the input. We obtain an $\mathcal{O}(1)$-approximation algorithm with $\mathcal{O}(1)$-resource augmentation, that is, we give an algorithm that returns a solution which exceeds the given leader's budget by $\mathcal{O}(1)$ times, and the objective value achieved by the solution is $\mathcal{O}(1)$ times the optimal objective value that respects the leader's budget.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Bilevel Integer Programming, Interdiction Constraints, Knapsack

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.39

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2204.11106` [7]

Funding *Lin Chen*: NSF No. 2004096
Guochuan Zhang: NSFC Key Program No. 12131003

1 Introduction

In recent years, there is an increasing interest in adopting the *Stackelberg competition model* [16] to address the critical security concern that arises in protecting our ports, airports, transportation, and other critical national infrastructures (see, e.g., [1, 29, 33]). In these

[1] Optional footnote, e.g. to mark corresponding author

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).
Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;
Article No. 39; pp. 39:1–39:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

problems, the attacker's target is to maximize the illicit gain, while the defender tries to mitigate the attack by minimizing the attacker's objective through deploying defending resources.

In this paper, we consider an abstract model for general defending problems called *interdiction with packing constraints* (IPC). In IPC, given are a set of items, together with a *leader* and a *follower*. Both the leader and the follower have their own (multi-dimensional) budgets, respectively. Each item is associated with a profit, which is the same to the leader and the follower, and will consume the leader's (follower's) budget if it is selected by the leader (follower). The leader and the follower will select items in a sequential way: First, the leader selects items within the leader's budget. Then the follower selects items from the remaining items within the follower's budget. The goal of the leader is to minimize the maximum profit that the follower can obtain. IPC captures the general setting where the follower is the attacker who gets profit by attacking items, and the leader is the defender who tries to minimize the attacker's gain by protecting a subset of items.

IPC can be formulated as a bilevel integer program (IP) as follows. Denote by $I = \{1, 2, \cdots, n\}$ the set of items. Each item $j \in I$ is associated with a *profit* $p_j \in \mathbb{Q}_{>0}$, an s_A-dimensional *cost vector* $\mathbf{A}_j \in \mathbb{Q}_{\geq 0}^{s_A}$ to the leader and an s_B-dimensional *weight vector* $\mathbf{B}_j \in \mathbb{Q}_{\geq 0}^{s_B}$ to the follower. The leader and the follower have their own *budget vectors*, denoted by $\mathbf{a} \in \mathbb{Q}_{\geq 0}^{s_A}$ and $\mathbf{b} \in \mathbb{Q}_{\geq 0}^{s_B}$, respectively. We introduce 0-1 variables x_j and y_j for each $j \in I$ as the decision variables for the leader and the follower. More precisely, if the leader chooses item j, then $x_j = 1$, otherwise $x_j = 0$. Similarly, $y_j = 1$ if the follower chooses item j and $y_j = 0$ otherwise. Denote by $\mathbf{x} = (x_1, x_2, \ldots, x_n)$, $\mathbf{y} = (y_1, y_2, \ldots, y_n)$, $\mathbf{p} = (p_1, p_2, \ldots, p_n)$ and $\mathbf{1} = (\underbrace{1, \ldots, 1}_{n})$. IPC can be formulated as a bilevel program $\mathbf{IPC}(I, \mathbf{a}, \mathbf{b})$ as follows:

$$\mathbf{IPC}(I, \mathbf{a}, \mathbf{b}) : \min_{\mathbf{x}} \ \mathbf{py} \tag{1a}$$

$$s.t. \ \mathbf{Ax} \leq \mathbf{a} \tag{1b}$$

$$\mathbf{x} \in \{0, 1\}^n \tag{1c}$$

where $\mathbf{y}$ solves the following:

$$\max_{\mathbf{y}} \quad \mathbf{py} \tag{1d}$$

$$s.t. \quad \mathbf{By} \leq \mathbf{b} \tag{1e}$$

$$\mathbf{x} + \mathbf{y} \leq \mathbf{1} \tag{1f}$$

$$\mathbf{y} \in \{0, 1\}^n \tag{1g}$$

where $\mathbf{A} = (\mathbf{A}_1, \cdots, \mathbf{A}_n)$ and $\mathbf{B} = (\mathbf{B}_1, \cdots, \mathbf{B}_n)$ are $s_A \times n$ and $s_B \times n$ non-negative rational matrices, respectively.

The most relevant prior work to our IPC model is the well-known knapsack interdiction problem introduced by DeNegre [14], which is the special case of IPC where $s_A = 1$ and $s_B = 1$. Very recently, Caprara et al. [4] proved that DeNegre's knapsack interdiction problem is $\sum_2^p$-complete and strongly NP-hard, which also implies the $\sum_2^p$-completeness and strongly NP-hardness for IPC. Caprara et al. showed a polynomial time approximation scheme (PTAS) for a special case of knapsack interdiction problem where the profit of an item is equal to its weight to the follower.

Except for the knapsack interdiction problem, we are not aware of any approximation algorithms for other special cases of IPC. However, if we relax the follower's problem by allowing $\mathbf{y}$ to take fractional value, then there are several research works in the literature. The most relevant work is the packing interdiction problem studied by Dinitz and Gupta [15], where the follower's problem is given by $\max\{\sum_j p_j(1-x_j)y_j : B\mathbf{y} \leq \mathbf{b}, \mathbf{y} \geq \mathbf{0}\}$, while the leader's constraints are the same as Equation 1b and Equation 1c except that $s_A = 1$. Dinitz and Gupta provided an approximation algorithm whose ratio depends on the sparsity of the matrix B. Their techniques crucially rely on the fact that $\mathbf{y}$ can take fractional value, and therefore duality theory can be applied to the follower's problem, allowing the bilevel problem to be transformed to a single level problem. Besides the packing interdiction problem, quite a few graph interdiction problems have been studied in the literature, where the follower's problem is a standard graph optimization problem, and the leader can remove edges or vertices to minimize the follower's optimal objective value on the graph after edge-removal or vertex-removal. On planar graphs, polynomial time approximation schemes (PTASs) were obtained for network flow interdiction [32, 37] and matching interdiction [28]. On general graphs, approximation algorithms were also obtained for, e.g., connectivity interdiction [38], minimum spanning tree interdiction [26, 39], matching interdiction [15, 36], network flow interdiction [3, 9, 10], etc. All of these algorithms crucially rely on the follower's specific graph optimization problem and do not apply directly to IPC.

Our Contributions

The main contribution of this paper is an $(s_B + \epsilon)$-approximation polynomial time algorithm for IPC when s_A and s_B are both constant. In particular, when $s_B = 1$, our algorithm is a PTAS. Since the knapsack interdiction problem is a special case of IPC when $s_A = s_B = 1$, our result gives the first $\mathcal{O}(1)$-approximation algorithm for this problem. To complement our result, we also show that IPC does not admit any $(4/3 - \epsilon)$-approximation algorithm even if $s_B = 2$ and $s_A = 1$, assuming $P \neq NP$. This implies that the PTAS for $s_B = 1$ cannot be further extended to the case of $s_B \geq 2$.

We also consider a natural variant of IPC where the leader's budget can be violated. For this variant we obtain a $(\frac{\rho}{1-\alpha}, \frac{1}{\alpha})$-bicriteria approximation algorithm for any $\alpha \in (0, 1)$, which runs in polynomial time when s_A and s_B are arbitrary (not necessarily polynomial in the input size). More precisely, the algorithm takes as input two oracles, a ρ-approximation algorithm to the follower's optimization problem $\max\{\mathbf{py} : \mathbf{By} \leq \mathbf{b}, \mathbf{y} \in \{0, 1\}^n\}$; and a separation oracle for the leader's problem that given any $\mathbf{x} = \mathbf{x}^0$, it either asserts that $\mathbf{Ax}^0 \leq \mathbf{a}$ or returns a violating constraint. Then in polynomial oracle time the algorithm returns a solution $\mathbf{x}^*$ for the leader such that $\mathbf{Ax}^* \leq \frac{1}{\alpha}\mathbf{a}$, and the objective value is at most $\frac{\rho T^*}{1-\alpha}$, where T^* is the optimal objective value with the leader's budget being $\mathbf{a}$. When we take, e.g., $\alpha = 1/2$, we achieve an objective of $2\rho T^*$ with the leader's budget augmented to $2\mathbf{a}$.

In terms of techniques, our main contribution is a general method for bilevel optimization problems where the leader's and follower's decision variables are both integral. Most prior works on bilevel optimization require follower's decision variables to take fractional values, which accommodates the application of LP duality to transform the bilevel optimization problem to a standard (single level) optimization, and are thus inapplicable when the follower's decision variables become integral. A common technique used in many single level optimization problems is to first classify items into large and small based on whether they can make a significant contribution to the objective value, then guess out large items via enumeration, and handle small items fractionally via LP (see, e.g., [5, 21, 22]. However, such a technique encounters a fundamental challenge in IPC: we can guess out all large items selected by the leader, however, the follower may still select arbitrarily from the remaining large items. In other word, the follower's choice on large items can never be guessed out, and therefore we cannot apply duality to the follower's problem. We overcome the challenge based on the following two ideas: First, we show that given leader's choice on large items, there is a fixed number of "dominant choices" such that the follower's choice on remaining large items always belong to the dominant choices. Second, we show that there exists a subset of "critical items" such that the follower's choice on small items can be characterized through linear constraints given that these critical items are known. The two observations allow us to transform the bilevel program for IPC to an LP without utilizing duality.

