

Tight Approximation Results for Matroid Optimization with a Linear Constraint

Ilan Doron-Arad   

MIT, Cambridge, MA, USA

Hadas Shachnai   

Technion – Israel Institute of Technology, Haifa, Israel

Gilad Shmerler  

Technion – Israel Institute of Technology, Haifa, Israel

Abstract

We study the following class of *matroid optimization problems with a linear constraint* ($\mathcal{P}$ -MOL). Given a matroid $\mathcal{M} = (E, \mathcal{I})$, two nonnegative weight functions $v, w : E \rightarrow \mathbb{R}_{\geq 0}$, and a threshold $L \in \mathbb{R}_{\geq 0}$, find $\text{opt } v(S)$ where S is either an independent set or a base of $\mathcal{M}$ satisfying a budget-type constraint: $w(S) \leq L$ or $w(S) \geq L$, and $\text{opt} \in \{\min, \max\}$. $\mathcal{P}$ -MOL provides a unified representation for a broad family of NP-hard optimization problems, including BUDGETED MATROID INDEPENDENT SET, CONSTRAINED MINIMUM-BASIS, and KNAPSACK-COVER variants with a matroid constraint. Also, it naturally extends to multiple matroid constraints. In particular, we consider the MATROID INTERSECTION COVER (MIC) problem, where feasibility is defined by the common independent sets of two matroids and one seeks minimum $v(S)$ subject to $w(S) \geq L$.

Our main result is a unified *efficient polynomial-time approximation scheme* (EPTAS) for all nontrivial $\mathcal{P}$ -MOL variants, obtained by generalizing a technique of Hassin and Levin (SIAM J. Comput., 2004) for solving the CONSTRAINED MINIMUM SPANNING TREE problem. Specifically, for any fixed $\varepsilon > 0$, we present an algorithm running in time $|E|^{O(1)} \cdot (1/\varepsilon^2)^{O(1/\varepsilon)}$ that outputs a feasible solution S whose value is at most $(1 + \varepsilon)\text{OPT}$ for minimization variants and at least $(1 - \varepsilon)\text{OPT}$ for maximization variants. This resolves the complexity status of all members of $\mathcal{P}$ -MOL, as none of these problems admits a *fully* polynomial-time approximation scheme (Doron-Arad, Kulik and Shachnai, ICALP’24). Finally, we separate the $\mathcal{P}$ -MOL family from its extension to matroid intersection. We show that an EPTAS is unlikely to exist for the matroid intersection variant of $\mathcal{P}$ -MOL under a covering constraint, whereas an EPTAS is known to exist under a budget constraint. This highlights a qualitative difference between these two types of linear constraints that does not arise in the single-matroid setting.

2012 ACM Subject Classification Theory of computation

Keywords and phrases Matroids, Knapsack, Linear Constraints, EPTAS, Budgeted Optimization, MOL Problems

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.18

Category APPROX

1 Introduction

Matroids offer a unifying abstraction for “independence” in combinatorial structures: they capture linear independence in vector spaces, forests in graphs, and many other systems constrained by an exchange principle. A *matroid* is a set system $(E, \mathcal{I})$, where E is a finite set and $\mathcal{I} \subseteq 2^E$ are the *independent sets* (IS) such that (i) $\emptyset \in \mathcal{I}$, (ii) for all $A \in \mathcal{I}$ and $B \subseteq A$ it holds that $B \in \mathcal{I}$, and (iii) for all $A, B \in \mathcal{I}$ where $|A| > |B|$, there is $e \in A \setminus B$ such that $B \cup \{e\} \in \mathcal{I}$.¹ A broad range of classical problems can be cast as *matroid*

¹ Properties (ii) and (iii) are known, respectively, as *hereditary property*, and *exchange property*.



© Ilan Doron-Arad, Hadas Shachnai, and Gilad Shmerler;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2026).

Editors: Mohit Singh and Tom Gur; Article No. 18; pp. 18:1–18:19



Leibniz International Proceedings in Informatics

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

optimization: one seeks a minimum-cost or maximum-value independent set or basis of a matroid. Canonical examples include minimum spanning trees, maximum-weight forests, and a variety of assignment and scheduling problems.

In many applications, however, matroid feasibility alone is insufficient. Solutions must also satisfy an additional *budget-type* constraint, such as an upper bound on total length or cost, or a lower bound on total coverage. Adding even a single linear constraint of this form typically turns an otherwise polynomial-time matroid optimization problem into an NP-hard one, and has motivated a long line of work on approximation algorithms.

1.1 The $\mathcal{P}$ -MOL Class and Our Framework

We start by defining the class of problems studied in this paper. Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid, and let $v, w : E \rightarrow \mathbb{R}_{\geq 0}$ be two nonnegative functions. The value v_e contributes to the main objective, and w_e contributes to a budget constraint. For the matroid $\mathcal{M} = (E, \mathcal{I})$ we define $\text{IS}(\mathcal{M}) := \mathcal{I}$ and $\text{bases}(\mathcal{M}) := \{S \in \mathcal{I} : |S| = \text{rank}(\mathcal{M})\}$, where $\text{rank}(\mathcal{M}) := \max\{|T| : T \in \mathcal{I}\}$ is the rank of $\mathcal{M}$, i.e., the maximum cardinality of an independent set.

We are interested in three design choices:

- **Objective direction.** We may want to minimize or maximize $v(S)$.²
 - **Matroid family.** We may optimize over independent sets or bases of $\mathcal{M}$.
 - **Budget direction.** The budget constraint may take the form $w(S) \leq L$ or $w(S) \geq L$.
- We encode these choices by three parameters

$$\text{opt} \in \{\min, \max\}, \quad \mathcal{F} \in \{\text{IS}, \text{bases}\}, \quad \triangleright \in \{\leq, \geq\}.$$

Given a threshold $L \geq 0$, we consider the optimization problem

$$\text{opt } v(S) \quad \text{s.t.} \quad S \in \mathcal{F}, \quad w(S) \triangleright L. \quad (1)$$

We refer to (1) as a $\mathcal{P}$ -MOL (Matroid Optimization with a Linear constraint) problem, with the descriptor $\mathcal{P} = (\text{opt}, \mathcal{F}, \triangleright)$. An instance of a $\mathcal{P}$ -MOL problem is specified by a tuple $\mathcal{A}_{\mathcal{P}} = (E, \mathcal{I}, v, w, L)$.

For a fixed descriptor $\mathcal{P}$, we denote the optimal value of (1) by $\text{OPT}_{\mathcal{P}}$, or simply OPT when the context is clear:

$$\text{OPT}_{\mathcal{P}} := \text{opt}\{v(S) : S \in \mathcal{F}(\mathcal{M}), w(S) \triangleright L\},$$

This formulation captures, in a uniform way, several extensively studied problems. For example:

- $(\max, \text{IS}, \leq)$ gives the BUDGETED MATROID INDEPENDENT SET problem, generalizing 0/1-KNAPSACK.
- $(\min, \text{bases}, \leq)$ gives CONSTRAINED MINIMUM-BASIS problems, including CONSTRAINED MINIMUM SPANNING TREE when $\mathcal{M}$ is a graphic matroid.
- $(\min, \text{IS}, \geq)$ corresponds to KNAPSACK-COVER WITH A MATROID.

Among the eight possible choices of $(\text{opt}, \mathcal{F}, \triangleright)$, only $(\min, \text{IS}, \leq)$ is trivial (its optimum is $S = \emptyset$); we focus on the remaining seven nontrivial variants.

The $\mathcal{P}$ -MOL viewpoint also naturally extends beyond a single matroid. If one lets $\mathcal{F}$ denote the intersection of two matroids, then (1) becomes matroid-intersection optimization with a linear constraint. A notable example is the MATROID INTERSECTION COVER (MIC)

² For any vector $p = (p_e)_{e \in E}$ and any $S \subseteq E$ we use the shorthand $p(S) := \sum_{e \in S} p_e$.

problem: given matroids $\mathcal{M}_1 = (E, \mathcal{I}_1)$ and $\mathcal{M}_2 = (E, \mathcal{I}_2)$ and weights $v, w : E \rightarrow \mathbb{R}_{\geq 0}$, find a set $S \in \mathcal{I}_1 \cap \mathcal{I}_2$ with $w(S) \geq L$ minimizing $v(S)$. In our terminology, MIC aligns with a (min, intersection, $\geq$)-MOL formulation, with feasibility governed by $\mathcal{I} = \mathcal{I}_1 \cap \mathcal{I}_2$.

1.2 Our Contribution

We present (in Section 2) a unified approximation framework for $\mathcal{P}$ -MOL problems by generalizing a technique of Hassin and Levin [20]. This allows to obtain an EPTAS for all non-trivial variants of $\mathcal{P}$. Specifically, for every fixed $\varepsilon > 0$ there exists an algorithm which, given an instance $\mathcal{A}_{\mathcal{P}}$, outputs a feasible solution S satisfying

$$\text{opt} = \min \implies v(S) \leq (1 + \varepsilon) \text{OPT}_{\mathcal{P}}, \quad \text{opt} = \max \implies v(S) \geq (1 - \varepsilon) \text{OPT}_{\mathcal{P}}.$$

The running time is $h(1/\varepsilon) \cdot \text{poly}(|\mathcal{A}_{\mathcal{P}}|)$ for an arbitrary function h . This resolves the complexity status of all members of $\mathcal{P}$ -MOL, as none of these problems admits a *fully* polynomial-time approximation scheme (FPTAS) [15]. Table 1 presents the best known results for the $\mathcal{P}$ -MOL class. Results obtained in this paper are marked in boldface.

We also consider the natural extension of $\mathcal{P}$ -MOL variants to matroid intersection. For BUDGETED MATROID INTERSECTION, an EPTAS was presented in [12]. In contrast to our positive result for the $\mathcal{P}$ -MOL class, where matroid optimization is over a *single* matroid constraint, we show that a unified EPTAS is unlikely to exist for *all* $\mathcal{P}$ -MOL variants, already when the single matroid constraint is replaced by a *matroid intersection* constraint. Specifically, we prove (in Section 3) that MIC does not admit an EPTAS unless $\text{W}[1] = \text{FPT}$.

