



Hardness and approximation of submodular minimum linear ordering problems

Majid Farhadi¹ · Swati Gupta² · Shengding Sun¹ · Prasad Tetali³ · Michael C. Wigal¹

Received: 23 January 2022 / Accepted: 6 November 2023

© The Author(s) 2023

Abstract

The minimum linear ordering problem (MLOP) generalizes well-known combinatorial optimization problems such as minimum linear arrangement and minimum sum set cover. MLOP seeks to minimize an aggregated cost $f(\cdot)$ due to an ordering σ of the items (say $[n]$), i.e., $\min_{\sigma} \sum_{i \in [n]} f(E_{i,\sigma})$, where $E_{i,\sigma}$ is the set of items mapped by σ to indices $[i]$. Despite an extensive literature on MLOP variants and approximations for these, it was unclear whether the graphic matroid MLOP was NP-hard. We settle this question through non-trivial reductions from minimum latency vertex cover and minimum sum vertex cover problems. We further propose a new combinatorial algorithm for approximating monotone submodular MLOP, using the theory of principal partitions. This is in contrast to the rounding algorithm by Iwata et al. (in: APPROX, 2012), using Lovász extension of submodular functions. We show a $(2 - \frac{1+\ell_f}{1+|E|})$ -approximation for monotone submodular MLOP where $\ell_f = \frac{f(E)}{\max_{x \in E} f(\{x\})}$ satisfies $1 \leq \ell_f \leq |E|$. Our theory provides new approximation bounds for special cases of

A part of this work was done while all the authors were at the Georgia Institute of Technology.

✉ Swati Gupta
swatig@mit.edu

Majid Farhadi
farhadi@gatech.edu

Shengding Sun
ssun313@gatech.edu

Prasad Tetali
ptetali@cmu.edu

Michael C. Wigal
wigal@gatech.edu

¹ Georgia Institute of Technology, Atlanta, GA, USA

² Massachusetts Institute of Technology, Cambridge, MA, USA

³ Carnegie Mellon University, Pittsburgh, PA, USA

the problem, in particular a $(2 - \frac{1+r(E)}{1+|E|})$ -approximation for the matroid MLOP, where $f = r$ is the rank function of a matroid. We further show that minimum latency vertex cover is $\frac{4}{3}$ -approximable, by which we also lower bound the integrality gap of its natural LP relaxation, which might be of independent interest.

Keywords Minimum linear ordering problem · Submodular · Principal partitions · Complexity · Approximation

Mathematics Subject Classification 90C59

1 Introduction

In the Minimum Linear Ordering Problem (MLOP), given a finite set of elements E , and a function over the subsets, $f : 2^E \rightarrow \mathbb{R}$, one seeks an ordering of the elements, i.e., a bijection $\sigma : E \rightarrow \{1, \dots, |E|\}$, that minimizes the aggregated cost over prefixes (or equivalently suffixes) of the ordering. In other words, MLOP is of the form $\min_{\sigma \in \mathcal{S}_E} \sum_{i=0}^{|E|} f(E_{i,\sigma})$, where $E_{i,\sigma} = \{e \in E : \sigma(e) \leq i\}$, and $\mathcal{S}_E$ is the set of permutations of E . This is in contrast to the classical phenomenon of minimizing a cost function over a combinatorial subset of the powerset of the elements, for example, as in the set cover problem or the minimum spanning tree problem.

It is known that the MLOP is NP-hard even with additional assumptions, for example, when the set function $f(\cdot)$ is monotone and submodular, or symmetric and submodular, or supermodular (see Table 1). Despite a rich literature on hardness of MLOP variants, it is unclear whether the problem remains NP-hard for many structured cases, for instance when $f(\cdot)$ is the rank function of a matroid (i.e., submodular, monotone, bounded by set size, and integral). Furthermore, much is still unknown about the related approximation guarantees. In this work, we push the envelope of hardness and approximability for variants of submodular MLOP. In particular, we show the following:

1. Matroid MLOP, graphic matroid MLOP, co-graphic matroid MLOP, and minimum latency vertex cover (MLVC) are NP-hard.
2. Graphic matroid is polynomially-solvable for some classes of graphs.
3. We improve the approximation factors for matroid MLOP to $2 - \frac{1+r(E)}{1+|E|}$ and minimum latency set cover (MLSC) to $2 - \theta$, where $\theta \geq \frac{2}{|E|+1}$, and it depends on the instance, by exploiting the theory of principal partitions. These results provide a refinement of the previously best-known factors for these problems [1].
4. We also show that MLVC can be approximated to $\frac{4}{3}$, improving upon the previously best approximation achieving a factor 2 [1]. We analyze the fractional dimension of a related poset to achieve this bound. We further lower bound the integrality gap of the natural LP relaxation for MLVC.

Here, matroid variants of MLOP are when $f(\cdot)$ is the rank function of the corresponding matroid, and the minimum latency set (vertex) cover problems are defined on a hypergraph (graph) where the vertices must be ordered so that the sum of the

maximum indices at which every hyperedge (edge) is covered must be minimized. We include precise definitions of these problems in Sect. 4.