The characterization of dominant choices and critical items become sophisticated in the general case when s_B is an arbitrary constant, but is much simpler in the special case $s_B = 1$. Hence, for ease of presentation, in the main part we present our algorithm for the special case to give an overview on the technical insights, and meanwhile provide a proof sketch towards generalizing the algorithm for the general case. Our techniques may be of separate interest to other bilevel optimization problems.

Related work

Our IPC problem lies generally in the area of *bilevel optimization*, which has received extensive research in the literature. Jeroslow [23] showed that in general, bilevel optimization problems are NP-hard even when the objectives and the constraints are linear. We refer readers to Colson et al. [11] for a comprehensive survey on bilevel optimization.

Within the area of bilevel optimization, *Mixed-Integer Bilevel Linear Problem* (MIBLP) is related to our IPC. MIBLP is a bilevel optimization problem where the objective functions and the constraints for the leader and follower are both linear. MIBLP has been studied extensively in the literature, see, e.g., [19, 20, 35]. We also refer the reader to [18, 25] for an overview on MIBLP solvers and related applications. Most of these algorithmic results are for finding exact solutions through, e.g., branch and bound based approach. For DeNegre's knapsack interdiction problem, an improved exact algorithm was derived by Federico Della Croce and Rosario Scatamacchia [12].

It is worth mentioning that besides DeNegre's knapsack interdiction problem (i.e., $s_A = s_B = 1$ in IPC), other variants of bilevel knapsack problems have also been studied in which the leader interferes the follower's program in a different way. One kind of bilevel knapsack problem was introduced by Dempe and Richter [13] where two players hold one knapsack, the leader determines the knapsack's capacity while the follower picks items into the knapsack to maximize his own total profit. The goal is to maximize the objective of the leader. Brotcorne et al. [2] gave a dynamic programming algorithm for both cases of this model. Chen and Zhang [8] proposed a bilevel knapsack variant where two players hold their own knapsacks and the leader can only influence the profit of the items. The follower is interested in his own revenue while the leader aims at maximizing the total profit of both players. The improved approximation results for this problem were derived by Xian Qiu and Walter Kern [34]. Another bilevel knapsack variant occurred in the work of Pferschy et al. [30] where the leader controls the weights of a subset of the follower's items and the follower aims at maximizing his own profit. The leader's payoff is the total weight of the items he controls and selected by the follower. Very recently, Pferschy et al. [31] tackled a "symmetrical" problem in which the leader can control the profits instead of item weights. In addition to these works, a matrix interdiction problem was studied by Kasiviswanathan and Pan [24].

It is also worth mentioning that the continuous version of DeNegre's knapsack interdiction problem, where the leader and the follower can both fractionally choose an item, has also been studied in recent years. Carvalho et al. [6] gave the first polynomial time optimal algorithm. Later on, a faster optimal algorithm was proposed by Woeginger and Fischer [17].

Some notations. We write column vectors in boldface, e.g. $\mathbf{x}, \mathbf{y}$, and their entries in normal font. For a vector $\mathbf{x}$, we either denote its entries by $\mathbf{x} = (x_1, x_2, \cdots, x_n)$, or by $\mathbf{x} = (\mathbf{x}[1], \mathbf{x}[2], \cdots, \mathbf{x}[n])$. Given two vectors $\mathbf{x}$ and $\mathbf{y}$ with the same dimension, we use $\mathbf{xy}$ to represent their dot product, i.e., $\mathbf{xy} = \sum_j x_j y_j$.

2 Hardness results

▶ **Theorem 1.** *Assuming $P \neq NP$, for arbitrary small $\epsilon > 0$, there does not exist a $(4/3 - \epsilon)$-approximation polynomial time algorithm for IPC when $s_A \geq 1$ and $s_B \geq 2$.*

Towards the proof, we need the 3 hitting set (3HS) problem.

Problem: 3 Hitting Set

Instance: A ground set $U = \{u_1, u_2, \cdots, u_n\}$; a collection C of m subsets $S_1, S_2, \cdots, S_m$ whose union is U, where each subset S_h contains exactly 3 elements; a positive integer k.

Question: Is there a hitting subset $S \subseteq U$ such that $|S| \leq k$, and S contains at least one element from each subset in C?

Proof. Recall that 3HS problem is a natural generalization of the well-known Vertex Cover problem, and both are NP-complete [27]. Our reduction is from the 3HS problem. Given an instance of the 3HS, we construct an instance of the IPC where $s_A = 1$, $s_B = 2$ as follows. Let $E = 10 \cdot \sum_{i=1}^{n} 10^i$, and Q be any sufficiently large integer, say, $Q = 10E$. Let $\mathbf{a} = k$ and $\mathbf{b} = (E, 4Q - E)$. The profit of every item constructed below is 1. For every element u_i, we construct an *element-item* (item i) whose interdiction cost is 1, and whose weight vector is $(10^i, Q - 10^i)$. For every subset $S_h = \{u_i, u_j, u_k\}$, we construct a *set-item* (item $n+h$) whose interdiction cost is $k+1$ (that is, the leader cannot interdict a set-item), and whose weight vector is $(E - 10^i - 10^j - 10^k, Q - E + 10^i + 10^j + 10^k)$. In total we construct $n+m$ items.

We first claim that the objective value of any feasible solution for the IPC instance is at most 4. Suppose on the contrary the claim is false, then the follower is able to select at least 5 items under the budget $\mathbf{b} = (E, 4Q - E)$. Notice that for any $1 \leq i \leq n$, $Q - 10^i > Q - E \geq 0.9Q$, and for any $1 \leq i, j, k \leq n$ we have $Q - E + 10^i + 10^j + 10^k > 0.9Q$, if we sum up the weight vectors of any 5 items, then the second coordinate is at least $4.5Q$, which exceeds the budget $4Q - E$, hence the claim is true.

Suppose the 3HS instance admits hitting set S of size at most k, we show that the optimal objective value of the IPC instance is at most 3. Let $S = \{u_{\ell_1}, u_{\ell_2}, \cdots, u_{\ell_k}\}$ (if S contains less than k elements, we simply add arbitrary elements to make it contain exactly k items), then we consider the solution $\mathbf{x}$ where $x_{\ell_i} = 1$ for $1 \leq i \leq k$, and $x_j = 0$ otherwise. We claim that for any $\mathbf{y}$ satisfying $\mathbf{x} + \mathbf{y} \leq \mathbf{1}$, $\mathbf{py} = \sum_j y_j \leq 3$. Suppose on the contrary that the claim is false, then the follower can select at least 4, and hence exactly 4 items (given our claim in the above paragraph that shows the objective value cannot exceed 4). Notice that for every item, if we add the first and second coordinate of its weight vector, then the sum is exactly Q. Hence, if we add up the weight vector of the 4 items, it must be $(z, 4Q - z)$ for some z, and meanwhile, we have $(z, 4Q - z) \leq (E, 4Q - E)$, that is $z \leq E$ and $4Q - z \leq 4Q - E$. Hence, $z = E$, which means the sum of the first coordinate of the weight vectors of the 4 items is exactly $E = 10 \sum_i 10^i$. We first observe that it is impossible for the 4 items to be all element-items, this is because the first coordinate of the weight vector for any element-item is at most $10^n < 0.1E$. We then observe that there cannot be two set-items among the 4 items, because the first coordinate of the weight vector for any set-item is at least $E - 0.1E = 0.9E$. Hence, among the 4 items, there must be exactly 1 set-item and 3 element-items. Let the 3 element-items be those corresponding to u_i, u_j, u_k and the set-item be the one corresponding to $\{u_{i'}, u_{j'}, u_{k'}\}$, then it follows that $10^i + 10^j + 10^k + E - 10^{i'} - 10^{j'} - 10^{k'} = E$, implying that $\{i, j, k\} = \{i', j', k'\}$. However, this is not possible because the hitting set S contains at least one element from $\{u_{i'}, u_{j'}, u_{k'}\} = \{u_i, u_j, u_k\}$, which implies that $x_i + x_j + x_k \geq 1$, and whereas y_i, y_j, y_k cannot be 1 simultaneously. Thus, the optimal objective value of the IPC instance is at most 3.

Suppose the optimal objective value of the IPC instance is at most 3, we show that the 3HS problem admits a hitting set of size at most k. Let $\mathbf{x}^*$ be the optimal solution for IPC. Consider the set $S^* = \{u_i : x_i^* = 1\}$. Given that $\sum_i x_i \leq k$, we know $|S^*| \leq k$. We claim that S^* is a hitting set. Suppose on the contrary that the claim is false, then there exists some subset $\{u_i, u_j, u_k\}$ such that $S^* \cap \{u_i, u_j, u_k\} = \emptyset$. Then we consider the 3 element-items whose weight vectors are $(10^i, Q - 10^i)$, $(10^j, Q - 10^j)$, $(10^k, Q - 10^k)$, and the set-item whose weight vector is $(E - 10^i - 10^j - 10^k, Q - E + 10^i + 10^j + 10^k)$. It is easy to see that the follower can select all the 4 items, leading to an objective value of 4, contradicting the fact that the optimal objective value is at most 3.