1.3 Techniques

A Unified EPTAS for $\mathcal{P}$ -MOL

For BUDGETED MATROID INDEPENDENT SET, as well as for its extension to BUDGETED MATROID INTERSECTION, recent works [13, 12] give EPTASs based on the *representative-set* methodology. Roughly, it computes a small candidate pool that is guaranteed to contain the high-value elements of some near-optimal solution, and then recovers the near-optimum by exhaustive enumeration over this pool. However, this approach inherently (i) loses the profit of $O(1)$ elements, and (ii) may drop $O(1)$ elements to restore feasibility. Thus, a solution output by the representative set approach may be a non-basis for the given matroid. It may also be infeasible w.r.t. a covering constraint. Therefore, the EPTASs of [13, 12] cannot be extended to EPTASs for the *entire* $\mathcal{P}$ -MOL class, which must handle both $\text{IS}(\mathcal{M})$ and $\text{bases}(\mathcal{M})$, and all four choices of $(\text{opt}, \triangleright)$.

Instead, we generalize (in Section 2) a technique of Hassin and Levin [20] to obtain a unified EPTAS for all members of $\mathcal{P}$ -MOL under a general matroid constraint. The following is a high-level overview of the algorithmic framework.

Given a constant-factor guarantee $\hat{V}$ for OPT, we bucket elements according to their values v_e . Elements with $v_e \leq \varepsilon \hat{V}$ are classified as *small*, while the remaining elements are partitioned into $O(1/\varepsilon^2)$ narrow buckets. We note that any optimal solution contains only $O(1/\varepsilon)$ *large* elements. This allows us to enumerate *configuration vectors* $\vec{k}$ that specify the number of selected elements from each bucket.

For a fixed configuration $\vec{k}$, we enforce these counts by intersecting the matroid with a partition matroid, yielding a matroid-intersection family $\mathcal{F}_{\vec{k}}$. We then use a Lagrangian relaxation of the linear constraint $w(S) \triangleright L$ and compute $\lambda_{\vec{k}}^*$, an optimal value of the multiplier λ .

The main technical challenge is to convert the Lagrangian optimum into a feasible solution. At λ_k^* , the family $\mathcal{O}_k$ of Lagrangian-optimal sets is *well connected*: any two of its members can be transformed into one another via a sequence of exchanges, each replacing at most one element per bucket. We compute sets $S^-, S^+ \in \mathcal{O}_k$ that respectively minimize and maximize the signed violation of $w(S) \triangleright L$, and hence lie on opposite sides of feasibility. Traversing the one-per-bucket exchange path from S^- to S^+ , we select the first feasible set S^* .

Because elements within each bucket have nearly identical v -values, each exchange affects the objective value by at most the total *bucket spread* – that is, the difference between the maximum and minimum v -value within a bucket. Consequently, the first feasible set S^* yields an objective value within an additive $O(\varepsilon \text{OPT})$ of the Lagrangian bound. Taking the best such S^* over all configurations yields an EPTAS.

We note that some variants in our descriptor-based formulation are related by standard reductions. Nevertheless, we keep the unified formulation for completeness and to make the scope of the algorithm explicit, since such reductions may not always preserve the approximation guarantees in a straightforward way. In particular, budget-type and covering-type constraints require careful treatment: the existence of an EPTAS for the budget setting under a matroid intersection constraint [12] does not, by itself, imply an analogous result for the covering setting. Our hardness result in Theorem 13 highlights this distinction. Therefore, we present the argument directly for the descriptor class $\mathcal{P} = (\text{opt}, \mathcal{F}, \triangleright)$.

Hardness of MIC

We establish a parameterized lower bound for MATROID INTERSECTION COVER (MIC). Via a reduction from k -Clique, we show that MIC admits no EPTAS unless $W[1] = \text{FPT}$ (Theorem 13). The novelty of our result lies in the fact that this reduction rules out an *EPTAS*, whereas a similar construction of [15] is used only to exclude an *FPTAS*. Our result complements recent applications of related reductions to 3-MATROID INTERSECTION [16]. Together, our findings demonstrate that for matroid intersection, different linear constraints (cover versus budget) lead to markedly different algorithmic behavior, in contrast to the case of a single matroid.

■ **Table 1** Our positive results for $\mathcal{P}$ -MOL problems and the matching negative results of [15].

Problem Type	Opt	$\mathcal{F}$	$\triangleright$	Positive Result	Negative Result
BUDGETED MATROID INDEPENDENT SET	max	IS	$\leq$	EPTAS [13]	No FPTAS
KNAPSACK-COVER WITH A MATROID	min	IS	$\geq$	EPTAS	No FPTAS
BUDGETED MATROID BASIS	max	Bases	$\leq$	EPTAS	No FPTAS
CONSTRAINED MINIMUM BASIS	min	Bases	$\leq$	EPTAS ³	No FPTAS
LOWER-BOUNDED MAX BASIS	max	Bases	$\geq$	EPTAS	No FPTAS
LOWER-BOUNDED MIN BASIS	min	Bases	$\geq$	EPTAS	No FPTAS
LOWER-BOUNDED INDEP. SET	max	IS	$\geq$	EPTAS	No FPTAS

³ Previously, an EPTAS was known for the special case of CONSTRAINED MINIMUM BASIS in a *graphic* matroid [20].

1.4 Related Work

Knapsack and approximation schemes

The 0/1-KNAPSACK problem is a canonical benchmark for approximation schemes. Early work already established a fully polynomial-time approximation scheme [24], and dynamic-programming-based FPTASs and their implications have since been extensively studied [8]. More recent progress has focused on fine-grained analyses and improved time-approximation tradeoffs [4, 9].

Matroid and matching constraints with a single budget

Adding a matroid or matching feasibility constraint to 0/1-KNAPSACK gives rise to the well-studied problems of BUDGETED MATCHING and BUDGETED MATROID INTERSECTION. A PTAS for both variants was obtained by Berger et al. [1]. Subsequent work strengthened these results to EPTAS guarantees: for BUDGETED MATROID INDEPENDENT SET [13], and for BUDGETED MATCHING and BUDGETED MATROID INTERSECTION via representative sets [12].

While an FPTAS is ruled out for general instances of BUDGETED MATROID INDEPENDENT SET [15], FPTASs are known for several special cases, including CARDINALITY-CONSTRAINED KNAPSACK [5], MULTIPLE-CHOICE KNAPSACK [22], and laminar matroids [14, 27]. Additional variants of the problem have also been studied, for example in [11, 10, 17].

$\mathcal{P}$ -MOL problems have also been studied in the covering setting (as opposed to a budget constraint), in conjunction with a matroid feasibility constraint. For LOWER-BOUNDED MINIMUM INDEPENDENT SET, the best known result prior to this work was a PTAS, due to [6]. For CONSTRAINED MINIMUM SPANNING TREE, a PTAS was given in [26], followed by an EPTAS in [20]. Furthermore, a multi-criteria FPTAS is known for this problem [21].

Multiple budgets

When multiple budget constraints are present, PTASs are known for variants with either matroid intersection or matching constraints, using techniques such as dependent randomized rounding [7] and structural properties of matroid polytopes [19]. On the other hand, even seemingly simple extensions face strong complexity barriers: for example, the MULTIDIMENSIONAL KNAPSACK problem with just two budget constraints already rules out an EPTAS under standard complexity assumptions [23].

2 An EPTAS for the $\mathcal{P}$ -MOL Class

Let $\mathcal{A}_{\mathcal{P}} = (E, \mathcal{I}, v, w, L)$ be an instance of $\mathcal{P}$ -MOL with descriptor $\mathcal{P} = (\text{opt}, \mathcal{F}, \triangleright)$ and fix $\varepsilon > 0$. Recall that $\mathcal{F}$ denotes the underlying combinatorial family: either all independent sets or all bases of $\mathcal{M}$. Our EPTAS combines a value-bucketing scheme with a Lagrangian relaxation applied to the intersection of $\mathcal{M}$ with a partition matroid (defined below). We begin by describing the bucketing and configuration step. Let $n := |E|$ denote the ground-set size.

2.1 Value Bucketing and Configuration Vectors

We first obtain a constant-factor estimate of OPT.

► **Lemma 1.** *Given an instance $\mathcal{A}_{\mathcal{P}}$, one can guess in $|E|^{O(1)}$ time a value $\hat{V}$ such that*

$$\frac{OPT}{2} \leq \hat{V} \leq OPT \quad \text{if } opt = \max, \quad \text{and} \quad OPT \leq \hat{V} \leq 2OPT \quad \text{if } opt = \min.$$

Proof. We obtain $\hat{V}$ by iterating over all powers of 2 in the domain $[0, v(E)]$. In one of these iterations we necessarily hit a value within a factor 2 of OPT , and we output this guess as $\hat{V}$. For simplicity of presentation, in the remainder we assume that we are given such a value $\hat{V}$ satisfying the stated 2-approximation guarantees.