We summarize our hardness and approximation results on MLOP in Tables 1 and 2. The paper is structured as follows: We give an overview of our results and techniques in Sect. 2, discuss related work on MLOP variants in Sect. 3, and present preliminaries in Sect. 4. We discuss detailed proofs of our results in Sects. 5–8. We finally conclude the paper with open problems in Sect. 9.

2 Overview of results and techniques

We next present an overview of our results and techniques.

1. Hardness of matroid MLOP:

We first show in Sect. 5 the NP-hardness of matroid MLOP, by observing the fact that a uniform matroid on a ground set E ($|E| = n$) with rank k has the unique property (up to isomorphism) of having $\binom{n}{k}$ independent sets of size k . We will show that *any* optimal matroid MLOP solution can detect this, thereby reducing the “uniform matroid isomorphism” problem (known to be NP-hard [2]) to matroid MLOP.

Theorem 1 *Matroid MLOP is NP-hard.*

Furthermore, we show that matroid MLOP in decision form on a family of matroids shares the same complexity class with matroid MLOP in decision form on the matroidal *dual* family. This observation will be useful for upcoming results.

2. Hardness of graphic matroid MLOP:

Next, in Sect. 6, we further restrict $f(\cdot)$ to the special case of the rank function of any *graphic* matroid. Tutte [3] gave a complete minor-free characterization for graphic matroids. In particular, graphic matroids are *regular*, i.e., representable using a totally unimodular matrix, and in particular, do not contain a rank-2 uniform matroid over 4 elements as a minor (e.g., see [4]). Therefore for graphic matroids, the reduction from uniform matroid isomorphism does not suffice. We show that it is NP-hard using a series of reductions beginning at the minimum sum vertex cover (MSVC) on simple graphs G , which we show reduces to the minimum latency vertex cover (MLVC) on the complement graph $\bar{G}$, which we show finally reduces to the graphic matroid MLOP.

Theorem 2 *Graphic matroid MLOP is NP-hard.*

To reduce MLVC to the graphic matroid MLOP with graph $G = (V, E)$, we first create an auxiliary graph H by adding a new vertex z to V , and a weighted star graph T centered at z , connected to each vertex in V . We choose the edge weights for T in such a way that they each induce a distinct flat in any optimal ordering for weighted graphic MLOP. This implies solving (weighted) graphic matroid MLOP for H is equivalent to solving MLVC on G . As we can keep the magnitude of the weights controlled, this allows us to reduce MLVC to graphic matroid MLOP, thereby showing hardness of the latter. As a by product, we also show that MLVC and the co-graphic matroid MLOP are NP-hard, which was not known before our work.

3. Improved approximation of MLSC:

We present two different approximation algorithms for minimum latency set cover problem, both of which refine the best-known constant 2-approximation with improvements in different settings, and use different techniques. The MLSC can be modeled as a covering problem on a hypergraph (see Sect. 3 for details), and in Sect. 7, we present our first randomized approximation algorithm based on scheduling theory, whose approximation factor depends on rank of the hypergraph, i.e. the maximum cardinality of its hyperedges (or in other words, the maximum cardinality of the candidate sets) (Theorem 3). Our second approach simply applies the approximation for monotone submodular MLOP to MLSC, as the latter is a special case of the former (Corollary 2). Both the resultant approximation factors depend on the properties of the instance, and none of them dominate the other on all instances. They both improve on the previous best-known approximation bound for MLSC of 2, using a reduction to the single machine scheduling problem with precedence constraints [5–7].

Theorem 3 *There is a randomized polynomial time algorithm that approximates MLSC within factor $2 - \frac{2}{1+\ell}$, where ℓ is the maximum cardinality among all hyperedges of H .*

The idea for achieving our improved approximation bound for MLSC is to exploit the structural complexity of the precedence constraints (corresponding to a poset) for the subsequent scheduling instance. Bounding the fractional dimension of this poset by $1 + \ell$ allows us to utilize a state-of-the-art scheduling algorithm by Ambühl et al. [8] to approximate the objective by a factor of $2 - \frac{2}{1+\ell}$. For the special case where the input is a graph, this algorithm gives a factor $\frac{4}{3}$ approximation for MLVC.

Corollary 1 *There exists a randomized polynomial time factor $\frac{4}{3}$ -approximation algorithm for MLVC.*

To the best of our knowledge, this is the current best approximation factor for MLVC. For ℓ -uniform regular hypergraphs, i.e., where each hyperedge has size ℓ and each vertex is contained in the same amount of hyperedges, we show that a simple LP relaxation also achieves the $2 - \frac{2}{1+\ell}$ approximation factor. In particular, the LP relaxation gives a factor $\frac{4}{3}$ -approximation algorithm for MLVC on regular graphs. From this result, we raise the question whether the LP relaxation for MLSC on ℓ -uniform hypergraphs has the same $2 - \frac{2}{1+\ell}$ approximation factor. Indeed, a better approximation factor seems unlikely, as we observe a lower bound of $2 - \frac{2}{1+\ell}$ on integrality gap of the LP relaxation for MLSC on ℓ -uniform hypergraphs, matching our current approximation result.

In Sect. 8.3, we discuss the use of principal partitions to obtain an approximation for MLSC as a special case.