Now suppose there exists a $(4/3 - \epsilon)$-approximation polynomial time algorithm for the IPC. We apply the algorithm to the IPC instance constructed from the 3HS instance. If the 3HS instance admits a hitting set of size at most k, then the approximation algorithm returns a solution with objective value at most $4 - \epsilon < 4$, which means it must return a solution with objective value at most 3. If the 3HS instance does not admit a hitting set of size at most k, then the approximation algorithm returns a solution with objective value at least 4. Hence the polynomial time approximation algorithm can be used to determine whether 3HS problem admits a feasible solution, contradicting the NP-hardness of 3HS problem. ◀

3 A PTAS for IPC where $s_B = 1$ and s_A is a fixed constant

The goal of this section is to prove the following Theorem 2. Theorem 2 is a special case of our main result, however, its proof shares similar key ideas as the general case (where s_A and s_B are arbitrary fixed constants). Therefore, we provide a full presentation to demonstrate the technical insights, and in the next section we will show how to extend the techniques when $s_B \geq 2$.

▶ **Theorem 2.** *When $s_B = 1$ and s_A is an arbitrary fixed constant, there exists a polynomial time approximation scheme for IPC.*

The rest of this section is dedicated to proving the following Lemma 3, which implies Theorem 2 directly by scaling item profits (here we write $\mathbf{IPC}(I, \mathbf{a}, b)$ instead of $\mathbf{IPC}(I, \mathbf{a}, \mathbf{b})$ as $\mathbf{b}$ becomes 1-dimensional given that $s_B = 1$).

▶ **Lemma 3.** *Let OPT be the optimal objective value of $\boldsymbol{IPC}(I, \mathbf{a}, b)$. If $OPT \leq 1$, then for an arbitrarily small number $\epsilon > 0$, there exists a polynomial time algorithm that returns a feasible solution to $\boldsymbol{IPC}(I, \mathbf{a}, b)$ with an objective value of at most $1 + \mathcal{O}(\epsilon)$.*

3.1 Preprocessing

From now on we assume $OPT \leq 1$. Without loss of generality, we further assume that $\max_j p_j \leq 1$.

Scaling. We scale the matrix $\mathbf{A}$ and $\mathbf{B}$ such that $\mathbf{a} = \mathbf{1}$ and $b = 1$. From now on we denote this IPC instance as $\mathbf{IPC}(I, \mathbf{1}, 1)$. Without loss of generality, we further assume that $\max_j B_j \leq 1$.

Rounding down the profits. We apply the standard geometric rounding. Let $\delta > 0$ be some small parameter to be fixed later (in particular, we can choose $\delta = \epsilon^2$). Consider each item profit p_j. If $p_j \leq \delta^2$, we keep it as it is; otherwise $p_j > \delta^2$, we round the profit down to the largest value of the form $\delta^2(1+\delta)^h$. For profits whose values are at least δ^2, simple calculation shows there are at most $\tilde{\mathcal{O}}(1/\delta)$ distinct rounded profits. This rounding scheme introduces an additive loss of at most $\mathcal{O}(\delta)$ times the objective value. For simplicity, we still denote the rounded profits by p_j's.

Item classification. Recall that each item j is associated with a profit p_j and a weight vector $\mathbf{B}_j$. Since $s_B = 1$, we write B_j as its weight.

Classifying Weights: We say an item j has a *large* weight if $B_j > \delta$; otherwise, it has a *small* weight.

Classifying Profits: We say an item j has a *large* profit if $p_j > \delta$; a *medium* profit if $\delta^2 < p_j \leq \delta$; and a *small* profit if $p_j \leq \delta^2$.

We say an item is *large* if it has a large-profit, or a large-weight. Otherwise, the item is *small*. Large items and small items will be handled separately.

Denote by S^* the items selected by the leader in an optimal solution of $\mathbf{IPC}(I, \mathbf{1}, 1)$.

3.2 Handling Large Items

3.2.1 Determining the leader's choice on large items

The goal of this subsection is to guess large items in S^* in polynomial time.

Large-profit small-weight items. Notice that if there are at least $1/\delta$ such items for the follower to select, then selecting any $1/\delta$ of them gives a solution with an objective value strictly larger than 1, contradicting to the assumption that $OPT \leq 1$. Thus S^* must include all except at most $1/\delta - 1$ such items, which can be guessed out via $n^{\mathcal{O}(1/\delta)}$ enumerations. Hence, we have the following observation.

▶ **Observation 4.** *With $n^{\mathcal{O}(1/\delta)}$ enumerations, we can guess out all large-profit small-weight items in S^*.*

Small-profit large-weight items. Notice that the follower can select at most $1/\delta$ items from this subgroup and their total profit is at most $\delta^2 * \frac{1}{\delta} = \delta$. Hence, even if the leader does not select any such item, the objective value can increase by at most δ, which leads to the following observation.

▶ **Observation 5.** *With $\mathcal{O}(\delta)$ additive error, we may assume that S^* does not contain small-profit large-weight items.*

Large/medium-profit large-weight items. Notice that the follower can select at most $1/\delta$ items from this group. Since we are considering the case of $s_B = 1$, if there are two items that are not selected by the leader, and they have the same profit, then the follower always prefers the one with a smaller weight. Hence, we have the following lemma.

▶ **Lemma 6.** *With $n^{\tilde{\mathcal{O}}(1/\delta^2)}$ enumerations, we can guess out all large/medium-profit large-weight items in S^*.*

Proof. Recall that there are at most $\tilde{\mathcal{O}}(1/\delta)$ distinct large/medium profits. Let S_h be the set of large-weight items whose profits are all $\delta^2(1+\delta)^h$. We observe two facts: (i). Among items in $S_h \setminus S^*$, the follower always selects the ones with the smallest weights; (ii). The follower can select at most $1/\delta$ items from $S_h \setminus S^*$. We claim that, $S^* \cap S_h$ can be determined through guessing out the following $1/\delta$ key items in S_h: among items in $S_h \setminus S^*$, which is the item that has the k-th smallest weight for $k = 1, 2, \cdots, 1/\delta$?[2] To see the claim, let w_h^{max} be the weight of the item in $S_h \setminus S^*$ that has the $1/\delta$-th smallest weight. Consider any item in S_h: if its weight is smaller than w_h^{max}, and is not one of the key items, then this item must belong to S^* (by the definition of key items); if its weight is larger than or equal to w_h^{max}, and is not one of the key items, then it is not in S^* (there is no need for the leader to select such an item since the follower will never select this item even if it is available). Thus via $n^{\mathcal{O}(1/\delta)}$ enumerations, we can guess out all large/medium-profit large-weight items in $S^* \cap S_h$. Moreover, via total $n^{\tilde{\mathcal{O}}(1/\delta^2)}$ enumerations, we can guess out all large/medium-profit large-weight items in S^*. ◀

To summarize, our above analysis leads to the following lemma:

▶ **Lemma 7.** *With $\mathcal{O}(\delta)$ additive error, we can guess out all the large items in the optimal solution S^*, i.e., all items that either have a large profit or a large weight, by $n^{\tilde{\mathcal{O}}(\frac{1}{\delta^2})}$ enumerations.*

Let $\bar{I} \subseteq I = \{1, 2, \cdots, n\}$ be the set of small items, i.e., items of medium/small-profit and small-weight. Then $I \setminus \bar{I}$ is the set of large items. Denote by $\mathbf{x}^*$ the optimal solution to $\mathbf{IPC}(I, \mathbf{1}, 1)$, which is corresponding to S^*. In the following we assume a correct guess on large items. Hence, the values of $\{x_j^* : j \in I \setminus \bar{I}\}$ are known. We let $\mathbf{a'}$ be the total cost of these guessed-out large items.

3.2.2 Finding the follower's dominant choices on large items

Consider all the large items. Even if the leader's choice on large items is fixed, the follower may still have exponentially many different choices on the remaining large items. The goal of this subsection is to show that, among these choices of the follower, it suffices to restrict our attention to a few "dominant" choices that always outperform other choices.

For simplicity, we re-index items such that $\bar{I} = \{1, 2, \cdots, \bar{n}\}$, where $\bar{n} \le n$.

We further assume that items in $\bar{I}$ are sorted in decreasing order of the profit-weight ratios p_j/B_j. For any $\bar{\mathbf{a}} \le \mathbf{1}$ and $\bar{b} \le 1$, denote by $\mathbf{IPC}(\bar{I}, \bar{\mathbf{a}}, \bar{b})$ the "residual instance" where the item set is $\bar{I}$, the budget vector of the leader is $\bar{\mathbf{a}}$ and the budget of the follower is $\bar{b}$.

Denote by $I' \subseteq I \setminus \bar{I}$ the subset of large items which are *not* selected by the leader. Note that due to the assumption $OPT \le 1$ and that we have guessed out correct large items in S^*, the follower cannot select items from I' with total profit larger than 1. Hence, for each integer $k \in [1, 1 + 1/\epsilon]$, we can define the following sub-problem: among items in I', find out a subset of items with minimal total weight such that their total profit is within $[(k-1)\epsilon, k\epsilon)$. Denote by $SP(k)$ this sub-problem and by $KP(k\epsilon)$ its optimal solution, if it exists. We claim that the follower can select at most $\mathcal{O}(1/\delta)$ items from I', thus via $n^{\mathcal{O}(1/\delta)}$ enumerations, we can return $KP(k\epsilon)$ or assert there does not exist a feasible solution to $SP(k)$. The claim is guaranteed by the following two facts: (i). The total profit the follower could obtain from I' is at most 1; (ii). Items in I' either have a large-profit, or a large-weight.