The number of guesses is $O(\log v(E))$. Since $OPT \leq v(E) \leq n v_{\max}$,

$$\log v(E) \leq \log(n v_{\max}) = O(\log n + \log v_{\max}).$$

Let b be the bit-length of $v_{\max}$. If $b = \text{poly}(n)$ we are done; otherwise b is exponential in n and we can solve the instance exactly by enumerating all 2^n subsets of E , which is still polynomial in $|E|$. ◀

We assume w.l.o.g. that $v_e \leq 2\hat{V}$ for all $e \in E$ (any e with $v_e > 2\hat{V}$ cannot belong to an optimal solution, since all values are nonnegative and $OPT \leq 2\hat{V}$ by Lemma 1). We partition the ground set E into bins according to the values v_e . We first define the bin of *small* elements:

$$B_0 := \{e \in E : v_e \leq \varepsilon \hat{V}\}.$$

All other elements are considered *large* and are further categorized into $T = \left\lceil \frac{2-\varepsilon}{\varepsilon^2} \right\rceil$ bins:

$$B_i := \{e \in E : (\varepsilon + (i-1)\varepsilon^2)\hat{V} < v_e \leq (\varepsilon + i\varepsilon^2)\hat{V}\}, \quad i = 1, \dots, T.$$

► **Lemma 2.** *In any optimal solution S^{opt} , at most $2/\varepsilon$ elements belong to $\bigcup_{i=1}^T B_i$.*

Proof. Suppose, for the sake of contradiction, that S^{opt} contains more than $2/\varepsilon$ elements from $B_1 \cup \dots \cup B_T$. By the definition of these bins, every such element has value strictly larger than $\varepsilon \hat{V}$, and hence

$$v(S^{\text{opt}}) \geq v(S^{\text{opt}} \cap (B_1 \cup \dots \cup B_T)) > \frac{2}{\varepsilon} \cdot \varepsilon \hat{V} = 2\hat{V}.$$

Hence $v(S^{\text{opt}}) > 2\hat{V}$. If $opt = \max$, then by Lemma 1 we have $OPT \leq 2\hat{V}$, contradicting $v(S^{\text{opt}}) = OPT > 2\hat{V}$. If $opt = \min$, then Lemma 1 gives $OPT \leq \hat{V}$ and hence $OPT \leq 2\hat{V}$, again contradicting $v(S^{\text{opt}}) = OPT > 2\hat{V}$. Thus S^{opt} contains at most $2/\varepsilon$ elements from $\bigcup_{i=1}^T B_i$. ◀

We therefore restrict our attention to solutions with at most $2/\varepsilon$ large elements. We encode such solutions via a *configuration vector* $\vec{k} = (k_0, k_1, \dots, k_T)$, where k_i denotes the number of elements that should be chosen from bin B_i . We only consider vectors satisfying

$$\sum_{i=1}^T k_i \leq \left\lceil \frac{2}{\varepsilon} \right\rceil, \quad \sum_{i=0}^T k_i \leq |E|. \quad (2)$$

For such a vector $\vec{k}$, a set $S \subseteq E$ respects the configuration $\vec{k}$ if $|S \cap B_i| = k_i$ for all $i = 0, \dots, T$.

► **Lemma 3.** *The number of distinct configuration vectors satisfying (2) is at most $O\left(|E| \cdot \left(\frac{1}{\varepsilon^2}\right)^{2/\varepsilon}\right)$.*

Proof. Let $m := \lceil 2/\varepsilon \rceil$. By (2), the number of valid tuples $(k_1, \dots, k_T)$ corresponds to the number of non-negative integer solutions to the inequality $\sum_{i=1}^T k_i \leq m$. By introducing a slack variable $k^* \geq 0$ such that $k^* + \sum_{i=1}^T k_i = m$, this count is equivalent to the number of weak compositions of m into $T+1$ parts. Applying the stars and bars formula, we obtain exactly

$$\binom{(T+1) + m - 1}{m} = \binom{T+m}{m}$$

configurations. For each fixed tuple, k_0 is bounded by $|E|$. For a fixed ε (and thus fixed m), the number of configurations is

$$|E| \cdot \binom{T+m}{m} \leq |E| \cdot (T+m)^m = O(|E| \cdot T^m) = O\left(|E| \cdot \left(\frac{1}{\varepsilon^2}\right)^{2/\varepsilon}\right). \quad \blacktriangleleft$$

Let $\vec{k}^{\text{opt}}$ denote the configuration vector of an optimal solution S^{opt} . By Lemma 2, $\vec{k}^{\text{opt}}$ satisfies (2), hence it is among the vectors considered by our algorithm.

2.2 Partition Matroid and Restricted Family

Fix a configuration vector $\vec{k} = (k_0, \dots, k_T)$ satisfying (2), and let $K := \sum_{i=0}^T k_i$. We can enforce the cardinality of $\mathcal{M}$ by considering its truncation to rank K , denoted $\mathcal{M}_{\leq K}$, whose independent sets are $\mathcal{I}_{\leq K} = \text{IS}(\mathcal{M}_{\leq K}) = \{S \in \mathcal{I} : |S| \leq K\}$.

Next, we encode the bin-wise upper bounds prescribed by $\vec{k}$ via a partition matroid. Let $B_0, \dots, B_T$ be the bins (a partition of E). Define the partition matroid $\mathcal{M}_k^P$ on E by

$$S \in \mathcal{I}_k^P = \text{IS}(\mathcal{M}_k^P) \iff |S \cap B_i| \leq k_i \text{ for all } i = 0, \dots, T.$$

We restrict the $\mathcal{P}$ -MOL problem to sets that satisfy both matroid constraints and have exactly K elements:

$$\mathcal{F}_{\vec{k}} := \{S \subseteq E : S \in \mathcal{I}_{\leq K} \cap \mathcal{I}_k^P \text{ and } |S| = K\}.$$

If $K > \text{rank}(\mathcal{M})$, then $\mathcal{F}_{\vec{k}} = \emptyset$ (since no independent set can have size K); in the algorithm we simply skip such configurations.

Equivalently, $\mathcal{F}_{\vec{k}}$ is the family of common bases of $\mathcal{M}_{\leq K}$ and $\mathcal{M}_k^P$ (since in each matroid, a set is a base iff it is independent and has size equal to the matroid rank, here K). If the original feasible family is $\mathcal{F} = \text{IS}(\mathcal{M})$, then the above definition applies for any $K \leq \text{rank}(\mathcal{M})$. If instead $\mathcal{F} = \text{bases}(\mathcal{M})$, then we must have $K = \text{rank}(\mathcal{M})$; otherwise no set of size K can be a base of $\mathcal{M}$. Accordingly, in the bases variant we only consider configuration vectors with $K = \text{rank}(\mathcal{M})$ (and for other K we take $\mathcal{F}_{\vec{k}} = \emptyset$).

Finally, note that membership in $\mathcal{F}_{\vec{k}}$ forces the bin constraints of $\mathcal{M}_k^P$ to be tight. Indeed, for any $S \in \mathcal{F}_{\vec{k}}$, we have

$$|S \cap B_i| \leq k_i \quad \forall i \in \{0, \dots, T\} \quad \text{and} \quad \sum_{i=0}^T |S \cap B_i| = |S| = K = \sum_{i=0}^T k_i,$$

implying that $|S \cap B_i| = k_i$ for all i . In particular, every $S \in \mathcal{F}_{\vec{k}}$ is a base of the partition matroid $\mathcal{M}_k^P$ and also a base of the truncated matroid $\mathcal{M}_{\leq K}$ since $|S| = K$.

For $\text{opt} = \max$, let $\text{OPT}_{\vec{k}}$ denote the maximum of $v(S)$ over all $S \in \mathcal{F}_{\vec{k}}$ satisfying $w(S) \triangleright L$; for $\text{opt} = \min$ we define $\text{OPT}_{\vec{k}}$ analogously as a minimum. If $\vec{k} = \vec{k}^{\text{opt}}$ is the configuration of an optimal solution, then clearly $\text{OPT}_{\vec{k}^{\text{opt}}} = \text{OPT}$.

2.3 Lagrangian Relaxation

Returning to our problem, we must satisfy both the matroid-intersection constraint and the linear constraint: $w(S) \triangleright L$. Since the linear constraint is typically the harder one to handle directly, we apply *Lagrangian relaxation* to incorporate it into the objective via a penalty term.

To express the different orientations in a unified form, we define two sign parameters:

$$s_{\text{opt}} = \begin{cases} +1 & \text{if } \text{opt} = \max, \\ -1 & \text{if } \text{opt} = \min, \end{cases} \quad s_{\triangleright} = \begin{cases} +1 & \text{if } \triangleright = \geq, \\ -1 & \text{if } \triangleright = \leq. \end{cases}$$

For $\lambda \geq 0$ and $S \subseteq E$, we define the Lagrangian objective

$$\phi_{\lambda}(S) := v(S) + s_{\text{opt}} s_{\triangleright} \lambda (w(S) - L).$$

The corresponding Lagrangian relaxation value for some fixed configuration vector $\vec{k}$ is

$$V_{\vec{k}}(\lambda) := \text{opt}_{S \in \mathcal{F}_{\vec{k}}} \phi_{\lambda}(S). \quad (3)$$

Finally, we define the optimal value of the Lagrangian relaxation by

$$\text{LR}_{\vec{k}} := \begin{cases} \min_{\lambda \geq 0} V_{\vec{k}}(\lambda) & \text{if } \text{opt} = \max, \\ \max_{\lambda \geq 0} V_{\vec{k}}(\lambda) & \text{if } \text{opt} = \min. \end{cases}$$

► **Lemma 4.** *For every configuration vector $\vec{k}$, the function $V_{\vec{k}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is piecewise linear. Moreover, if $\text{opt} = \max$, then $V_{\vec{k}}$ is convex, and if $\text{opt} = \min$, then $V_{\vec{k}}$ is concave.*

Proof. Fix $\vec{k}$ and consider any $S \in \mathcal{F}_{\vec{k}}$. The function $\phi_{\lambda}(S)$ is affine in λ . Therefore, $V_{\vec{k}}(\lambda)$ is the pointwise optimum of finitely many affine functions (one for each $S \in \mathcal{F}_{\vec{k}}$), and hence $V_{\vec{k}}$ is piecewise linear.