Corollary 2 *There is a deterministic factor $(2 - \frac{\Delta+|E|}{\Delta(1+|V|)})$ -approximation algorithm for MLSC, where Δ is the maximum degree of hypergraph $H = (V, E)$.*

Note that $\Delta = \max_{v \in V} |\{e \in E : v \in e\}|$. Together Theorem 3 and Corollary 2 imply that MLSC can be approximated within factor $2 - \theta$, where $\theta = \max\{\frac{2}{1+\ell}, \frac{\Delta+|E|}{\Delta(1+|V|)}\}$. Note θ can be very small, for example, for ℓ -uniform hypergraphs where ℓ is large. However, since $\theta \geq \frac{2}{n+1}$, we get a slight improvement over 2.

4. Polynomially solvable instances of matroid MLOP:

In Sect. 6.4, we propose a novel characterization of matroid MLOP, wherein one can search through bases and permutations of bases, rather than permutations of the ground set. In particular, whenever the number of bases of a matroid is small (polynomial in $|E|$) and the rank of the matroid is also small (constant), we show that matroid MLOP becomes polynomial time solvable.

Theorem 4 *Let $\mathcal{X}$ be a family of matroids such that for all $M = (E, r_M) \in \mathcal{X}$ with $|E| = m$, the number of bases of M is $|\mathcal{B}(M)| \in O(g(m))$, and the rank of M is $r_M(m) \in O(h(m))$, for some $g, h : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$. Then, every matroid MLOP instance in $\mathcal{X}$ can be solved in time $O(g(m) \cdot \text{poly}(m, g(m)) \cdot (h(m))!)$. In particular, if g is polynomial in m and h is bounded by a constant, then matroid MLOP for $\mathcal{X}$ is in P .*

For the special case of graphic matroid MLOP on cactus graphs,¹ we show that an optimal MLOP ordering can be found by fixing any spanning tree of the cactus graph. To find an ordering of the edges of the spanning tree, we show that a greedy ordering on the cycles of the cactus graph suffices (even though the size of the basis may not be logarithmic in size with respect to the ground set).

Theorem 5 *Given a simple cactus graph G , there is a polynomial time algorithm that solves graphic matroid MLOP on G .*

Furthermore, in Sect. 7.2, we show how if a graph is regular, then the optimal objective values for MLA, MSVC, and MLVC are all related by linear shifts in the objective parameterized by the number and degree of the vertices. As many instances of regular graphs have polynomial time algorithms (e.g. see [9–12]) this leads to many new polynomial time algorithms for MSVC and MLVC for many instances of regular graphs.

5. Improved approximation for monotone submodular MLOP:

For monotone submodular MLOP, Iwata, Tetali, and Tripathi [1] provided a $(2 - \frac{2}{|E|+1})$ -approximation algorithm for monotone submodular MLOP based on Lovász extension in 2012. Another natural approach is to use the theory of principal partitions induced by a given submodular function [13, 14]. The principal partition of a ground set E of a monotone submodular function is a chain of subsets $\mathcal{C} = \emptyset \subseteq S_1 \subseteq \dots \subseteq S_k = E$, such that each S_i is the unique maximal minimizer of $f(S) - \lambda_i |S|$ for some $\lambda_i \in \mathbb{R}_+$. As early as 1992, Pinaruk considered completing the chain $\mathcal{C}$ randomly to add subsets of missing cardinality ([15], c.f. [16]). Later in 2019, Fokkink et al. [16] considered the same algorithm for the submodular search problem, which includes monotone submodular MLOP as a special case. They showed that this algorithm has an approximation ratio based on the total curvature² of the submodular function, and is always at most 2.

It was not known how these two results compare, as they use very different techniques. We show that the algorithm based on principal partitions always has better approximation guarantee than the $(2 - \frac{2}{|E|+1})$ bound of Lovász extension relaxation proven in [1].

¹ A graph G is a cactus graph if every maximal 2-connected subgraph of G is a cycle or an edge.

² The total curvature of a set function f is defined to be $\max_{x \in E} \frac{f(\{x\}) + f(E-x) - f(E)}{f(\{x\})}$, e.g., see [16].




Fig. 1 Overview of related problems. A solid arrow from problem A to B indicates that B generalizes A . A dashed arrow from problem A to B denotes that computation of A can be polynomially reduced to computation of B , using our gadgets

Theorem 6 Let $f : 2^E \rightarrow \mathbb{R}$ be a non-trivial, normalized and monotone submodular function. There exists a factor $(2 - \frac{1+\ell_f}{1+|E|})$ -approximation algorithm to MLOP with $f(\cdot)$ in polynomial time, where $\ell_f = \frac{f(E)}{\max_{x \in E} f(\{x\})}$.

As ℓ_f is bounded below by 1, the above result is a refinement of the previous $(2 - \frac{2}{|E|+1})$ -approximation [1]. Our result is also independent from the analysis in Fokkink et al. [16] using total curvature, and leads to nice approximation bounds for some classes of matroids where ℓ_f is large. For example, for graphic matroid MLOP on connected graphs of bounded maximum degree Δ with $\Delta > 1$, we obtain a $(2 - \frac{2}{\Delta})$ -approximation asymptotically. This constant factor improvement from 2 cannot be obtained using either the Lovász extension bound in [1] or the total curvature bound in [16].