[2] If there are less than $1/\delta$ items in S_h, we can simply guess out all items in $S_h \setminus S^*$ via $n^{\mathcal{O}(\frac{1}{\delta})}$ enumerations.

Let $\Theta = \{KP(k\epsilon) : k \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\}\}$, which contains the follower's $\mathcal{O}(1/\epsilon)$ possible choices on I'. For $\ell \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\}$, we let b_ℓ and P_ℓ be the total weight and the total profit of the items selected by the follower, respectively[3]. Then the follower has a residual budget of $1 - b_\ell$ for items in $\bar{I}$. Recall that the leader has a residual budget vector of $\mathbf{1} - \mathbf{a}'$ for items in $\bar{I}$.

Define $\mathbf{y}[\bar{I}] = (y_1, y_2 \cdots, y_{\bar{n}})$. Recall that by guessing we already know the value of x_j^* for $j \in I \setminus \bar{I}$. Consider the following bilevel program:

$$\textbf{Bi-IP}(I, \mathbf{1}, 1) : \min_{\mathbf{x}} P_\ell + \sum_{j=1}^{\bar{n}} p_j y_j$$

$$s.t. \sum_{j=1}^{\bar{n}} \mathbf{A}_j x_j \leq \mathbf{1} - \mathbf{a}' \tag{2a}$$

$$x_j = x_j^*, \quad \forall j \in I \setminus \bar{I} \tag{2b}$$

$$x_j \in \{0, 1\}, \quad \forall j \in \bar{I} \tag{2c}$$

where integer $\ell, \mathbf{y}[\bar{I}]$ solves the following:

$$\max_{1 \leq \ell \leq 1 + \frac{1}{\epsilon}} \max_{\mathbf{y}[\bar{I}]} \quad P_\ell + \sum_{j=1}^{\bar{n}} p_j y_j \tag{2d}$$

$$s.t. \qquad \sum_{j=1}^{\bar{n}} B_j y_j \leq 1 - b_\ell \tag{2e}$$

$$y_j \leq 1 - x_j, \quad \forall j \in \bar{I} \tag{2f}$$

$$y_j \in \{0, 1\}, \quad \forall j \in \bar{I} \tag{2g}$$

What is the difference between Bi-IP$(I, \mathbf{1}, 1)$ and IPC$(I, \mathbf{1}, 1)$, assuming the correct guess of x_j^* for $j \in I \setminus \bar{I}$? In **Bi-IP**$(I, \mathbf{1}, 1)$, the follower's choices on remaining large items are restricted to the $\mathcal{O}(1/\epsilon)$ choices in Θ, while in **IPC**$(I, \mathbf{1}, 1)$, the follower can choose any remaining large items. However, we observe that Θ contains all the follower's "dominant choices of remaining large items" in the sense that the follower uses the smallest budget to achieve a profit within $[(k-1)\epsilon, k\epsilon)$. Consequently, the objective value of **Bi-IP**$(I, \mathbf{1}, 1)$ differs by at most ϵ to that of **IPC**$(I, \mathbf{1}, 1)$. A formal description is given below.

▶ **Lemma 8.** *Let $\bar{\mathbf{x}}$ be any feasible solution to **Bi-IP**$(I, \mathbf{1}, 1)$. Then $\bar{\mathbf{x}}$ is also feasible to **IPC**$(I, \mathbf{1}, 1)$. Let $Obj_{Bi}(\bar{\mathbf{x}})$ and $Obj(\bar{\mathbf{x}})$ be the objective values of **Bi-IP**$(I, \mathbf{1}, 1)$ and **IPC**$(I, \mathbf{1}, 1)$ for $\mathbf{x} = \bar{\mathbf{x}}$, respectively. We have*

$$Obj_{Bi}(\bar{\mathbf{x}}) \leq Obj(\bar{\mathbf{x}}) \leq Obj_{Bi}(\bar{\mathbf{x}}) + \epsilon.$$

*Furthermore, let OPT_{Bi} and OPT be the optimal objective values of **Bi-IP**$(I, \mathbf{1}, 1)$ and **IPC**$(I, \mathbf{1}, 1)$, respectively, then we have*

$$OPT_{Bi} \leq OPT \leq OPT_{Bi} + \epsilon.$$

Proof. Compare the follower's possible choices in **Bi-IP**$(I, \mathbf{1}, 1)$ and **IPC**$(I, \mathbf{1}, 1)$ when the leader's solution is fixed to $\bar{\mathbf{x}}$. It is easy to see that in **IPC**$(I, \mathbf{1}, 1)$, the follower's feasible choices on the remaining large items contain Θ, it thus follows that $Obj_{Bi}(\bar{\mathbf{x}}) \leq Obj(\bar{\mathbf{x}})$. Particularly, since the optimal solution $\mathbf{x}^*$ of **IPC**$(I, \mathbf{1}, 1)$ is a feasible solution of **Bi-IP**$(I, \mathbf{1}, 1)$ and the optimal solution of **Bi-IP**$(I, \mathbf{1}, 1)$ may achieve an even smaller value, it follows that $OPT_{Bi} \leq OPT$. It remains to prove that $Obj(\bar{\mathbf{x}}) \leq Obj_{Bi}(\bar{\mathbf{x}}) + \epsilon$ and $OPT \leq OPT_{Bi} + \epsilon$.

[3] If there is no feasible solution to $SP(\ell)$, we let $b_\ell = 1$ and $P_\ell = 0$.

Note that $Obj(\bar{\mathbf{x}})$ is exactly the optimal objective value of the following integer program:

$$
\begin{aligned}
\mathbf{IP}(\bar{\mathbf{x}}) : \max_{\mathbf{y}} \quad & \mathbf{py} \\
s.t. \quad & \sum_{j=1}^{n} B_j y_j \leq 1 \\
& \mathbf{y} \leq \mathbf{1} - \bar{\mathbf{x}} \\
& \mathbf{y} \in \{0,1\}^n
\end{aligned}
$$

Let $\bar{\mathbf{y}}$ be an optimal solution of $\mathbf{IP}(\bar{\mathbf{x}})$, then $Obj(\bar{\mathbf{x}}) = \sum_{j\in I\setminus\bar{I}} p_j\bar{y}_j + \sum_{j\in\bar{I}} p_j\bar{y}_j$. Recall that $\mathbf{x}^*$ is an optimal solution of $\mathbf{IPC}(I,\mathbf{1},1)$, and $\bar{\mathbf{x}}_j = \mathbf{x}^*_j$ for $j \in I\setminus\bar{I}$ by (2b). Consequently, if we compare the follower in $\mathbf{IPC}(I,\mathbf{1},1)$ and the follower in **Bi-IP**$(I,\mathbf{1},1)$, the subset of items in $I\setminus\bar{I}$ available for the two followers to select is the same, and we let this subset be $R = \{j : x^*_j = 0, j \in I\setminus\bar{I}\}$. Given that we assume $OPT \leq 1$, the maximal profit the follower could obtain from R is at most 1, thus there exists some integer $\bar{\ell} \in [1, 1+1/\epsilon]$ such that $\sum_{j\in I\setminus\bar{I}} p_j\bar{y}_j \in [(\bar{\ell}-1)\epsilon, \bar{\ell}\epsilon)$. By the definitions of $P_{\bar{\ell}}$ and $b_{\bar{\ell}}$, we have $\sum_{j\in I\setminus\bar{I}} p_j\bar{y}_j \leq P_{\bar{\ell}} + \epsilon$ and $b_{\bar{\ell}} \leq \sum_{j\in I\setminus\bar{I}} B_j\bar{y}_j$. Define $\mathbf{y}' \in \{0,1\}^n$ such that the $\mathbf{y}'$ is a combination of two partial solutions: in $I\setminus\bar{I}$, $\mathbf{y}'$ is the same as $KP(\bar{\ell}\epsilon)$; and in $\bar{I}$, $\mathbf{y}'$ is the same as $\bar{\mathbf{y}}$. Then $\mathbf{y}'$ is a feasible solution of the following program:

$$
\begin{aligned}
\overline{\mathbf{IP}}(\bar{\mathbf{x}}) : \max_{\ell} \max_{\mathbf{y}[\bar{I}]} \quad & P_\ell + \sum_{j=1}^{\bar{n}} p_j y_j \\
s.t. \quad & \sum_{j=1}^{\bar{n}} B_j y_j \leq 1 - b_\ell \\
& y_j \leq 1 - \bar{x}_j, \ \ \forall j \in \bar{I} \\
& y_j \in \{0,1\}, \ \ \forall j \in \bar{I}
\end{aligned}
$$