Moreover, if $\text{opt} = \max$ then $V_{\vec{k}}$ is the pointwise maximum of affine functions and hence convex; if $\text{opt} = \min$ then $V_{\vec{k}}$ is the pointwise minimum of affine functions and hence concave. ◀

► **Lemma 5.** *For every configuration vector $\vec{k}$, if $\text{opt} = \max$, then $\text{LR}_{\vec{k}} \geq \text{OPT}_{\vec{k}}$, and if $\text{opt} = \min$, then $\text{LR}_{\vec{k}} \leq \text{OPT}_{\vec{k}}$.*

Proof. Let S be any feasible solution with configuration $\vec{k}$, i.e., $S \in \mathcal{F}_{\vec{k}}$ and $w(S) \triangleright L$. Then $s_{\triangleright}(w(S) - L) \geq 0$. For any $\lambda \geq 0$,

$$s_{\text{opt}} \phi_{\lambda}(S) = s_{\text{opt}} v(S) + s_{\triangleright} \lambda (w(S) - L) \geq s_{\text{opt}} v(S).$$

Optimizing over all such S yields $s_{\text{opt}} V_{\vec{k}}(\lambda) \geq s_{\text{opt}} \text{OPT}_{\vec{k}}$. If $\text{opt} = \max$ (so $s_{\text{opt}} = +1$), this implies $V_{\vec{k}}(\lambda) \geq \text{OPT}_{\vec{k}}$ for all λ , hence $\text{LR}_{\vec{k}} = \min_{\lambda \geq 0} V_{\vec{k}}(\lambda) \geq \text{OPT}_{\vec{k}}$. The minimization case is analogous. ◀

We now argue that the Lagrangian relaxation (3) can be solved efficiently.

► **Lemma 6.** *For every configuration vector $\vec{k}$ and any $\lambda \geq 0$, the value $V_{\vec{k}}(\lambda)$ and a set $S_{\vec{k}, \lambda}^* \in \mathcal{F}_{\vec{k}}$ attaining this value can be computed in polynomial time. Moreover, a multiplier $\lambda_{\vec{k}}^*$ satisfying $V_{\vec{k}}(\lambda_{\vec{k}}^*) = \text{LR}_{\vec{k}}$ can be computed in polynomial time.*

Proof. Fix a configuration vector $\vec{k}$ and $\lambda \geq 0$, and define the modified weight function v^λ as follows

$$v_e^\lambda := v_e + s_{\text{opt}S_\triangleright} \lambda w_e \quad \text{for all } e \in E.$$

Then, for any $S \subseteq E$, we can write

$$\phi_\lambda(S) = v(S) + s_{\text{opt}S_\triangleright} \lambda (w(S) - L) = \sum_{e \in S} v_e^\lambda - s_{\text{opt}S_\triangleright} \lambda L.$$

Hence, for fixed λ , optimizing ϕ_λ over $S \in \mathcal{F}_{\vec{k}}$ is equivalent (up to an additive constant independent of S) to solving the following optimization problem

$$\text{opt}_{S \in \mathcal{F}_{\vec{k}}} \sum_{e \in S} v_e^\lambda.$$

For both $\mathcal{F} \in \{\text{IS}, \text{bases}\}$, the restriction to $\mathcal{F}_{\vec{k}}$ imposes a fixed cardinality constraint $|S| = K$ (where $K = \sum_i k_i$, and in the bases case necessarily $K = \text{rank}(\mathcal{M})$, otherwise $\mathcal{F}_{\vec{k}} = \emptyset$). Consequently, the problem of optimizing a linear objective over $\mathcal{F}_{\vec{k}}$ is a cardinality-constrained weighted matroid intersection instance between $\mathcal{M}$ (or $\mathcal{M}_{\leq K}$) and the partition matroid $\mathcal{M}_k^P$. This problem is solvable in polynomial time by the algorithm of Brezovec et al. [2]. Therefore, for any fixed $\vec{k}$ and $\lambda \geq 0$, we can compute $V_{\vec{k}}(\lambda)$ in polynomial time.

It remains to find a multiplier $\lambda_{\vec{k}}^* \geq 0$ such that $V_{\vec{k}}(\lambda_{\vec{k}}^*) = \text{LR}_{\vec{k}}$. Since the objective weights v_e^λ depend linearly on λ , we can apply Megiddo's parametric-search framework [25] to compute such an optimal multiplier $\lambda_{\vec{k}}^*$ in polynomial time. $\blacktriangleleft$

Let $\lambda_{\vec{k}}^*$ be such an optimal multiplier, and define the family of Lagrangian-optimal solutions

$$\mathcal{O}_{\vec{k}} := \{S \in \mathcal{F}_{\vec{k}} : \phi_{\lambda_{\vec{k}}^*}(S) = \text{LR}_{\vec{k}}\}.$$

For every $S \in \mathcal{O}_{\vec{k}}$ we have $\phi_{\lambda_{\vec{k}}^*}(S) = \text{LR}_{\vec{k}}$, which can be rewritten as

$$v(S) = \text{LR}_{\vec{k}} - s_{\text{opt}S_\triangleright} \lambda_{\vec{k}}^* (w(S) - L). \quad (4)$$

We next consider the degenerate edge case in which $\lambda_{\vec{k}}^* = 0$. In this case, the Lagrangian objective reduces to the original objective over $\mathcal{F}_{\vec{k}}$. Let

$$S_{\vec{k}}^0 \in \arg \text{opt}_{S \in \mathcal{F}_{\vec{k}}} v(S)$$

be an unconstrained optimum chosen to maximize $s_{\triangleright}(w(S) - L)$ among all unconstrained optima. If $S_{\vec{k}}^0$ is feasible, we add it directly to the candidate pool. If no unconstrained optimum is feasible, this configuration contributes no candidate in the case $\lambda_{\vec{k}}^* = 0$. Henceforth, we assume that $\lambda_{\vec{k}}^* > 0$.

2.4 Spread and Adjacency

We introduce a bound on the change in $v(\cdot)$ under a one-per-bin exchange, and then prove an adjacency property of $\mathcal{O}_{\vec{k}}$ showing that any two solutions in $\mathcal{O}_{\vec{k}}$ are connected by a sequence of such exchanges. For each bin B_i define its *spread*

$$\Delta_i := \max_{e \in B_i} v_e - \min_{e \in B_i} v_e,$$

18:10 Tight Approximation for Matroid Optimization

with the convention that $\Delta_i = 0$ if $B_i = \emptyset$. For a configuration vector $\vec{k}$, define the total spread

$$\Sigma_{\Delta}(\vec{k}) := \sum_{i: k_i > 0} \Delta_i.$$

► **Lemma 7.** *Fix a configuration vector $\vec{k}$ satisfying (2).*

1. *If $S, S' \in \mathcal{F}_{\vec{k}}$ differ by exchanging at most one element per bin, then $|v(S) - v(S')| \leq \Sigma_{\Delta}(\vec{k})$.*
2. *It holds that $\Sigma_{\Delta}(\vec{k}) \leq 8\varepsilon \text{OPT}$.*

Proof. *Part (1).* Fix a bin B_i with $k_i > 0$. By assumption, passing from S to S' removes at most one element from B_i and therefore (to maintain $|S' \cap B_i| = |S \cap B_i|$, since $S, S' \in \mathcal{F}_{\vec{k}}$) it also adds at most one element to B_i . Hence the contribution of bin B_i to the value can change by at most Δ_i , i.e., $|v(S \cap B_i) - v(S' \cap B_i)| \leq \Delta_i$. Summing over all bins with $k_i > 0$ yields

$$|v(S) - v(S')| \leq \sum_{i: k_i > 0} |v(S \cap B_i) - v(S' \cap B_i)| \leq \sum_{i: k_i > 0} \Delta_i = \Sigma_{\Delta}(\vec{k}).$$

Part (2). By the bucketing construction, $\Delta_0 \leq \varepsilon \hat{V}$ and $\Delta_i \leq \varepsilon^2 \hat{V}$ for all $i \geq 1$. Moreover, (2) implies that the number of indices $i \geq 1$ with $k_i > 0$ is at most $\lceil 2/\varepsilon \rceil$. Therefore,

$$\Sigma_{\Delta}(\vec{k}) = \Delta_0 \cdot \mathbf{1}[k_0 > 0] + \sum_{i \geq 1: k_i > 0} \Delta_i \leq \varepsilon \hat{V} + \left\lceil \frac{2}{\varepsilon} \right\rceil \cdot \varepsilon^2 \hat{V} \leq \varepsilon \hat{V} + 3\varepsilon \hat{V} \leq 4\varepsilon \hat{V}.$$

Finally, by Lemma 1, we have $\hat{V} \leq \text{OPT}$ when $\text{opt} = \max$, and $\hat{V} \leq 2\text{OPT}$ when $\text{opt} = \min$; thus, in both cases $\Sigma_{\Delta}(\vec{k}) \leq 8\varepsilon \text{OPT}$. ◀

We next establish the adjacency property of $\mathcal{O}_{\vec{k}}$, namely that any two sets in $\mathcal{O}_{\vec{k}}$ can be connected by a sequence of exchanges that swaps at most one element per bin. Our proof uses the following theorem of Brezovec et al. [3].

► **Lemma 8** (Theorem 4 from [3]). *Let $\mathcal{M}_1 = (E, \mathcal{I}_1)$ be a matroid and let $\mathcal{M}_2 = (E, \mathcal{I}_2)$ be a partition matroid on the same ground set E , given by a partition $E = B_0 \cup \dots \cup B_T$ and capacities $(k_0, \dots, k_T)$. Let $c: E \rightarrow \mathbb{R}$ be a weight function, and let S be a common base of $\mathcal{M}_1$ and $\mathcal{M}_2$.*

Define the directed bipartite exchange graph $B_S = (U \cup V, A)$ by $U := S, V := E \setminus S$, (i.e., one vertex for each element), and for each $e \in U$ and $f \in V$ define arcs as follows:

- *if $S - e + f \in \mathcal{I}_1$, add an arc (e, f) with cost $c(f) - c(e)$;*
- *if $S - e + f \in \mathcal{I}_2$, add an arc (f, e) with cost 0.*

Let D_S be the directed graph obtained from B_S by compressing, for each $i \in \{0, \dots, T\}$, all vertices in $V \cap B_i$ into a single vertex v_i , and then deleting parallel arcs while keeping, for each ordered pair of vertices, only the arc of minimum length.

For the graph D_S the following hold:

- *There are no negative cycles in D_S if and only if S is optimal for the weighted matroid intersection instance.*
- *If D_S contains a directed cycle C of negative total length, then there exists a common base T of $\mathcal{M}_1$ and $\mathcal{M}_2$ that is obtained from S by exchanging elements along C and satisfies $c(T) < c(S)$.*

Using this lemma, we obtain the following exchange-path guarantee.