Our results have led to multiple open questions which may be of independent interest, and are discussed in Sect. 9.

3 Related work

MLOP was formally introduced by Iwata et al. [1], generalizing many well-known combinatorial optimization problems. In this section, we describe related work in combinatorial optimization that can be viewed as different instances of MLOP. Some of these MLOP variants (e.g., minimum latency set cover) will be utilized in our proof that the graphic MLOP is NP-hard, as depicted in Fig. 1.

Table 1 Previously known results and our results on NP-hardness of different MLOP variants

MLOP class	Problem	Hardness	Source
General	Multiple intents ranking (MIR)	NP-hard	Azar et al. [25]
Monotone supermodular	Minimum sum set cover (MSSC)	NP-hard	Feige et al. [26]
	Minimum sum vertex cover (MSVC)	NP-hard	Feige et al. [26]
Monotone submodular	matroid MLOP	NP-hard	Theorem 1
	Graphic matroid MLOP	NP-hard	Theorem 2
	Co-graphic matroid MLOP	NP-hard	Corollary 5
	Graphic matroid MLOP for cactus graphs	P	Theorem 5
	Minimum latency set cover (MLSC)	NP-hard	Hassin and Levin [27]
	Minimum latency vertex cover (MLVC)	NP-hard	Theorem 9
Submodular	Sum cut (SUMCUT)	NP-hard	[28, 29]
Symmetric submodular	Minimum linear arrangement (MLA)	NP-hard	[20, 30]

Bold values indicate the result of this paper

Minimum linear arrangement (MLA)

Motivated by applications in coding theory, Harper [17] introduced minimum linear arrangement (MLA) in 1964, which seeks to find an arrangement of the vertices of a given graph $G = (V, E)$ such that the total “stretch” of each edge is minimized, i.e., MLA on a graph $G = (V, E)$ is the following,

$$\min_{\pi \in \mathcal{S}_V} \sum_{(u,v) \in E} |\pi(u) - \pi(v)|.$$

Note that any permutation $\pi \in \mathcal{S}_V$ naturally induces a chain on V with prefix sets $V_{i,\pi} = \{v \in V : \pi(v) \leq i\}$. Let ϕ be the cut function of the graph, i.e., for all $S \subseteq V$, $\phi(S)$ is the number of edges with exactly one end in S . Note then for any permutation $\pi \in \mathcal{S}_V$ if an edge $(u, v) \in E$ is stretched to a value $k = |\pi(u) - \pi(v)|$, it must cross the cut of exactly k prefix sets in the chain $V_{0,\pi} \subsetneq V_{1,\pi} \subsetneq \cdots \subsetneq V_{n-1,\pi} \subsetneq V_{n,\pi}$ where $|V(G)| = n$. Thus, MLA on a graph $G = (V, E)$ is equivalent to

$$\min_{\pi \in \mathcal{S}_V} \sum_{i=0}^n \phi(V_{i,\pi}),$$

which is an instance of MLOP with ϕ being a symmetric submodular function. Solving MLA for specific instances of graphs has received considerable attention due to its many applications, see surveys [9–12]. While MLA is polynomial time solvable for some classes of graphs, for example trees [18, 19], its decision form has been known to be NP-complete since 1974 [20]. The best known approximation bound for MLA is $O(\sqrt{\log n} \log \log n)$ [21, 22]. Under the exponential time hypothesis [23] that there does not exist a randomized algorithm to solve SAT in time 2^{n^ε} where n is the instance size and $\varepsilon > 0$ is arbitrarily small, it is also known that MLA is inapproximable to some constant [24].

Table 2 Summary of approximation factors known for MLOP variants

Problem	Approximation	Source
Matroid MLOP	$2 - \frac{1+r(E)}{1+ E }$	Corollary 8
Graphic MLOP [‡]	$2 - \frac{ V(G) }{1+ E(G) }$	Corollary 8
Monotone Submodular MLOP	$2 - \frac{1+\ell_f}{1+ E }$	Theorem 6
MLA	$O(\sqrt{\log n} \log \log n)$	Feige and Lee [21]
MLVC	$\frac{4}{3}$	Theorem 3
MLSC	$2 - \theta$	Theorem 3 and Corollary 2
MIR	4.642	Bansal et al. [31]
SUMCUT	$O(\log n)$	Rao and Richa [32]
Supermodular MLOP	4	Iwata et al. [1]
MSSC	4	Feige et al. [26]
MSVC	$\frac{16}{9}$	Bansal et al. [31]

Bold values indicate the result of this paper

For MLSC, $\theta = \max(\frac{2}{1+\ell}, \frac{\Delta+|E|}{\Delta(1+|V|)})$, where ℓ is the maximum cardinality of hyperedges, and Δ is the maximum degree in the graph. For graphic MLOP, we assume that the graph is connected

Minimum latency set cover (MLSC)

MLSC was introduced by Hassin and Levin [27] with motivations from problems in job scheduling, and they provided an e -approximation. The best known approximation constant for MLSC is 2 [25, 27]. Later in our work, we show that MLSC can be viewed as an instance of monotone submodular MLOP, for which Iwata, Tetali and Tripathi [1] gave a factor $(2 - \frac{2}{|E|+1})$ approximation algorithm using the Lovász extension. We give a more refined approximation algorithm for monotone submodular MLOP using principal partitions, which applies to MLSC as well.