Notice that the optimal objective value of $\overline{\mathbf{IP}}(\bar{\mathbf{x}})$ is $Obj_{Bi}(\bar{\mathbf{x}})$, thus $\mathbf{py}' = P_{\bar{\ell}} + \sum_{j=1}^{\bar{n}} p_j\bar{y}_j \leq Obj_{Bi}(\bar{\mathbf{x}})$. To conclude, we have

$$
Obj(\bar{\mathbf{x}}) = \sum_{j\in I\setminus\bar{I}} p_j\bar{y}_j + \sum_{j=1}^{\bar{n}} p_j\bar{y}_j \leq P_{\bar{\ell}} + \epsilon + \sum_{j=1}^{\bar{n}} p_j\bar{y}_j \leq Obj_{Bi}(\bar{\mathbf{x}}) + \epsilon
$$

Particularly, given an optimal solution $\bar{\mathbf{x}}^*$ of **Bi-IP**$(I,\mathbf{1},1)$, we have $Obj(\bar{\mathbf{x}}^*) \leq OPT_{Bi} + \epsilon$. Since the optimal solution of $\mathbf{IPC}(I,\mathbf{1},1)$ may achieve an even smaller objective value, it follows that $OPT \leq OPT_{Bi} + \epsilon$. Hence Lemma 8 is proved. ◀

3.3 Handling Small Items

According to Lemma 8, to solve $\mathbf{IPC}(I,\mathbf{1},1)$, it suffices to solve **Bi-IP**$(I,\mathbf{1},1)$, which is the goal of this subsection. Towards this, we first obtain a linear relaxation of **Bi-IP**$(I,\mathbf{1},1)$ where both the leader and the follower can select items fractionally. Then we reformulate this bilevel linear relaxation as a single level linear program and find an extreme point optimal fractional solution. Finally we round this fractional solution to obtain a feasible solution to **Bi-IP**$(I,\mathbf{1},1)$ with an objective value of at most $1+\mathcal{O}(\epsilon)$, which is thus also a feasible solution to $\mathbf{IPC}(I,\mathbf{1},1)$ with an objective value of at most $1+\mathcal{O}(\epsilon)$.

Replace (2c) and (2g) in $\textbf{Bi-IP}(I,\mathbf{1},1)$ with $x_j \in [0,1](\forall j \in \bar{I})$ and $y_j \in [0,1](\forall j \in \bar{I})$, respectively, we obtain a relaxation of $\textbf{Bi-IP}(I,\mathbf{1},1)$ as follows.

$$\textbf{Bi-IP}_r(I,\mathbf{1},1) : \min_{\mathbf{x}} P_\ell + \sum_{j=1}^{\bar{n}} p_j y_j$$

$$s.t. \sum_{j=1}^{\bar{n}} \mathbf{A}_j x_j \leq \mathbf{1} - \mathbf{a}' \tag{3a}$$

$$x_j = x_j^*, \quad \forall j \in I \setminus \bar{I} \tag{3b}$$

$$x_j \in [0,1], \quad \forall j \in \bar{I} \tag{3c}$$

where integer $\ell, \mathbf{y}[\bar{I}]$ solves the following:

$$\max_{1\leq \ell \leq 1+\frac{1}{\epsilon}} \max_{\mathbf{y}[\bar{I}]} \quad P_\ell + \sum_{j=1}^{\bar{n}} p_j y_j \tag{3d}$$

$$s.t. \qquad \sum_{j=1}^{\bar{n}} B_j y_j \leq 1 - b_\ell \tag{3e}$$

$$y_j \leq 1 - x_j, \quad \forall j \in \bar{I} \tag{3f}$$

$$y_j \in [0,1], \quad \forall j \in \bar{I} \tag{3g}$$

Denote by OPT^r_{Bi} the optimal objective value of $\textbf{Bi-IP}_r(I,\mathbf{1},1)$. Note that items in $\bar{I}$ are sorted in decreasing order of the profit-weight ratios p_j/B_j. Consider any fixed leader's solution $\mathbf{x} \in [0,1]^n$ and any fixed ℓ, the follower is solving a knapsack problem in the remaining (fractional) items. The maximal objective value of the follower, given $\mathbf{x} \in [0,1]^n$ and ℓ, is obtained by a simple greedy algorithm that selects remaining fractional items in $\bar{I}$ in the natural order of indices (recall that items are re-indexed in non-increasing order of ratios), until the budget $1 - b_\ell$ is exhausted. Note that the greedy algorithm will stop at some (fractional) item when the budget $1 - b_\ell$ is exhausted[4]. We say this item is *critical* and let its index be c_ℓ. Given any fixed $\mathbf{x} \in [0,1]^n$ and ℓ, the maximal objective value of program (3d)-(3g) for $\mathbf{x}$ is

$$P_\ell + \sum_{j=1}^{c_\ell - 1} p_j(1 - x_j) + p_{c_\ell} \frac{1 - b_\ell - \sum_{j=1}^{c_\ell - 1} B_j(1 - x_j)}{B_{c_\ell}}, \tag{4a}$$

where c_ℓ is the critical item given $\mathbf{x}$ and ℓ. The following two formulas are directly given by the definition of critical.

$$\sum_{j=1}^{c_\ell - 1} B_j(1 - x_j) \leq 1 - b_\ell \tag{5a}$$

$$B_{c_\ell} + \sum_{j=1}^{c_\ell - 1} B_j(1 - x_j) \geq 1 - b_\ell \tag{5b}$$

We first show that the optimal objective value of the $\textbf{Bi-IP}_r(I,\mathbf{1},1)$ is at most $OPT_{Bi} + \delta$.

[4] The greedy algorithm may pack all remaining (fractional) items without using up the budget $1 - b_\ell$. To patch this case, we add a dummy item whose profit is 0, cost vector is $\mathbf{0}$ and weight is sufficiently large. We assume the last item $\bar{n}$ is the dummy item.

▶ **Lemma 9.** *Let OPT_{Bi} and OPT^r_{Bi} be the optimal objective values of **Bi-IP**$(I, \mathbf{1}, 1)$ and **Bi-IP**$_r(I, \mathbf{1}, 1)$, respectively, then $OPT^r_{Bi} \leq OPT_{Bi} + \delta$.*

Proof. It is not straightforward if we compare **Bi-IP**$(I, \mathbf{1}, 1)$ with **Bi-IP**$_r(I, \mathbf{1}, 1)$ directly, as both the follower and the leader become stronger in the relaxation (in the sense they can pack items fractionally). Towards this, we introduce an intermediate bilevel program **Bi-IP**$_{in}(I, \mathbf{1}, 1)$, which is obtained by replacing (3c) in **Bi-IP**$_r(I, \mathbf{1}, 1)$ with $x_j \in \{0, 1\}(\forall j \in \bar{I})$, that is, we only allow the follower to select items fractionally but not the leader. Denote by OPT_{in} the optimal objective value of **Bi-IP**$_{in}(I, \mathbf{1}, 1)$.

First, we compare **Bi-IP**$_{in}(I, \mathbf{1}, 1)$ with **Bi-IP**$_r(I, \mathbf{1}, 1)$. We see that in **Bi-IP**$_r(I, \mathbf{1}, 1)$ the follower is facing a stronger leader who is allowed to fractionally pack items, and it thus follows that $OPT^r_{Bi} \leq OPT_{in}$.

Next, we compare **Bi-IP**$_{in}(I, \mathbf{1}, 1)$ and **Bi-IP**$(I, \mathbf{1}, 1)$. Note that the leader's solution must be integral in both programs. Any feasible solution of **Bi-IP**$_{in}(I, \mathbf{1}, 1)$ is a feasible solution of **Bi-IP**$(I, \mathbf{1}, 1)$, and vice versa. Let $\bar{\mathbf{x}}^*$ be an optimal solution of **Bi-IP**$(I, \mathbf{1}, 1)$, then $\bar{\mathbf{x}}^*$ is also a feasible solution of **Bi-IP**$_{in}(I, \mathbf{1}, 1)$. Once the leader fixes his solution as $\bar{\mathbf{x}}^*$ in **Bi-IP**$_{in}(I, \mathbf{1}, 1)$, there exist ℓ and $c_\ell \in \bar{I}$, such that the objective value of **Bi-IP**$_{in}(I, \mathbf{1}, 1)$ is

$$Obj_{in} = P_\ell + \sum_{j=1}^{c_\ell - 1} p_j(1 - \bar{x}^*_j) + p_{c_\ell} \frac{1 - b_\ell - \sum_{j=1}^{c_\ell - 1} B_j(1 - \bar{x}^*_j)}{B_{c_\ell}},$$

where c_ℓ is the critical item corresponding to $\bar{\mathbf{x}}^*$ and ℓ. We have the following two observations:

- $OPT_{Bi} \geq Obj_{in} - p_{c_\ell} \geq Obj_{in} - \delta$. This is because the follower in **Bi-IP**$(I, \mathbf{1}, 1)$ can guarantee an objective value of $P_\ell + \sum_{j=1}^{c_\ell - 1} p_j(1 - \bar{x}^*_j)$, and $p_j \leq \delta$ for $j \in \bar{I}$;
- $OPT_{in} \leq Obj_{in}$. This is because $\bar{\mathbf{x}}^*$ is just a feasible solution of **Bi-IP**$_{in}(I, \mathbf{1}, 1)$, while the optimal solution of the leader may achieve an even smaller objective value.