► **Lemma 9.** *For any $S, S' \in \mathcal{O}_{\vec{k}}$, we can find in polynomial time a sequence of solutions $S = S_0, S_1, \dots, S_\ell = S'$ such that:*

1. $S_j \in \mathcal{O}_{\vec{k}}$ for all $j = 0, \dots, \ell$;
2. For each $j = 0, \dots, \ell - 1$ and for each bin B_i ,

$$|(S_j \setminus S_{j+1}) \cap B_i| = |(S_{j+1} \setminus S_j) \cap B_i| \leq 1.$$

Proof. Fix a configuration vector $\vec{k}$, and let $v^* = v^{\lambda_{\vec{k}}^*}$. W.l.o.g., assume that $\mathcal{O}_{\vec{k}}$ consists of *minimum-cost* bases with respect to v^* (otherwise, replace v^* by $-v^*$).

We prove the statement by induction on $x := |S' \setminus S|$. The case $x = 0$ is trivial. Assume that $x \geq 1$ and the lemma holds for every pair $T, T' \in \mathcal{O}_{\vec{k}}$ satisfying $|T' \setminus T| < x$. We prove the claim for the given pair S, S' with $|S' \setminus S| = x$.

Fix $\delta > 0$ and define $v' : E \rightarrow \mathbb{R} \cup \{+\infty\}$ by

$$v'(e) = \begin{cases} v^*(e) + \delta, & e \in S \setminus S', \\ v^*(e) - \delta, & e \in S' \setminus S, \\ v^*(e), & e \in S \cap S', \\ +\infty, & e \notin S \cup S'. \end{cases}$$

Then $v'(S) = v^*(S) + x\delta$ and $v'(S') = v^*(S') - x\delta$, and hence $v'(S') = v^*(S') - x\delta < v^*(S) + x\delta = v'(S)$, where we used $v^*(S) = v^*(S')$ since $S, S' \in \mathcal{O}_{\vec{k}}$. To avoid choosing an explicit infinitesimal δ , we implement this perturbation lexicographically. Let U be an integer upper bound on $\sum_{e \in S} |v^*(e)|$ over all $S \in \mathcal{F}_{\vec{k}}$ (e.g., $U := K \cdot \max_{e \in E} |v^*(e)|$), and set $M := 2U + 1$. Define an auxiliary weight function $\tilde{v} : E \rightarrow \mathbb{R}$ by

$$\tilde{v}(e) = \begin{cases} M \cdot v^*(e) + 1, & e \in S \setminus S', \\ M \cdot v^*(e) - 1, & e \in S' \setminus S, \\ M \cdot v^*(e), & e \in S \cap S', \\ M \cdot v^*(e) + M^2 + K, & e \notin S \cup S'. \end{cases}$$

We claim that $\tilde{v}(S') < \tilde{v}(S)$. Indeed, since $S, S' \in \mathcal{O}_{\vec{k}}$ we have $v^*(S) = v^*(S')$. Moreover, letting $x := |S' \setminus S| = |S \setminus S'|$, the definition of $\tilde{v}$ implies

$$\tilde{v}(S) = \sum_{e \in S} \tilde{v}(e) = M \sum_{e \in S} v^*(e) + |S \setminus S'| = M v^*(S) + x,$$

and similarly

$$\tilde{v}(S') = \sum_{e \in S'} \tilde{v}(e) = M \sum_{e \in S'} v^*(e) - |S' \setminus S| = M v^*(S') - x.$$

Therefore,

$$\tilde{v}(S) - \tilde{v}(S') = M(v^*(S) - v^*(S')) + 2x = 2x > 0,$$

and hence $\tilde{v}(S') < \tilde{v}(S)$.

Since the weights $v^*(e)$ may be rational, we first scale them to integers. Let D be a common denominator of $\{v^*(e) : e \in E\}$, and define $\bar{v}^*(e) := D \cdot v^*(e)$ for all $e \in E$. Then $\bar{v}^*(e) \in \mathbb{Z}$ for all e , and for any two common bases T_1, T_2 we have $\bar{v}^*(T_1) - \bar{v}^*(T_2) \in \mathbb{Z}$.

We redefine $\tilde{v}$ using $\bar{v}^*$ as follows. Let $U := K \cdot \max_{e \in E} |\bar{v}^*(e)|$ and set $M := 2U + 1$. Define $\tilde{v} : E \rightarrow \mathbb{Z}$ by $\tilde{v}(e) = M \cdot \bar{v}^*(e) + 1$ for $e \in S \setminus S'$, $\tilde{v}(e) = M \cdot \bar{v}^*(e) - 1$ for $e \in S' \setminus S$, $\tilde{v}(e) = M \cdot \bar{v}^*(e)$ for $e \in S \cap S'$, and $\tilde{v}(e) = M \cdot \bar{v}^*(e) + M^2 + 2MU + 2K$ for $e \notin S \cup S'$.

Our redefinition preserves the computation for S and S' . In particular, using $\bar{v}^*(S) = \bar{v}^*(S')$ we have $\tilde{v}(S) = M \cdot \bar{v}^*(S) + x$ and $\tilde{v}(S') = M \cdot \bar{v}^*(S') - x$, so $\tilde{v}(S') < \tilde{v}(S)$.

18:12 Tight Approximation for Matroid Optimization

▷ **Claim 10.** Any $\tilde{v}$ -minimum common base is (i) a v^* -minimum common base (and hence belongs to $\mathcal{O}_{\vec{k}}$), and (ii) consists solely of elements from $S \cup S'$.

Proof. We first show that minimizing $\tilde{v}$ enforces minimization of v^* (lexicographically). Fix any common base $T \in \mathcal{F}_{\vec{k}}$. Write $\tilde{v}(T) = M \cdot \bar{v}^*(T) + \tau(T)$, where $\tau(T)$ is the additive tie-breaking term induced by the ± 1 offsets (and 0 offsets) in the definition of $\tilde{v}$. Since $|T| = K$, we have the uniform bound $|\tau(T)| \leq K$ whenever $T \subseteq S \cup S'$. Moreover, if T contains some element $e \notin S \cup S'$, then

$$\tilde{v}(T) \geq M \cdot \bar{v}^*(T) - K + (M^2 + 2MU + 2K) \geq -MU - K + (M^2 + 2MU + 2K) = M^2 + MU + K.$$

On the other hand, for every common base $T \subseteq S \cup S'$ we have

$$\tilde{v}(T) = M \cdot \bar{v}^*(T) + \tau(T) \leq MU + K.$$

Therefore any $\tilde{v}$ -minimum common base must be contained in $S \cup S'$.

Now let $T_1, T_2 \subseteq S \cup S'$ be any two common bases. Then

$$\tilde{v}(T_1) - \tilde{v}(T_2) = M(\bar{v}^*(T_1) - \bar{v}^*(T_2)) + (\tau(T_1) - \tau(T_2)),$$

and $|\tau(T_1) - \tau(T_2)| \leq 2K$. Choose U so that $U \geq K \cdot \max_{e \in E} |\bar{v}^*(e)|$, and set $M := 2U + 1$ as above. Then $M > 2K$. Moreover, since $\bar{v}^*(T_1), \bar{v}^*(T_2) \in \mathbb{Z}$, if $\bar{v}^*(T_1) > \bar{v}^*(T_2)$ then $\bar{v}^*(T_1) - \bar{v}^*(T_2) \geq 1$, and hence

$$\tilde{v}(T_1) - \tilde{v}(T_2) = M(\bar{v}^*(T_1) - \bar{v}^*(T_2)) + (\tau(T_1) - \tau(T_2)) \geq M - 2K > 0.$$

Therefore $\tilde{v}(T_1) > \tilde{v}(T_2)$. Consequently, among all common bases contained in $S \cup S'$, every $\tilde{v}$ -minimum base is also v^* -minimum. Moreover, as shown above, if a common base T contains an element outside $S \cup S'$, then $\tilde{v}(T) \geq M^2 + MU + K$, whereas every common base $T \subseteq S \cup S'$ satisfies $\tilde{v}(T) \leq MU + K$. Hence no $\tilde{v}$ -minimum common base uses elements outside $S \cup S'$. Therefore, any $\tilde{v}$ -minimum common base belongs to $\mathcal{O}_{\vec{k}}$ and is contained in $S \cup S'$. $\triangleleft$

Apply Lemma 8 to the pair $(\mathcal{M}_{\leq K}, \mathcal{M}_{\vec{k}}^P)$, the weight function $\tilde{v}$, and the current base S . Since $\tilde{v}(S') < \tilde{v}(S)$, the base S is not $\tilde{v}$ -optimal, so the compressed graph D_S contains a negative cycle. Let C be a negative directed cycle in D_S with the minimum number of arcs. Lemma 8 yields a common base S'' obtained from S by exchanging along C such that $\tilde{v}(S'') < \tilde{v}(S)$. By the choice of $\tilde{v}(e)$ for $e \notin S \cup S'$, we have $S'' \subseteq S \cup S'$.

Since D_S is bipartite, any directed cycle C alternates between vertices $e \in S$ and contracted bin vertices v_i . We chose C with the minimum number of arcs, hence C cannot visit the same v_i twice; otherwise C would split into two directed cycles, and at least one of them would still have negative total length but fewer arcs, contradicting minimality. Thus, C visits each bin at most once, so the corresponding exchange inserts at most one element from each bin B_i , i.e. $|(S'' \setminus S) \cap B_i| \leq 1$. Finally, S and S'' are both bases of the partition matroid $\mathcal{M}_{\vec{k}}^P$, so in each bin B_i an insertion must be matched by a deletion from the same bin in order to keep $|S \cap B_i| = k_i$. Hence, the numbers of inserted and deleted elements match in every bin:

$$|(S \setminus S'') \cap B_i| = |(S'' \setminus S) \cap B_i| \leq 1 \quad \forall i \in \{0, \dots, T\}.$$

Since $S \in \mathcal{O}_{\vec{k}}$ is v^* -minimum, we have $v^*(T) \geq v^*(S)$ for every common base T . In particular, for every common base $T \subseteq S \cup S'$ we have $\bar{v}^*(T) \geq \bar{v}^*(S)$. Now $S'' \subseteq S \cup S'$ and $\tilde{v}(S'') < \tilde{v}(S)$. If $\bar{v}^*(S'') > \bar{v}^*(S)$ then $\tilde{v}(S'') - \tilde{v}(S) \geq M - 2K > 0$, a contradiction. Hence $\bar{v}^*(S'') = \bar{v}^*(S)$, i.e., $v^*(S'') = v^*(S)$, and therefore $S'' \in \mathcal{O}_{\vec{k}}$.