Minimum sum set cover

Minimum sum set cover (MSSC) was introduced by Feige, Lovász, and Tetali [26], who also presented a greedy algorithm that provides a 4-approximate solution to MSSC, and showed it is NP-hard to do better. Later, Iwata et al. [1] showed that MSSC is an instance of supermodular MLOP, and the greedy algorithm for MSSC can be generalized to approximate supermodular MLOP within factor 4.

MSSC can be formulated as follows, using the notation of hypergraphs: given a hypergraph $H = (V(H), E(H))$, MSSC seeks to find a permutation of vertices that minimizes the total costs of all hyperedges, where the cost of each hyperedge is the minimum of its vertex labels, i.e.,

$$\min_{\pi \in S_{V(H)}} \sum_{e \in E(H)} \min_{v \in e} \pi(v).$$

The special case when H is a graph is the well-known minimum sum vertex cover (MSVC). Independent from MSSC, MSVC was introduced earlier by Burer and Mon-

teiro [33] as a heuristic in solving semidefinite relaxation of the Max-Cut problem. Feige, Lovász, and Tetali [26] later showed MSVC has a 2-approximation based on linear programming rounding, and also showed that it is NP-hard to approximate for an unknown constant ϵ , where $1 < \epsilon < 2$. Later, Barenholz, Feige and Peleg [34] improved to a 1.9946-approximation, and recently, Bansal et al. [31] gave a $\frac{16}{9}$ -approximation for MSVC. The best possible approximation constant for MSVC is still unknown. For the special case of regular graphs, Feige, Lovász, and Tetali [26] gave a $\frac{4}{3}$ -approximation. This approximation guarantee for regular graphs was later improved by Stanković [35] to 1.225.

These problems concern supermodular functions, but we only consider submodular functions in this work.

Other variants of MLOP

Another variant of MLOP is called the multiple intents ranking (MIR), and has been studied in [25, 31, 36–38]. Azar, Gamzu, and Yin [25] gave a 2-approximation for MIR, for the case when the weight vector for each hyperedge is monotonically non-decreasing. This variant of MIR includes MLSC as a special case. These problems have found a broad spectrum of applications in query results diversification [39], motion planning for robots [40], cost-minimizing search [41], and optimal scheduling [42], among others.

Another example of an instance of submodular MLOP is the sum cut problem (SUMCUT). The problem was independently introduced by Díaz et al. [28] and also Yixun and Jinjiang [29] to study circuit layouts. SUMCUT is NP-complete [28, 29] and Rao and Richa [32] gave a $O(\log n)$ -approximation algorithm for SUMCUT using a divide-and-conquer approach.

Recently, Happach et al. [43] viewed MLOP under the umbrella of minimum sum ordering/permutation problem, and generalized results of Feige et al. [26].

4 Preliminaries

We now present notation and background useful for parsing this work. We refer an interested reader to [44] for further reading.

1. Submodular set functions

For a set of elements S and elements $x \notin S, y \in S$, we use $S + x, S - y$ to denote $S \cup \{x\}, S \setminus \{y\}$ respectively. Let $f : 2^E \rightarrow \mathbb{R}$ be a set function. We say f is *submodular* if for all $S, T \subseteq E$, $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$. An equivalent definition is $f(S + e) - f(S) \geq f(T + e) - f(T)$ for all $S \subseteq T$ and $e \notin T$. This property is sometimes called the diminishing return property. A set function f is *supermodular* if $-f$ is submodular, and is *symmetric* if $f(S) = f(E \setminus S)$ for all $S \subseteq E$ and is *monotone* if $f(S) \leq f(T)$ for all $S \subseteq T \subseteq E$. We say f is *normalized* if $f(\emptyset) = 0$. A normalized monotone submodular function f is *non-trivial* if $f(E) \neq 0$, i.e., f is not identically zero. For a normalized non-trivial monotone submodular function f , we define the *steepness* of f as $\kappa_f = \max_{x \in E} f(\{x\})$, which is the maximum function value of any singleton. We further define the *linearity* of f as $\ell_f = \frac{f(E)}{\kappa_f}$.

For a normalized symmetric submodular function f , and $s, t \in E$, an $s - t$ cut is a subset $X \subseteq E$ such that $s \in X$, $t \notin X$ or $s \notin X$, $t \in X$. The cut is said to have value $f(X) = f(E \setminus X)$.

For a finite set E with size $m > 0$, we define $\mathcal{S}_E$ to be the set of all bijective functions $\sigma : E \rightarrow \{1, \dots, |E|\}$. For every $\sigma \in \mathcal{S}_E$, we define the prefix sets $E_{i,\sigma} = \{e \in E : \sigma(e) \leq i\}$. Note $|E_{i,\sigma}| = i$ for all $1 \leq i \leq m$ and $\emptyset \subsetneq E_{1,\sigma} \subsetneq E_{2,\sigma} \subsetneq \dots \subsetneq E_{m,\sigma} = E$.