To summarize, we know $OPT^r_{Bi} \leq OPT_{in} \leq Obj_{in} \leq OPT_{Bi} + \delta$. Lemma 9 is proved. ◀

Given Lemma 9, we are still facing two questions: how can we solve **Bi-IP**$_r(I, \mathbf{1}, 1)$; and even if we obtain a fractional solution to **Bi-IP**$_r(I, \mathbf{1}, 1)$, how can we transform it to an integral solution without incurring a huge loss. Towards this, consider the optimal solution $\mathbf{x}^r$ to **Bi-IP**$_r(I, \mathbf{1}, 1)$. Note that leader's choice on large items is guessed out in **Bi-IP**$_r(I, \mathbf{1}, 1)$. Consider the scenario when the follower adopts the ℓ-th dominant choice on the remaining large items, and recall the definition of critical items (see Equation 4a). Given solution $\mathbf{x}^r$, for any $\ell \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\}$ there must exist a critical item. Therefore, there are $1 + \frac{1}{\epsilon}$ critical items corresponding to $\mathbf{x}^r$. The crucial fact is that, while we cannot guess out $\mathbf{x}^r$ directly, we can guess out all the critical items corresponding to $\mathbf{x}^r$. More precisely, with $\bar{n}^{\mathcal{O}(1/\epsilon)}$ enumerations, we can guess out the critical item c^r_ℓ for $\mathbf{x}^r$ and each ℓ. Suppose we have guessed out the correct c^r_ℓ's corresponding to the optimal solution $\mathbf{x}^r$, we consider the following LP:

$$
\begin{aligned}
\mathbf{LP}_{\text{Bi-IP}}: \quad & \min_{\mathbf{x},M} \; M \\
& \sum_{j=1}^{\bar{n}} \mathbf{A}_j x_j \leq \mathbf{1} - \mathbf{a}' \\
& P_\ell + \sum_{j=1}^{c_\ell^r - 1} p_j(1-x_j) + p_{c_\ell^r} \frac{1 - b_\ell - \sum_{j=1}^{c_\ell^r-1} B_j(1-x_j)}{B_{c_\ell^r}} \leq M, \quad \forall \; \ell \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\} \\
& \sum_{j=1}^{c_\ell^r - 1} B_j(1-x_j) \leq 1 - b_\ell, \quad \forall \;\; \ell \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\} \\
& B_{c_\ell^r} + \sum_{j=1}^{c_\ell^r - 1} B_j(1-x_j) \geq 1 - b_\ell, \quad \forall \;\; \ell \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\} \\
& x_j = x_j^*, \quad j \in I \setminus \bar{I} \\
& x_j \in [0,1], \quad j \in \bar{I}
\end{aligned}
$$

We have the following simple observation.

▶ **Observation 10.** *Let M^* and OPT_{Bi}^r be the optimal objective values of $\boldsymbol{LP}_{Bi\text{-}IP}$ and $\boldsymbol{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$, respectively, then $M^* \leq OPT_{Bi}^r$.*

Let $\mathbf{x}^r$ be an optimal solution to $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$. The observation follows directly as $\mathbf{x}^r$ together with OPT_{Bi}^r form a feasible solution to $\mathbf{LP}_{\text{Bi-IP}}$.

In the meantime, we also have the following observation.

▶ **Observation 11.** *Let $\{\mathbf{x}^{ex}, M^*\}$ be an extreme point optimal solution to $\boldsymbol{LP}_{Bi\text{-}IP}$, then $\mathbf{x}^{ex}$ is also a feasible solution to $\boldsymbol{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$ whose objective value is at most M^*.*

The observation follows since by the definition of critical, (4a) is the largest profit the follower can achieve. Hence, when $\mathbf{x} = \mathbf{x}^{ex}$ in $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$, the objective value is bounded by M^*. Given the two observations above, we know $\mathbf{x}^{ex}$ is an optimal solution to $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$ and we have $M^* = OPT_{Bi}^r$. Finally, a near-optimal solution to $\mathbf{IPC}(I, \mathbf{1}, 1)$ can be obtained through the optimal solution to $\mathbf{LP}_{\text{Bi-IP}}$, as implied by the following lemma.

▶ **Lemma 12.** *Let $\{\mathbf{x}^{ex}, M^*\}$ be an extreme point optimal solution to $\boldsymbol{LP}_{Bi\text{-}IP}$. Define $\tilde{\mathbf{x}}$ such that $\tilde{x}_j = 1$ if $x_j^{ex} = 1$, and $\tilde{x}_j = 0$ otherwise. Then $\tilde{\mathbf{x}}$ is a feasible solution of $\boldsymbol{IPC}(I, \mathbf{1}, 1)$ with an objective value of at most $OPT + \mathcal{O}(\epsilon)$, where OPT is the optimal objective value of $\boldsymbol{IPC}(I, \mathbf{1}, 1)$.*

Proof. The feasibility of $\tilde{\mathbf{x}}$ to $\mathbf{IPC}(I, \mathbf{1}, 1)$ is straightforward since

$$\sum_{j=1}^{n} \mathbf{A}_j \tilde{x}_j = \sum_{j \in I \setminus \bar{I}} \mathbf{A}_j x_j^* + \sum_{j \in \bar{I}} \mathbf{A}_j \tilde{x}_j \leq \mathbf{a}' + \sum_{j \in \bar{I}} \mathbf{A}_j x_j^{ex} \leq \mathbf{a}' + \mathbf{1} - \mathbf{a}' \leq \mathbf{1}.$$

Notice that $\tilde{\mathbf{x}}$ is also feasible for $\mathbf{Bi\text{-}IP}(I, \mathbf{1}, 1)$ and $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$. Let $Obj(\tilde{\mathbf{x}})$, $Obj_{Bi}(\tilde{\mathbf{x}})$ and $Obj_{Bi}^r(\tilde{\mathbf{x}})$ be the objective values of $\mathbf{IPC}(I, \mathbf{1}, 1)$, $\mathbf{Bi\text{-}IP}(I, \mathbf{1}, 1)$ and $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$ by taking $\mathbf{x} = \tilde{\mathbf{x}}$, respectively. Since in $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$, the leader is facing a stronger follower who can select fractional items in $\bar{I}$, it follows that $Obj_{Bi}(\tilde{\mathbf{x}}) \leq Obj_{Bi}^r(\tilde{\mathbf{x}})$.

Now we compare the objective values of two solutions to $\mathbf{Bi\text{-}IP}_r(I, \mathbf{1}, 1)$, $\mathbf{x}^{ex}$ and $\tilde{\mathbf{x}}$. It is easy to see that in $\mathbf{x}^{ex}$, there are at most $(s_A + \frac{3(1+\epsilon)}{\epsilon})$ variables taking fractional values, and all these variables are in $\{x_j^{ex} : j \in \bar{I}\}$. So the leader in $\mathbf{x}^{ex}$ selects at most $(s_A + \frac{3(1+\epsilon)}{\epsilon})$

more items in $\bar{I}$, compared with the leader in $\tilde{\mathbf{x}}$. Consequently, the follower in $\mathbf{x}^{ex}$ may select at most $(s_A + \frac{3(1+\epsilon)}{\epsilon})$ less items in $\bar{I}$, compared with the follower in $\tilde{\mathbf{x}}$. Given that the objective value of $\mathbf{x}^{ex}$ is $OPT^r_{Bi} = M^*$, we have that

$$Obj^r_{Bi}(\tilde{\mathbf{x}}) \leq OPT^r_{Bi} + (s_A + \frac{3(1+\epsilon)}{\epsilon})\delta.$$

According to Lemma 8 and Lemma 9, we have $Obj(\tilde{\mathbf{x}}) \leq Obj_{Bi}(\tilde{\mathbf{x}}) + \epsilon$ and $OPT^r_{Bi} \leq OPT_{Bi} + \delta$. In conclusion, we have

$$Obj(\tilde{\mathbf{x}}) \leq OPT_{Bi} + (s_A + 1 + \frac{3(1+\epsilon)}{\epsilon})\delta + \epsilon.$$

Furthermore, $OPT_{Bi} \leq OPT$ by Lemma 8. By setting $\delta = \epsilon^2$, Lemma 12 is proved. ◀

Hence, if $OPT \leq 1$, then a feasible solution with objective value of at most $OPT + \mathcal{O}(\epsilon) \leq 1 + \mathcal{O}(\epsilon)$ is found, thus Lemma 3 is proved, and Theorem 2 follows.

4 Approximation algorithm for IPC where s_B and s_A are constant

In this section, we prove our main result – Theorem 13.

▶ **Theorem 13.** *When s_B and s_A are fixed constants, for an arbitrarily small number $\epsilon > 0$, there exists an $(s_B + \mathcal{O}(\epsilon))$-approximation polynomial time algorithm for IPC.*

Similar to the special case, by scaling item profits it suffices to show the following:

▶ **Lemma 14.** *Let OPT be the optimal objective value of **IPC**$(I, \mathbf{a}, \mathbf{b})$. If $OPT \leq 1$, then for an arbitrarily small number $\epsilon > 0$, there exists a polynomial time algorithm that returns a feasible solution to **IPC**$(I, \mathbf{a}, \mathbf{b})$ with an objective value of at most $s_B + \mathcal{O}(\epsilon)$.*

By further scaling the cost vectors and weight vectors, it suffices to find a near-optimal solution for **IPC**$(I, \mathbf{1}, \mathbf{1})$.