Finally, since the inserted elements come from $E \setminus S$ and, by the definition of $\tilde{v}$, must lie in $S' \setminus S$, we add at least one element of $S' \setminus S$; thus $|S' \setminus S''| < |S' \setminus S| = x$.

The induction hypothesis applies to the pair (S'', S') . Hence there exists a sequence $S'' = S_0, S_1, \dots, S_m = S'$ with $S_j \in \mathcal{O}_{\vec{k}}$ for all j and such that each transition $S_j \rightarrow S_{j+1}$ exchanges at most one element per bin. Appending the initial step $S \rightarrow S''$ to this sequence yields the desired path from S to S' .

The above inductive argument yields an explicit polynomial-time procedure to construct the sequence. Given $S, S' \in \mathcal{O}_{\vec{k}}$, form the perturbed weight function $\tilde{v}$ as in the proof. Since $\tilde{v}(S') < \tilde{v}(S)$, the base S is not $\tilde{v}$ -optimal, so by Lemma 8 the compressed exchange graph D_S contains a negative directed cycle, which can be found in polynomial time (e.g. by Bellman–Ford; extracting a *simple* negative cycle ensures no bin-vertex v_i is repeated). Exchanging along this cycle produces $S'' \in \mathcal{O}_{\vec{k}}$ that differs from S by at most one element per bin and satisfies $|S' \setminus S''| < |S' \setminus S|$. Repeating for at most $|S' \setminus S| \leq K$ iterations constructs $S = S_0, S_1, \dots, S_\ell = S'$. ◀

2.5 Deriving an EPTAS

We now show that for every feasible configuration vector $\vec{k}$, there is a Lagrangian-optimal solution that is both feasible for the original linear constraint and nearly optimal in value.

► **Lemma 11.** *Fix a configuration vector $\vec{k}$ with $\mathcal{F}_{\vec{k}} \neq \emptyset$, and let $\lambda_{\vec{k}}^*$, $\text{LR}_{\vec{k}}^*$, and $\mathcal{O}_{\vec{k}}$ be as defined above. Then, there exists a set $S^* \in \mathcal{O}_{\vec{k}}$ such that $w(S^*) \triangleright L$ and the following holds:*

- if $\text{opt} = \min$, then $v(S^*) \leq \text{LR}_{\vec{k}} + \Sigma_{\Delta}(\vec{k})$,
- if $\text{opt} = \max$, then $v(S^*) \geq \text{LR}_{\vec{k}} - \Sigma_{\Delta}(\vec{k})$.

Proof. Recall that $\mathcal{O}_{\vec{k}} = \arg \text{opt}_{S \in \mathcal{F}_{\vec{k}}} \phi_{\lambda_{\vec{k}}^*}(S)$. Choose $S^-, S^+ \in \mathcal{O}_{\vec{k}}$ such that $s_{\triangleright}(w(S^-) - L)$ is minimized over $\mathcal{O}_{\vec{k}}$ and $s_{\triangleright}(w(S^+) - L)$ is maximized over $\mathcal{O}_{\vec{k}}$ (These can be computed in polynomial time by lexicographic optimization over $\mathcal{F}_{\vec{k}}$: first optimize $\phi_{\lambda_{\vec{k}}^*}$, and among ties minimize/maximize $s_{\triangleright}w(S)$, implemented via weighted matroid intersection).

We claim that $s_{\triangleright}(w(S^-) - L) \leq 0 \leq s_{\triangleright}(w(S^+) - L)$. Otherwise, if $s_{\triangleright}(w(S) - L) > 0$ for all $S \in \mathcal{O}_{\vec{k}}$ (or < 0 for all such S), then perturbing $\lambda_{\vec{k}}^*$ slightly in the appropriate direction would improve $V_{\vec{k}}(\lambda)$, contradicting the optimality of $\lambda_{\vec{k}}^*$ in the definition of $\text{LR}_{\vec{k}}^*$.

For simplicity, we present the remainder of the proof for the minimization case; the maximization case is analogous. Using Eq. (4) with $\text{opt} = \min$ (so $s_{\text{opt}} = -1$), we have $v(S) - \text{LR}_{\vec{k}}^* = s_{\triangleright}\lambda_{\vec{k}}^*(w(S) - L)$. Since $\lambda_{\vec{k}}^* \geq 0$, it follows that $s_{\triangleright}(w(S^-) - L) \leq 0$ implies $v(S^-) \leq \text{LR}_{\vec{k}}^*$, and $s_{\triangleright}(w(S^+) - L) \geq 0$ implies $v(S^+) \geq \text{LR}_{\vec{k}}^*$.

By Lemma 9, there is a sequence $S^- = S_0, S_1, \dots, S_\ell = S^+$ inside $\mathcal{O}_{\vec{k}}$ such that each step exchanges at most one element per bin. Since $v(S_0) \leq \text{LR}_{\vec{k}}^*$ and $v(S_\ell) \geq \text{LR}_{\vec{k}}^*$, there exists an index j such that $v(S_j) \leq \text{LR}_{\vec{k}}^* \leq v(S_{j+1})$.

Combining this inequality with Lemma 7 yields

$$v(S_{j+1}) \leq v(S_j) + \Sigma_{\Delta}(\vec{k}) \leq \text{LR}_{\vec{k}}^* + \Sigma_{\Delta}(\vec{k}).$$

Moreover, since $v(S_{j+1}) \geq \text{LR}_{\vec{k}}^*$ and $v(S_{j+1}) - \text{LR}_{\vec{k}}^* = s_{\triangleright}\lambda_{\vec{k}}^*(w(S_{j+1}) - L)$ with $\lambda_{\vec{k}}^* \geq 0$, it follows that $s_{\triangleright}(w(S_{j+1}) - L) \geq 0$, i.e., $w(S_{j+1}) \triangleright L$. Thus $S^* := S_{j+1}$ is feasible and satisfies the desired bound, which completes the proof.

All objects used above can be computed in polynomial time. First, $\lambda_{\vec{k}}^*$ is computed as in Lemma 6, and S^-, S^+ are obtained by two lexicographic weighted-matroid-intersection calls over $\mathcal{F}_{\vec{k}}$ (optimize $\phi_{\lambda_{\vec{k}}^*}$, and among ties minimize/maximize $s_{\triangleright}w(S)$). Next, Lemma 9 yields an explicit sequence $S^- = S_0, \dots, S_\ell = S^+$ inside $\mathcal{O}_{\vec{k}}$ with one-per-bin exchanges; scanning this sequence finds an index j with $v(S_j) \leq \text{LR}_{\vec{k}}^* \leq v(S_{j+1})$, and we output $S^* := S_{j+1}$. ◀

■ **Algorithm 1** EPTAS for $\mathcal{P}$ -MOL .

Input: An instance $\mathcal{A}_P = (E, \mathcal{I}, v, w, L)$, a descriptor $\mathcal{P} = (\text{opt}, \mathcal{F}, \triangleright)$ and $\varepsilon > 0$.
Output: A feasible set S which is $(1 \pm \varepsilon)$ -approximation of OPT .

- 1: Compute a feasible set $\hat{S}$ and set $\hat{V} := v(\hat{S})$, as guaranteed by Lemma 1.
- 2: Define the buckets $B_0, \dots, B_T \subseteq E$.
- 3: $\mathcal{K} \leftarrow \{\vec{k} = (k_0, \dots, k_T) \in \mathbb{Z}_{\geq 0}^{T+1} : \sum_{i=1}^T k_i \leq \lceil 2/\varepsilon \rceil, \sum_{i=0}^T k_i \leq n\}$.
- 4: $S^{\text{best}} \leftarrow \hat{S}$
- 5: **for** $\vec{k} \in \mathcal{K}$ **do**
- 6: Construct the partition matroid $\mathcal{M}_{\vec{k}}^P$ from $\vec{k}$ and define $\mathcal{F}_{\vec{k}} \subseteq 2^E$.
- 7: Compute $\lambda_{\vec{k}}^*$ and $\text{LR}_{\vec{k}}$ (Lemma 6), and compute two sets $S_{\vec{k}}^-, S_{\vec{k}}^+ \in \mathcal{O}_{\vec{k}}$ as in Lemma 11 (with $s_{\triangleright}(w(S_{\vec{k}}^-) - L) \leq 0 \leq s_{\triangleright}(w(S_{\vec{k}}^+) - L)$).
- 8: Using Lemmas 7, 9 and 11, compute $S_{\vec{k}}^* \in \mathcal{F}_{\vec{k}}$ with $w(S_{\vec{k}}^*) \triangleright L$ and $s_{\text{opt}}(\text{LR}_{\vec{k}} - v(S_{\vec{k}}^*)) \leq \Sigma_{\Delta}(\vec{k})$.
- 9: **if** $\text{opt} = \min$ and $v(S_{\vec{k}}^*) < v(S^{\text{best}})$ **then**
- 10: $S^{\text{best}} \leftarrow S_{\vec{k}}^*$
- 11: **if** $\text{opt} = \max$ and $v(S_{\vec{k}}^*) > v(S^{\text{best}})$ **then**
- 12: $S^{\text{best}} \leftarrow S_{\vec{k}}^*$
- 13: **return** S^{best}

Our main result is the following.