2. Matroids

A rank function (for a matroid) is an integer-valued nonnegative monotone submodular function $r : 2^E \rightarrow \mathbb{Z}_{\geq 0}$, such that $r(A) \leq |A|$ for all $A \subseteq E$. A pair $M = (E, r)$ where r is a rank function on E is a *matroid*. There are multiple equivalent definitions for matroids and we refer to [4] for other equivalent definitions and basic theory. Note if r is a rank function of a non-trivial matroid (where $r(E) > 0$) then $\kappa_r = 1$ and $\ell_r = r(E)$. The set E is the *ground set* of the matroid M , which we also denote $E(M)$. A set $I \subseteq E$ is *independent* if $r(I) = |I|$, and is *dependent* otherwise. A maximal independent set is a *basis*, and the set of all bases of M is denoted as $\mathcal{B}(M)$. A *circuit* of M is a minimally dependent set. Let $e, e' \in E(M)$ for some matroid M , then e is a *loop* if $\{e\}$ is a circuit and e and e' are *parallel* if $\{e, e'\}$ is a circuit. Let B be a basis for M and note for all $e \in E \setminus B$, $B + e$ is a dependent set. It is well known that $B + e$ contains a unique circuit, called the *fundamental circuit* of e with respect to B , which we will denote $C(B, e)$. A *flat* of a matroid is a subset $X \subset E$ that is maximal with respect to its rank. The *closure* of a set $S \subseteq E$, is $\text{cl}(S) = \{x \in E : r(S \cup \{x\}) = r(S)\}$.

A matroid is *uniform* of rank k if its bases consists of all subsets of size k . We denote a uniform matroid of size m and of rank k as U_k^m . If the independent sets of a matroid M is the family of acyclic sets of a graph G , then M is a *graphic matroid*, which we denote $M = M[G]$. If M is a matroid with rank function r , then its *corank* function is the following, $r^*(X) = |X| - r(M) + r(E \setminus X)$. It is well known, see [4], that r^* is also a rank function for a matroid M^* on $E(M)$, and we let M^* denote the *dual matroid* of M . An element is a *coloop* if it is a loop in the dual matroid. A *cographic matroid* is a matroid whose dual matroid is graphic.

For a positive integer m , let $[m] = \{1, 2, \dots, m\}$. Given a $k \times m$ matrix A with integer entries, the vector matroid of A , denoted by $M[A]$, is defined as follows: the ground set is $[m]$, and the rank function of $J \subset [m]$ is the (matrix) rank of $k \times |J|$ submatrix A_J , which is obtained from A by deleting columns whose index is not in J .

Given a matroid $M = (E, r)$, matroid MLOP solves $\min_{\sigma \in \mathcal{S}_E} \sum_{i=0}^{|E|} r(E_{i,\sigma})$, where $E_{i,\sigma} = \{e \in E : \sigma(e) \leq i\}$, and $\mathcal{S}_E$ is the set of permutations of E .

3. Graphs, hypergraphs and partial orders

A graph G over a set of vertices $V(G)$, can be defined by a multiset of edges $E(G) \subseteq V \times V$. We allow graphs to have multiedges and loops, and a graph is *simple* if it does not have multiedges or loops. A graph is a *clique* if every pair of distinct vertices has a single edge joining them. The clique or complete graph on n vertices is denoted K_n . The *complement* of a simple graph G , denoted $\overline{G}$, is the graph where $V(\overline{G}) = V(G)$ and for all distinct $u, v \in V(G)$, $(u, v) \in E(\overline{G})$ if and only if $(u, v) \notin E(G)$. A *block* of a graph G is a maximal connected subgraph, without a

cut vertex. Note that the blocks of G are either an edge or a 2-connected subgraph. It is well known that every pair of distinct blocks are either disjoint or intersect at a cut vertex of G . A *cactus graph* is graph G in which every block of G is an edge or a circuit. A hypergraph $H = (V, E)$ is a generalization of graphs that allows each edge $e \in E$ to be a subset of vertices V , where each such subset is referred to as a hyperedge. A graph is a special case of a hypergraph where all hyperedges have size 2.

We now define the minimum latency set cover (MLSC). Given a hypergraph, MLSC asks to find a permutation on its vertices that minimizes the aggregated cost of the hyperedges, where the cost of an hyperedge is the maximum label of its vertices.

$$\min_{\pi \in \mathcal{S}_{V(H)}} \sum_{e \in E(H)} \max_{v \in e} \pi(v)$$

The minimum latency vertex cover (MLVC) is an instance of MLSC where the input is restricted to being a graph.

A *partially ordered set* or *poset* is a pair $(P, <_P)$ where P is a set and $<_P$ is an antisymmetric and transitive relation on P , i.e., such that for all distinct $x, y, z \in P$, we have that

1. if $x <_P y$ and $y <_P z$ then $x <_P z$, and
2. if $x <_P y$ then $y \not<_P x$.

For any $x, y \in P$ we say x and y are *comparable* if $x <_P y$, $y <_P x$ or $x = y$. A *chain* is a subset $S \subseteq P$ of pairwise comparable elements of $(P, <_P)$. A partial order $(P, <_P)$ is a *total order* if P is a chain. A poset $(P, <_P)$ is an *extension* of a poset $(P, <_P)$ if for all $x, y \in P$, if $x <_P y$ implies $x <_P y$. An extension $(P, <_P)$ is *linear* if $(P, <_P)$ is a total order.