Major technical challenge. Recall that the key to solving IPC for the special case of $s_B = 1$ is the establishment of **Bi-IP**$(I, \mathbf{1}, 1)$, which is essentially equivalent to **IPC**$(I, \mathbf{1}, 1)$. **Bi-IP**$(I, \mathbf{1}, 1)$ is built upon the observation that the follower admits only $\mathcal{O}(\frac{1}{\epsilon})$ dominant choices on large items, where each dominant choice corresponds to the minimal budget needed by the follower to ensure a profit of $[(k-1)\epsilon, k\epsilon)$ where $k \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\}$. Because the number of follower's choices on large items is small, its relaxation **Bi-IP**$_r(I, \mathbf{1}, 1)$ has a small number of constraints (see Equation 3e), and therefore we can further transform **Bi-IP**$_r(I, \mathbf{1}, 1)$ to **LP**$_{\text{Bi-IP}}$ with a small number of constraints whose extreme point solution promises a good rounding. We aim to follow a similar method, however, when $s_B \geq 2$, we can no longer bound the follower's choices on large items. This is because to achieve a profit of $[(k-1)\epsilon, k\epsilon)$ where $k \in \{1, 2, \cdots, 1 + \frac{1}{\epsilon}\}$, the follower may have a huge number of different choices utilizing different budgets, where the budgets are now vectors instead of numbers, and are thus incomparable. To handle the problem, we use the idea of rounding: let $\mathbf{b} = (\mathbf{b}[1], \mathbf{b}[2], \cdots, \mathbf{b}[s_B])$ and $\mathbf{B}_j = (\mathbf{B}_j[1], \mathbf{B}[2], \cdots, \mathbf{B}[s_B])$. We call the first dimension ($\mathbf{b}[1]$ and $\mathbf{B}_j[1]$'s) the principal dimension. The principal dimension will be treated the same as the special case and will *not* be rounded. The coordinates of other dimensions ($\mathbf{b}[h]$ and $\mathbf{B}_j[h]$'s for $2 \leq h \leq s_B$) will be rounded. Then, we will be able to compare follower's different choices on large items: if there are two choices both achieving profit within $[(k-1)\epsilon, k\epsilon)$ for the same $k \in \mathcal{O}(\frac{1}{\epsilon})$, and furthermore, the summation of their weight vectors share the same rounded value in each dimension $h \in [2, s_B]$, then the choice with smaller value in the principal dimension of the summed weight vectors dominates the other choice. By doing so, our argument for the special case can be carried over to the principal dimension.

There is one problem with the idea of rounding above, that is, if we round up weight vectors and meanwhile enlarge the follower's budget vector in dimension $h \in [2, s_B]$ (to accommodate the rounding up), the optimal objective value of IPC may increase. However, we are able to show that the optimal objective value only increases by a factor of s_B (see Lemma 15). This explains our approximation ratio of $s_B + \epsilon$.

Below we give a very brief walk through and the reader is referred to the full version [7] for details.

Step 1. We pick a small parameter δ as the rounding precision, keep the coordinates of weight vectors on principal dimension intact, and round the coordinates on other dimensions. We also round the profits. By doing so we obtain a rounded instance $\tilde{I}_\delta$. Then we pick another small parameter τ and enlarge the weight budget on dimension $h \in [2, s_B]$ by a factor of $1 + \tau$. By doing so we obtain:

$$\begin{aligned}
\mathbf{IPC}_\tau(\tilde{I}_\delta, \mathbf{1}, \mathbf{1}) : \min_{\mathbf{x}} \quad & \tilde{\mathbf{p}}\mathbf{y} \\
s.t. \quad & \mathbf{A}\mathbf{x} \le \mathbf{1} \\
& \mathbf{x} \in \{0,1\}^n \\
& \text{where } \mathbf{y} \text{ solves the following:} \\
& \max_{\mathbf{y}} \quad \tilde{\mathbf{p}}\mathbf{y} \\
& s.t. \quad \sum_{j=1}^{n} \tilde{\mathbf{B}}_j[1] y_j \le 1 \\
& \qquad\ \sum_{j=1}^{n} \tilde{\mathbf{B}}_j[i] y_j \le 1 + \tau, \quad \forall 2 \le i \le s_B \\
& \qquad\ \mathbf{x} + \mathbf{y} \le \mathbf{1} \\
& \qquad\ \mathbf{y} \in \{0,1\}^n
\end{aligned}$$

where $\tilde{\mathbf{B}}_j[1] = \mathbf{B}_j[1]$, $\tilde{\mathbf{B}}_j[h]$, $2 \le h \le s_B$ and $\tilde{\mathbf{p}}$ are rounded weights and profits. We are able to prove the following lemma which ensures that solving $\mathbf{IPC}_\tau(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ gives a good approximate solution to $\mathbf{IPC}(I, \mathbf{1}, \mathbf{1})$:

▶ **Lemma 15.** *Let $0 < \tau \le 1/2$. Let $\tilde{\mathbf{x}}$ be any feasible solution to $\mathbf{IPC}_\tau(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$. Then $\tilde{\mathbf{x}}$ is a feasible solution to $\mathbf{IPC}(I, \mathbf{1}, \mathbf{1})$. Let $\widetilde{Obj}_\tau(\tilde{\mathbf{x}})$ and $Obj(\tilde{\mathbf{x}})$ be the objective values of $\mathbf{IPC}_\tau(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ and $\mathbf{IPC}(I, \mathbf{1}, \mathbf{1})$ for $\mathbf{x} = \tilde{\mathbf{x}}$, respectively. If $2\delta \le \tau \le 1/2$, we have*

$$Obj(\tilde{\mathbf{x}}) \le (1+\delta)\widetilde{Obj}_\tau(\tilde{\mathbf{x}}) \le s_B(1+\delta)Obj(\tilde{\mathbf{x}}).$$

Furthermore, let $\widetilde{OPT}_\tau$ and OPT be the optimal objective values of $\mathbf{IPC}_\tau(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ and $\mathbf{IPC}(I, \mathbf{1}, \mathbf{1})$, respectively. We have

$$OPT \le (1+\delta)\widetilde{OPT}_\tau \le s_B(1+\delta)OPT.$$

Step 2. We handle large items. We first classify item profits into large, medium and small. We then classify item weights into large and small based on the largest coordinate in the weight vector, i.e., $\|\mathbf{B}_j\|_\infty$. We say an item is large if it has a large weight or a large profit, and small otherwise. Let S^* be the leader's optimal solution in $\mathbf{IPC}_\tau(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$. Using a similar argument as the special case, we can prove the following.

▶ **Lemma 16.** *With $\mathcal{O}(s_B\delta)$ additive error, we can guess out all items in S^* that have a large weight or a large profit by $n^{\tilde{\mathcal{O}}(s_B/\delta^{s_B+1})}$ enumerations.*

Utilizing the fact that coordinates in dimension $h \in [2, s_B]$ are all rounded, we can show that the follower only has a small number (i.e. $\tilde{\mathcal{O}}(s_B/\epsilon^{s_B})$) of dominant choices on large items, denoted as Θ. By restricting the follower's choices to Θ, we can obtain a new bilevel integer programming $\textbf{MBi-IP}(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$. Similar to $\textbf{Bi-IP}(I, \mathbf{1}, \mathbf{1})$ in the special case, $\textbf{MBi-IP}(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ has a small number of constraints.

Step 3. We handle small items. Since $\textbf{MBi-IP}(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ has a small number of constraints, we remove the integral constraint to obtain a relaxation $\textbf{MBi-IP}_r(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$. Next, we transform this bilevel LP $\textbf{MBi-IP}_r(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ to a standard (single level) LP, denoted as $\textbf{cen-LP}_\lambda$. Note that here the transformation is much more complicated than that in the special case: in the special case we know that if the follower can choose items fractionally, then its optimal fractional solution is always obtained greedily with respect to the ratio (i.e., profit to weight), whereas it suffices to guess one single critical item. In the general case, if the follower can choose items fractionally, we can only guarantee that among all items whose rounded weight vector are the same except for the principal dimension (i.e., $\mathbf{B}_j[h]$'s have the same rounded value for every $2 \le h \le s_B$), the follower selects items greedily with respect to the principal ratio (i.e., profit to weight coordinate in the principal dimension). Therefore, we need to guess a subset of critical items, and the subscript λ in $\textbf{cen-LP}_\lambda$ corresponds to a set of parameters characterizing the subset of critical items. The most technical part is to show that the optimal solution to $\textbf{cen-LP}_\lambda$ gives a good approximation to $\textbf{MBi-IP}_r(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ (see Lemma 31 in the full version [7]), where we need to create a sequence of "intermediate" LPs. Finally, we obtain an extreme point solution to $\textbf{MBi-IP}_r(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ by solving $\textbf{cen-LP}_\lambda$, and round it to an integral solution. The rounding error can be bounded due to that $\textbf{MBi-IP}_r(\tilde{I}_\delta, \mathbf{1}, \mathbf{1})$ contains a small number of constraints.

5 Approximation algorithm for IPC where s_B and s_A are arbitrary

We consider the most general setting of IPC where s_A and s_B are arbitrary (not necessarily polynomial in the input size).