► **Theorem 12.** *For every fixed $\varepsilon > 0$, there exists an EPTAS for the $\mathcal{P}$ -MOL problem. Given an instance $\mathcal{A}_P$, the algorithm runs in time $|E|^{O(1)} \cdot (1/\varepsilon^2)^{O(1/\varepsilon)}$ and outputs a feasible solution S satisfying*

$$\text{opt} = \min \implies v(S) \leq (1 + \varepsilon) \text{OPT}, \quad \text{opt} = \max \implies v(S) \geq (1 - \varepsilon) \text{OPT}.$$

Proof. We first describe our EPTAS for the $\mathcal{P}$ -MOL class (see Algorithm 1). Using a constant-factor estimate $\hat{V}$ of OPT , we first bucket the elements by their v -values so that elements within the same bucket have nearly equal value. We then enumerate all feasible configuration vectors $\vec{k}$; each such $\vec{k}$ defines a partition matroid $\mathcal{M}_{\vec{k}}^P$ and a restricted family $\mathcal{F}_{\vec{k}}$. For every configuration, we solve the Lagrangian relaxation of the linear constraint over $\mathcal{F}_{\vec{k}}$ to obtain $\lambda_{\vec{k}}^*$, $\text{LR}_{\vec{k}}$ and the set $\mathcal{O}_{\vec{k}}$ of Lagrangian-optimal solutions. Finally, using the adjacency property of $\mathcal{O}_{\vec{k}}$ together with a bound on the per-bucket value spread, we find a set $S_{\vec{k}}^* \in \mathcal{F}_{\vec{k}}$ that satisfies the original linear constraint and is within $\Sigma_{\Delta}(\vec{k})$ of the Lagrangian bound. The algorithm returns the best such solution over all configurations.

We now proceed to prove the approximation guarantee of Algorithm 1. Let S^{opt} be an optimal solution and $\vec{k}^{\text{opt}}$ its configuration vector. Lemma 2 establishes that $\vec{k}^{\text{opt}}$ satisfies Eq. (2), so the algorithm considers this vector. In that iteration, we get $\text{OPT}_{\vec{k}^{\text{opt}}} = \text{OPT}$. It holds that

$$\text{opt} = \min \implies v(S_{\vec{k}^{\text{opt}}}^*) \leq \text{LR}_{\vec{k}^{\text{opt}}} + \Sigma_{\Delta}(\vec{k}^{\text{opt}}) \leq \text{OPT} + \Sigma_{\Delta}(\vec{k}^{\text{opt}}) \leq \text{OPT} + 8\varepsilon \text{OPT} = (1 + 8\varepsilon) \text{OPT},$$

$$\text{opt} = \max \implies v(S_{\vec{k}^{\text{opt}}}^*) \geq \text{LR}_{\vec{k}^{\text{opt}}} - \Sigma_{\Delta}(\vec{k}^{\text{opt}}) \geq \text{OPT} - \Sigma_{\Delta}(\vec{k}^{\text{opt}}) \geq \text{OPT} - 8\varepsilon \text{OPT} = (1 - 8\varepsilon) \text{OPT},$$

where the first inequality in each row is by Lemma 11, the second is due to Lemma 5 and the last holds by Lemma 7. Therefore, running the algorithm with parameter $\varepsilon/8$ internally, we obtain the claimed $(1 \pm \varepsilon)$ guarantees. The feasibility of $S_{\vec{k}^{\text{opt}}}^*$ is immediate due to Lemma 11.

The running time is dominated by the enumeration of configuration vectors (Lemma 3) and, for each $\vec{k}$, a polynomial number of calls to (cardinality-constrained) weighted matroid intersection and local exchange operations. This yields a total running time of $|E|^{O(1)} \cdot (1/\varepsilon^2)^{O(1/\varepsilon)}$, as required. ◀

3 Hardness Proof for MIC

► **Theorem 13.** *The MATROID INTERSECTION COVER problem admits no EPTAS, unless $W[1] = \text{FPT}$.*

Proof. We show a reduction from the k -clique problem. Given a graph $G = (V, E)$ and $k \in \mathbb{N}$ we need to decide if there is a clique of size k in G . Let $n = |V|$, and denote by $C(G)$ the set of all cliques in G . We define below a MIC instance, where we assume, without loss of generality, that $0 < k \leq n$ and $n \geq 1$; otherwise, the problem becomes trivial.

Let $V = \{v_1, \dots, v_n\}$, and for each vertex $v_i \in V$ define a copy u_i , forming the set $U = \{u_1, \dots, u_n\}$, where $U \cap V = \emptyset$. We refer to u_i and v_i as the two *copies* of the same vertex. Define a function $f : V \cup U \rightarrow [n]$ satisfying the following properties:

1. The restriction of f to V is a bijection onto $[n]$.
2. For every $i \in [n]$, it holds that $f(u_i) = f(v_i)$.

For a set $T \subseteq V \cup U$, we write $f(T) := \{f(x) \mid x \in T\}$ for the image of T under f . We also write $\sigma_f(T) := \sum_{x \in T} f(x)$ for the sum of the labels of the elements in T . Using the function f , we define a matroid $M_1 = (V \cup U, \mathcal{F}_1)$, where the independent sets are given by $\mathcal{F}_1 = \mathcal{I}_1 \cup \mathcal{I}_2 \cup \mathcal{I}_3 \cup \mathcal{I}_4$, with the four families defined as follows:

$$\begin{aligned} \mathcal{I}_1 &= \{T \subseteq V \cup U \mid |T| < n\} \\ \mathcal{I}_2 &= \{T \subseteq V \cup U \mid |T| = n, f(T) \neq f(V)\} \\ \mathcal{I}_3 &= \{T \subseteq V \cup U \mid |T| = n, |T \cap V| \neq k\}. \\ \mathcal{I}_4 &= \{T \subseteq V \cup U \mid |T| = n, |T \cap V| = k, f(T) = f(V), T \cap V \in C(G)\}. \end{aligned} \tag{5}$$

▷ **Claim 14.** M_1 is a matroid.

Proof. We verify the three matroid axioms. First, $\emptyset \in \mathcal{I}_1$ since $0 < n$, and therefore $\emptyset \in \mathcal{F}_1$. Second, for the hereditary property, let $A \in \mathcal{F}_1$ and let $B \subseteq A$. Since $|B| < |A| \leq n$, we have $|B| < n$, hence $B \in \mathcal{I}_1 \subseteq \mathcal{F}_1$. Third, for the exchange property, let $A, B \in \mathcal{F}_1$ such that $|A| > |B|$. We distinguish cases based on the size of B .

- $|B| < n - 1$. Then, for all $e \in A \setminus B$ it holds that $|B + e| < (n - 1) + 1 = n$, and hence $B + e \in \mathcal{I}_1 \subseteq \mathcal{F}_1$. Such an element $e \in A \setminus B$ exists because $|A| > |B|$.
- $|B| = n - 1$ and $|A| = n$. We further distinguish two subcases.
 1. $B \subseteq A$. Then, $A \setminus B = \{e\}$ for some e , and thus $B + e = A \in \mathcal{F}_1$.
 2. $B \not\subseteq A$. Since $|B| = n - 1$ and $|A| = n$, we have $|B \cap A| \leq |B| - 1 = n - 2$. Hence, $|A \setminus B| = |A| - |A \cap B| \geq n - (n - 2) = 2$, so there exist $x, y \in A \setminus B$, such that $x \neq y$. We now consider two cases according to their images under f .
 - a. $f(x) \neq f(y)$. Therefore, there is $q \in \{x, y\}$ such that $f(B + q) \neq f(V)$. We conclude that $B + q \in \mathcal{I}_2 \subseteq \mathcal{F}_1$, as required.
 - b. $f(x) = f(y)$. By the definition of f , there exists $i \in [n]$ such that $\{x, y\} = \{v_i, u_i\}$ (or vice versa). Hence, adding x or y to B changes the number of selected vertices from V , i.e., $|(B + x) \cap V| \neq |(B + y) \cap V|$. Thus, it follows that for at least one $q \in \{x, y\}$, $|(B + q) \cap V| \neq k$, implying $B + q \in \mathcal{I}_3 \subseteq \mathcal{F}_1$. ◁

We now define a second matroid, $M_2 = (U \cup V, \mathcal{F}_2)$, by

$$\mathcal{F}_2 = \{T \subseteq V \cup U \mid |T \cap \{u_i, v_i\}| \leq 1 \ \forall i \in [n]\}. \tag{6}$$

18:16 Tight Approximation for Matroid Optimization

Note that M_2 is a partition matroid, in which any independent set contains at most one copy of each vertex. We define a weight function $w : (U \cup V) \rightarrow \mathbb{R}_{\geq 0}$ by

$$w(x) = \begin{cases} 1 + \frac{f(x)}{2n^2}, & x \in V, \\ \frac{f(x)}{2n^2} & x \in U. \end{cases}$$

and a valuation function $v : (U \cup V) \rightarrow \mathbb{R}_{\geq 0}$, given by

$$v(x) = \begin{cases} 1, & x \in V, \\ 0 & x \in U. \end{cases}$$

Finally, we set the threshold $L = k + \frac{\sigma_f(V)}{2n^2}$ and define the MIC instance $\mathcal{A} = (V \cup U, \mathcal{I}, v, w, L)$, where $\mathcal{I} = \mathcal{F}_1 \cap \mathcal{F}_2$.

▷ **Claim 15.** G contains a clique of size k if and only if $\mathcal{A}$ admits a feasible solution of value k .

Proof. Let Q_V be a clique of size k in G . Define $Q_U = \{u_i \mid i \in [n], v_i \notin Q_V\}$ and $K = Q_V \cup Q_U$. By construction, K contains exactly one copy of each vertex, either v_i or u_i . Therefore, $K \cap V = Q_V \in C(G)$, $|K \cap V| = |Q_V| = k$, and $f(K) = f(V)$, since $f(u_i) = f(v_i)$ for every $i \in [n]$ (and K contains exactly one copy from each vertex, either from U or from V).