4. Principal partitions

We refer the readers to [14, 45] for the general theory on principal partitions. Here we state some properties of principal partitions on monotone submodular functions.

Theorem 7 [45] *Let f be a monotone submodular function such that $f(A) = 0$ if and only if $A = \emptyset$. Then there exist positive integer $s \geq 1$ and nested sets $\emptyset = \Pi_0 \subsetneq \cdots \subsetneq \Pi_s = E$, called principal partitions of f , as well as real numbers $\lambda_0 < \lambda_1 < \cdots < \lambda_{s+1}$, called critical values, such that for all $0 \leq i \leq s$, Π_i is the unique maximal optimal solution to $\min_{X \subseteq E} f(X) - \lambda|X|$, for all $\lambda \in (\lambda_i, \lambda_{i+1})$. Furthermore, $\{\Pi_i\}_{0 \leq i \leq s}$ as well as $\{\lambda_i\}_{1 \leq i \leq s}$ can be computed in polynomial time.*

Some authors refer to $\{\Pi_i\}_{0 \leq i \leq s}$ as the principal sequence of partitions, and/or use the minimal (which is also unique) instead of maximal optimal solution. Note that the principal partitions minimize the function value among subsets of the same size.

5 Matroid MLOP is as hard as uniform matroid isomorphism

Before we show NP-hardness of graphic matroid MLOP, we will first show in this section that the more general case that matroid MLOP is indeed NP-hard using a

reduction to the uniform matroid isomorphism problem. Although the uniform matroid is one of the simplest matroids, it turns out that determining whether a given matroid is uniform is NP-hard [2]. Formally, the uniform matroid isomorphism problem is the following:

Given a $k \times m$ matrix A with integer entries, is $M[A]$ isomorphic to U_m^k ?

where $M[A]$ denotes the vector matroid of A .

In the following lemma we argue that the optimal matroid MLOP value is unique for each uniform matroid. This provides a reduction to the uniform matroid isomorphism problem.

Lemma 1 *Let $M = (E, r)$ be a matroid, of size $|E| = m$, and rank at most $k \geq 1$. We have*

$$\min_{\sigma \in \mathcal{S}_E} \sum_{i=1}^m r(E_{i,\sigma}) = \binom{k+1}{2} + k(m-k),$$

if and only if M is isomorphic to the uniform matroid U_k^m .

Proof For any uniform matroid U_k^m on ground set E , and for any ordering σ of elements in E we have, $\sum_{i=1}^m r(E_{i,\sigma}) = \binom{k+1}{2} + k(m-k)$, for prefix sets $E_{i,\sigma}$ of the ordering. If M is not isomorphic to the rank- k uniform matroid U_k^m , then it must have some subset $S \subseteq E$ of k elements with rank less than k . As E has rank at most k , ordering elements in S first, followed by elements in $E \setminus S$ arbitrarily constructs a solution with the matroid MLOP value less than the optimal solution for U_k^m . The claim follows. $\square$

Note if A is a $k \times m$ matrix, then $M[A]$ is a matroid of size m and rank at most k . By Lemma 1, if we solve matroid MLOP for $M[A]$, we can determine if $M[A]$ is isomorphic to U_m^k . By NP-hardness of uniform matroid, we have the following theorem.

Theorem 1 *Matroid MLOP is NP-hard.*

For matroid MLOP, the next lemma shows that solving matroid MLOP for any matroid $M = (E, r)$ is as hard as solving matroid MLOP for the dual matroid $M^* = (E, r^*)$. This will be useful to show the hardness of matroid MLOP for cographic matroids.

Lemma 2 *Let $M = (E, r)$ be a matroid with $|E| = m$ and consider an ordering $\sigma \in \mathcal{S}_E$, then*

$$\sum_{i=1}^m r^*(E_{i,\sigma}) = \binom{m+1}{2} - r(M)|E(M)| + \sum_{i=1}^m r(E_{i,\sigma^*}),$$

where σ^ is the reverse permutation, i.e., $\sigma^* = |E| + 1 - \sigma \in \mathcal{S}_E$.*

Proof One can easily verify that $\sigma^* \in \mathcal{S}_E$. As $r^*(X) = |X| - r(M) + r(E \setminus X)$ it follows,

$$\begin{aligned} \sum_{i=1}^m r^*(E_{i,\sigma}) &= \sum_{i=1}^m (|E_{i,\sigma}| - r(M) + r(E \setminus E_{i,\sigma})) \\ &= \binom{m+1}{2} - r(M)|E(M)| + \sum_{i=1}^m r(E \setminus E_{i,\sigma}) \\ &= \binom{m+1}{2} - r(M)|E(M)| + \sum_{i=1}^m r(E_{i,\sigma^*}). \end{aligned}$$

□

Therefore, any optimal ordering of E for matroid MLOP for a given matroid $M = (E, r)$ also gives an optimal ordering for matroid MLOP on the dual matroid $M^* = (E, r^*)$.