We define $\max\{\mathbf{py} : \mathbf{By} \le \mathbf{b}, \mathbf{y} \in \{0,1\}^n\}$ as the *follower's problem.* A separation oracle for the *leader's problem* is an oracle such that given any $\mathbf{x} = \mathbf{x}^0 \in [0,1]^n$, it either asserts that $\mathbf{x}^0 \in \{\mathbf{x} : \mathbf{Ax} \le \mathbf{a}, \mathbf{x} \in [0,1]^n\}$, or returns a violating constraint. The goal of this section is to prove the following theorem.

▶ **Theorem 17.** *Given a separation oracle O_L for the leader's problem, and an oracle O_F for the follower's problem that returns a ρ-approximation solution, there exists a $(\frac{\rho}{1-\alpha}, \frac{1}{\alpha})$-bicriteria approximation algorithm for any $\alpha \in (0,1)$ that returns a solution $\mathbf{x}^* \in \{0,1\}^n$ such that $\mathbf{Ax}^* \le \frac{1}{\alpha} \cdot \mathbf{a}$, and*

$$\max\{\mathbf{py} : \mathbf{By} \le \mathbf{b}, \mathbf{y} \le \mathbf{1} - \mathbf{x}^*, \mathbf{y} \in \{0,1\}^n\} \le \frac{\rho T^*}{1-\alpha},$$

where T^* is the optimal objective value of $\textbf{IPC}(I, \mathbf{a}, \mathbf{b})$. Furthermore, the algorithm runs in polynomial oracle time.

We omit the proof of Theorem 17 here, and refer the reader to the full version [7].

6 Conclusions

In this paper, we consider a general two-player zero-sum Stackelberg game in which the leader interdicts some items to minimize the total profit that the follower could obtain from the remaining items. We obtain an $(s_B + \epsilon)$-approximation algorithm when s_A and s_B are both constant, and show that there does not exist any $(4/3 - \epsilon)$-approximation algorithm when $s_B \geq 2$. Our algorithm is the best possible when $s_B = 1$, however, it is not clear whether it is the best possible when s_B is larger than or equal to 2. In particular, it is not clear whether an approximation algorithm with a ratio independent of s_B can be obtained. Furthermore, can we hope for a PTAS if $s_B \geq 2$ but the constraints of the leader or the follower are not given by general inequalities but follow from common optimization problems? For example, what if the follower's optimization problem is a bin packing problem? It would be interesting to investigate the bilevel generalization of well-known optimization problems, e.g., scheduling and bin packing.

References

1 Bo An, Fernando Ordóñez, Milind Tambe, Eric Shieh, Rong Yang, Craig Baldwin, Joseph DiRenzo III, Kathryn Moretti, Ben Maule, and Garrett Meyer. A deployed quantal response-based patrol planning system for the us coast guard. *Interfaces*, 43(5):400–420, 2013.

2 Luce Brotcorne, Saïd Hanafi, and Raïd Mansi. One-level reformulation of the bilevel knapsack problem using dynamic programming. *Discrete Optimization*, 10(1):1–10, 2013.

3 Carl Burch, Robert Carr, Sven Krumke, Madhav Marathe, Cynthia Phillips, and Eric Sundberg. A decomposition-based pseudoapproximation algorithm for network flow inhibition. In *Network Interdiction and Stochastic Integer Programming*, chapter 3, pages 51–68. Springer, 2003.

4 Alberto Caprara, Margarida Carvalho, Andrea Lodi, and Gerhard J. Woeginger. A complexity and approximability study of the bilevel knapsack problem. In *Proceedings of the 16th International Conference on Integer Programming and Combinatorial Optimization*, pages 98–109, 2013.

5 Alberto Caprara, Hans Kellerer, Ulrich Pferschy, and David Pisinger. Approximation algorithms for knapsack problems with cardinality constraints. *European Journal of Operational Research*, 123(2):333–345, 2000.

6 Margarida Carvalho, Andrea Lodi, and Patrice Marcotte. A polynomial algorithm for a continuous bilevel knapsack problem. *Operations Research Letters*, 46(2):185–188, 2018.

7 Lin Chen, Xiaoyu Wu, and Guochuan Zhang. Approximation algorithms for interdiction problem with packing constraints. *arXiv*, 2022.

8 Lin Chen and Guochuan Zhang. Approximation algorithms for a bi-level knapsack problem. *Theoretical Computer Science*, 497:1–12, 2013.

9 Stephen R Chestnut and Rico Zenklusen. Hardness and approximation for network flow interdiction. *Networks*, 69(4):378–387, 2017.

10 Stephen R. Chestnut and Rico Zenklusen. Interdicting structured combinatorial optimization problems with {0, 1}-objectives. *Mathematics of Operations Research*, 42(1):144–166, 2017.

11 Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153(1):235–256, 2007.

12 Federico Della Croce and Rosario Scatamacchia. An exact approach for the bilevel knapsack problem with interdiction constraints and extensions. *Mathematical Programming*, 183(1):249–281, 2020.

13 Stephan Dempe and Klaus Richter. Bilevel programming with knapsack constraints. *Central European Journal of Operations Research*, 8(2):93–108, 2000.

14 Scott DeNegre. *Interdiction and discrete bilevel linear programming.* PhD thesis, Lehigh University, 2011.

15 Michael Dinitz and Anupam Gupta. Packing interdiction and partial covering problems. In *Proceedings of the 16th international conference on Integer Programming and Combinatorial Optimization*, pages 157–168, 2013.

16 Avinash Dixit. The role of investment in entry-deterrence. *The economic journal*, 90(357):95–106, 1980.

17 Dennis Fischer and Gerhard J. Woeginger. A faster algorithm for the continuous bilevel knapsack problem. *Operations Research Letters*, 48(6):784–786, 2020.

18 Matteo Fischetti, Ivana Ljubic, Michele Monaci, and Markus Sinnl. A new general-purpose algorithm for mixed-integer bilevel linear programs. *Operations Research*, 65(6):1615–1637, 2017.

19 Matteo Fischetti, Ivana Ljubic, Michele Monaci, and Markus Sinnl. Interdiction games and monotonicity, with application to knapsack problems. *INFORMS Journal on Computing*, 31(2):390–410, 2019.

20 Matteo Fischetti, Michele Monaci, and Markus Sinnl. A dynamic reformulation heuristic for generalized interdiction problems. *European Journal of Operational Research*, 267(1):40–51, 2018.

21 Klaus Jansen. An eptas for scheduling jobs on uniform processors: using an milp relaxation with a constant number of integral variables. *SIAM Journal on Discrete Mathematics*, 24(2):457–485, 2010.

22 Klaus Jansen. Parameterized approximation scheme for the multiple knapsack problem. *SIAM Journal on Computing*, 39(4):1392–1412, 2010.

23 Robert G. Jeroslow. The polynomial hierarchy and a simple model for competitive analysis. *Mathematical Programming*, 32(2):146–164, 1985.

24 Shiva Prasad Kasiviswanathan and Feng Pan. Matrix interdiction problem. In *Proceedings of the 7th international conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pages 219–231, 2010.

25 Thomas Kleinert, Martine Labbé, Ivana Ljubic, and Martin Schmidt. A survey on mixed-integer programming techniques in bilevel optimization. *EURO Journal on Computational Optimization*, 2021.

26 André Linhares and Chaitanya Swamy. Improved algorithms for mst and metric-tsp interdiction. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming*, pages 32:1–32:14, 2017.

27 D. johnson) M. Garey. Computers and intractability: a guide to the theory of np-completeness. *Siam Review*, 24(1):90, 1982.

28 Feng Pan and Aaron Schild. Interdiction problems on planar graphs. *Discrete Applied Mathematics*, 198:215–231, 2016.

29 Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Efficient algorithms to solve bayesian stackelberg games for security applications. In *Proceedings of the 23th AAAI Conference on Artificial Intelligence*, pages 1559–1562, 2008.

30 Ulrich Pferschy, Gaia Nicosia, and Andrea Pacifici. A stackelberg knapsack game with weight control. *Theoretical Computer Science*, 799:149–159, 2019.

31 Ulrich Pferschy, Gaia Nicosia, Andrea Pacifici, and Joachim Schauer. On the stackelberg knapsack game. *European Journal of Operational Research*, 291(1):18–31, 2021.

32 Cynthia A. Phillips. The network inhibition problem. In *Proceedings of the 25th annual ACM symposium on Theory of computing*, pages 776–785, 1993.

33 James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 125–132, 2008.

34 Xian Qiu and Walter Kern. Improved approximation algorithms for a bilevel knapsack problem. *Theoretical Computer Science*, 595:120–129, 2015.

35 Yen Tang, Jean-Philippe P. Richard, and J. Cole Smith. A class of algorithms for mixed-integer bilevel min-max optimization. *Journal of Global Optimization*, 66(2):225–262, 2016.
36 Rico Zenklusen. Matching interdiction. *Discrete Applied Mathematics*, 158(15):1676–1690, 2010.
37 Rico Zenklusen. Network flow interdiction on planar graphs. *Discrete Applied Mathematics*, 158(13):1441–1455, 2010.
38 Rico Zenklusen. Connectivity interdiction. *Operations Research Letters*, 42(6-7):450–454, 2014.
39 Rico Zenklusen. An o(1)-approximation for minimum spanning tree interdiction. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science*, pages 709–728, 2015.