Thus, by (5) we have that $K \in \mathcal{I}_4$ implying that $K \in \mathcal{F}_1$. Moreover, since K includes exactly one element from each pair $\{u_i, v_i\}$, it also satisfies the partition constraint and hence $K \in \mathcal{F}_2$. We therefore have $K \in \mathcal{F}_1 \cap \mathcal{F}_2$. In addition, the weight of K satisfies

$$w(K) = w(Q_V) + w(Q_U) = k + \frac{\sigma_f(Q_V)}{2n^2} + \frac{\sigma_f(Q_U)}{2n^2} = k + \frac{\sigma_f(V)}{2n^2} = L,$$

where we use that $|Q_V| = k$, $f(u_i) = f(v_i)$ for all $i \in [n]$, and K contains exactly one copy of each vertex. Hence, K is a feasible solution for $\mathcal{A}$ with value $v(K) = |K \cap V| = |Q_V| = k$.

For the other direction, let Q be a solution for $\mathcal{A}$ with a value k . By the definition of the valuation function v , it follows that $|Q \cap V| = k$. In addition, the total weight of Q is bounded by

$$w(Q) = k + \frac{\sigma_f(Q)}{2n^2} \leq k + \frac{\sigma_f(V)}{2n^2} = L. \quad (7)$$

The first equality follows from $|Q \cap V| = k$, and the inequality holds because $Q \in \mathcal{F}_2$, so Q contains at most one of $\{v_i, u_i\}$ for each $i \in [n]$, so f contributes at most once per vertex of V . Since Q is feasible for $\mathcal{A}$, it must also satisfy $w(Q) \geq L$, and therefore $w(Q) = L$. Combined with (7), this implies $\sigma_f(Q) = \sigma_f(V)$.

Now, notice that if $|Q| < n$, then $\sigma_f(Q) < \sigma_f(V)$, since M_2 ensures that each vertex contributes at most once. Thus we must have $|Q| = n$. We have now established that

$$|Q \cap V| = k, \quad |Q| = n, \quad \sigma_f(Q) = \sigma_f(V).$$

Since $Q \in \mathcal{F}_2$, it contains at most one element from each pair $\{u_i, v_i\}$, and therefore the labels in $f(Q)$ are distinct. As $|Q| = n$ and $\sigma_f(Q) = \sigma_f(V) = 1 + \dots + n$, it follows that $f(Q) = f(V)$. Since Q is feasible, it belongs to one of the families defining $\mathcal{F}_1$; the only possible family is $\mathcal{I}_4$. By (5), we conclude that $Q \cap V \in C(G)$. Hence, $Q \cap V$ is a clique of size k in G . ◁

To complete the proof, assume toward a contradiction that there exists an EPTAS for MIC. Using this scheme, we can decide whether G contains a clique of size k in FPT time. Consider the instance $\mathcal{A}$ constructed above, and run the EPTAS with $\varepsilon = \frac{1}{2k}$. Observe that any feasible solution for $\mathcal{A}$ must include at least k vertices from V . The weight constraint encoded by w and L enforces this requirement, as each vertex from V contributes 1 to the value function v .

By Claim 15, G has a clique of size k if and only if $\mathcal{A}$ admits a feasible solution of value exactly k . If such a solution for $\mathcal{A}$ exists, then the EPTAS would yield a solution whose value is at most $(1 + \varepsilon)k = k + \frac{1}{2} < k + 1$. Since the valuation function v is integral, the only integer value below $k + 1$ is k , and therefore the scheme must return a solution of value exactly k . Thus, the EPTAS correctly determines whether $\mathcal{A}$ has a solution of value k and, by Claim 15, whether G contains a clique of size k .

The running time of the EPTAS is $g(\frac{1}{\varepsilon}) \cdot m^{O(1)}$ for some computable function g , where m is the encoding size of the k -clique instance. Since $\varepsilon = \Theta(\frac{1}{k})$, our reduction decides the k -clique problem in time $g(k) \cdot m^{O(1)}$. As the k -clique problem is known to be W[1]-Hard [18], it implies that there is no EPTAS for MIC unless $W[1] = \text{FPT}$. ◀

4 Discussion

In this paper, we present tight approximation results for a broad class of matroid optimization problems with a linear constraint, along with a hardness result for MIC. Our unified EPTAS for all $\mathcal{P}$ -MOL problems resolves the complexity status of this class, and the hardness of MIC suggests that such a unified EPTAS is unlikely to extend to *all* $\mathcal{P}$ -MOL variants once the single matroid constraint is replaced by a *matroid intersection* constraint.

A natural direction for future work is to characterize the approximability landscape of $\mathcal{P}$ -MOL variants under matroid intersection constraints. In particular, the existence of a PTAS for MIC remains open.

Finally, it would be interesting to identify settings in which stronger guarantees are attainable. For instance, for graphic matroids and for linear matroids given by an explicit representation, we obtain an EPTAS, while the existence of an FPTAS in these cases remains open.

References

- 1 André Berger, Vincenzo Bonifaci, Fabrizio Grandoni, and Guido Schäfer. Budgeted matching and budgeted matroid intersection via the gasoline puzzle. *Mathematical Programming*, 128(1):355–372, 2011. doi:10.1007/S10107-009-0307-4.
- 2 Carl Brezovec, Gerard Cornuéjols, and Fred Glover. Two algorithms for weighted matroid intersection. *Mathematical Programming*, 36(1):39–53, 1986. doi:10.1007/BF02591988.
- 3 Carl Brezovec, Gérard Cornuéjols, and Fred Glover. A matroid algorithm and its application to the efficient solution of two optimization problems on graphs. *Mathematical Programming*, 42(1):471–487, 1988. doi:10.1007/BF01589417.
- 4 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)*, pages 1797–1815. SIAM, 2021. doi:10.1137/1.9781611976465.108.
- 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. doi:10.1016/S0377-2217(99)00261-1.

- 6 Venkatesan T. Chakaravarthy, Anamitra Roy Choudhury, Sivaramakrishnan R. Natarajan, and Sambuddha Roy. Knapsack Cover Subject to a Matroid Constraint. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2013)*, volume 24 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 275–286. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013. doi:10.4230/LIPIcs.FSTTCS.2013.275.
- 7 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Multi-budgeted matchings and matroid intersection via dependent rounding. In *Proc. SODA*, pages 1080–1097, 2011. doi:10.1137/1.9781611973082.82.
- 8 Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, 2022.
- 9 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)*, pages 2961–2979. SIAM, 2023. doi:10.1137/1.9781611977554.CH113.
- 10 Ilan Doron-Arad, Fabrizio Grandoni, and Ariel Kulik. Unsplittable Flow on a Short Path. In *19th International Symposium on Parameterized and Exact Computation (IPEC 2024)*, 2024.
- 11 Ilan Doron-Arad, Ariel Kulik, and Hadas Shachnai. Budgeted matroid maximization: a parameterized viewpoint. *IPEC*, 2023.
- 12 Ilan Doron-Arad, Ariel Kulik, and Hadas Shachnai. An EPTAS for budgeted matching and budgeted matroid intersection via representative sets. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, pages 49:1–49:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ICALP.2023.49.
- 13 Ilan Doron-Arad, Ariel Kulik, and Hadas Shachnai. An EPTAS for budgeted matroid independent set. In *Proc. SODA*, pages 69–83, 2023. doi:10.1137/1.9781611977585.CH7.
- 14 Ilan Doron-Arad, Ariel Kulik, and Hadas Shachnai. An FPTAS for budgeted laminar matroid independent set. *Operations Research Letters*, 51(6):632–637, 2023. doi:10.1016/J.ORL.2023.10.005.
- 15 Ilan Doron-Arad, Ariel Kulik, and Hadas Shachnai. Lower bounds for matroid optimization problems with a linear constraint. In *51st International Colloquium on Automata, Languages, and Programming (ICALP), July 8–12, 2024*, pages 56:1–56:20, 2024. doi:10.4230/LIPIcs.ICALP.2024.56.
- 16 Ilan Doron-Arad, Ariel Kulik, and Hadas Shachnai. You (almost) can’t beat brute force for 3-matroid intersection. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5151–5171. SIAM, 2026. doi:10.1137/1.9781611978971.186.
- 17 Ilan Doron-Arad, Hadas Shachnai, Gilad Shmerler, and Inbal Talgam-Cohen. An algorithm-to-contract framework without demand queries. *arXiv preprint arXiv:2507.20038*, 2025. doi:10.48550/arXiv.2507.20038.
- 18 Rod G Downey and Michael R Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1-2):109–131, 1995. doi:10.1016/0304-3975(94)00097-3.
- 19 Fabrizio Grandoni and Rico Zenklusen. Approximation schemes for multi-budgeted independence systems. In *Proc. ESA*, pages 536–548, 2010. doi:10.1007/978-3-642-15775-2_46.
- 20 Refael Hassin and Asaf Levin. An efficient polynomial time approximation scheme for the constrained minimum spanning tree problem using matroid intersection. *SIAM Journal on Computing*, 33(2):261–268, 2004. doi:10.1137/S0097539703426775.
- 21 Sung-Pil Hong, Sung-Jin Chung, and Bum Hwan Park. A fully polynomial bicriteria approximation scheme for the constrained spanning tree problem. *Operations Research Letters*, 32(3):233–239, 2004. doi:10.1016/J.ORL.2003.06.003.
- 22 Hans Kellerer, Ulrich Pferschy, and David Pisinger. The multiple-choice knapsack problem. In *Knapsack Problems*, pages 317–347. Springer, 2004.
- 23 Ariel Kulik and Hadas Shachnai. There is no EPTAS for two-dimensional knapsack. *Information Processing Letters*, 110(16):707–710, 2010. doi:10.1016/J.IPL.2010.05.031.

- 24 Eugene L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4(4):339–356, 1979. doi:10.1287/MOOR.4.4.339.
- 25 Nimrod Megiddo. Combinatorial optimization with rational objective functions. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, pages 1–12, 1978. doi:10.1145/800133.804326.
- 26 Ram Ravi and Michel X Goemans. The constrained minimum spanning tree problem. In *Scandinavian Workshop on Algorithm Theory*, pages 66–75. Springer, 1996.
- 27 Yang Yang, Mikhail Y Kovalyov, and Yong Zhang. An improved fptas for budgeted laminar matroid independent set. *Operations Research Letters*, page 107307, 2025. doi:10.1016/J.ORL.2025.107307.