Corollary 3 *Matroid MLOP is NP-hard on a family of matroids $\mathcal{X}$ if and only if matroid MLOP is NP-hard on the dual family $\mathcal{X}^* = \{X^* : X \in \mathcal{X}\}$.*

6 Graphic matroid MLOP is NP-hard

We next consider the complexity of graphic matroid MLOP. This turns out to be non-trivial, involving a series of reductions from minimum sum vertex cover, to minimum latency vertex cover, to weighted graphic matroid MLOP, to matroid MLOP.

To show these reductions, we first argue that an optimal chain of matroid MLOP has a useful structure of flats of the matroid, in Lemma 3. Next, in Lemma 4, we reduce weighted matroid MLOP to matroid MLOP. In section 6.2, we provide a reduction from minimum latency vertex cover (MLVC) to weighted graphic matroid MLOP. Finally in Section 6.3, we argue that MLVC and minimum set vertex cover (MSVC) are equivalent in decision form, thus completing the proof that graphic matroid MLOP is NP-hard. In Sect. 6.4 we give an alternative characterization to matroid MLOP. In matroid MLOP we optimize over permutations of the ground set, while in this new formulation, we optimize over bases and then permutations of those bases. Using this characterization, we argue that graphic matroid MLOP for cactus graphs has a polynomial time algorithm.

6.1 Weighted graphic matroid MLOP

In this section, we first argue that any optimal matroid MLOP solution on a ground set E of size m has a nice “flat-like” structure, i.e., for any optimal permutation $\sigma \in \mathcal{S}_E$, the set $\bigcup \{E_{j,\sigma} : r(E_{j,\sigma}) \leq i\}$ is a flat for all $i \in [m]$. This is a useful structural result for optimal solutions and is necessary step towards showing the hardness of graphic matroid MLOP.

Lemma 3 Let $M = (E, r)$ be a matroid of size m and rank k and let $\sigma \in \mathcal{S}_E$ be a permutation that minimizes matroid MLOP. Then, there exists a basis $B = \{b_1, \dots, b_k\} \in \mathcal{B}(M)$ and a partition $\{X_0, X_1, \dots, X_k\}$ of E such that (i) $b_i \in X_i$, and (ii) $\bigcup_{i=0}^j X_i$ is a flat for all $0 \leq j \leq k$, and (iii) $\sigma(e) < \sigma(e')$ for $e \in X_i, e' \in X_l$, and $i < l$.

Proof We may suppose $k \geq 1$, as otherwise the statement is trivial. Let $\sigma \in \mathcal{S}_E$ be a permutation that minimizes matroid MLOP and $X_i := \{e_j : r(E_{j,\sigma}) = i\}$ for $i \geq 0$ and note that $\{X_0, X_1, \dots, X_k\}$ partitions the ground set E . Furthermore $X_i \neq \emptyset$ for all $1 \leq i \leq k$ as for all $e \in E$ and $X \subseteq E$, we have $r(X + e) \leq r(X) + 1$. For each $1 \leq i \leq k$, let $b_i \in X_i$ be the element e in X_i with the lowest index $\sigma(e)$.

For each $1 \leq i \leq k$, we claim $\{b_1, \dots, b_i\}$ is an independent set. For $i = 1$, this is clear. Suppose the claim holds for all positive integers less than j , and $r(\{b_1, \dots, b_j\}) = j - 1$. Note that $b_j \in \text{cl}(\{b_1, \dots, b_{j-1}\}) = \text{cl}(\bigcup_{i=0}^{j-1} X_i)$. As $r(\text{cl}(\{b_1, \dots, b_{j-1}\})) = j - 1$, this contradicts the fact that $r(\bigcup_{i=0}^{j-1} X_i \cup \{b_j\}) = j$. In particular, this implies that $\{b_1, \dots, b_k\}$ is a basis of M and (i) holds.

We now show that $\bigcup_{i=0}^j X_i$ is a flat for each $j < k$. Suppose for $e' \in X_{j'}$ for $j' > j$, that $r(\bigcup_{i=0}^j X_i \cup \{e'\}) = j$. Let $\sigma' \in \mathcal{S}_E$ be the permutation where we place e' before b_{j+1} in σ . That is,

$$\sigma'(e) = \begin{cases} \sigma(e) & \text{if } \sigma(e) < \sigma(b_{j+1}), \\ \sigma(b_{j+1}) & \text{if } e = e', \\ \sigma(e) + 1 & \text{if } \sigma(e) \geq \sigma(b_{j+1}) \text{ and } \sigma(e) < \sigma(e'), \\ \sigma(e) & \text{if } \sigma(e) \geq \sigma(e'). \end{cases}$$

Note $\sum_{i=1}^m r(E_{i,\sigma'}) < \sum_{i=1}^m r(E_{i,\sigma})$. This contradicts the optimality of σ , thus such an e' cannot exist. It follows each $\bigcup_{i=0}^j X_i$ is a flat for each j , and hence (ii) holds. As for all $e \in X_i$ and $e' \in X_l$ with $i < l$, we have $\sigma(e) < \sigma(b_l) \leq \sigma(e')$, thus (iii) holds as well. $\square$

We next introduce a weighted matroid MLOP, which given positive integer costs $c : E \rightarrow \mathbb{Z}_+$ to the elements E of a matroid $M = (E, r)$, checks if there exists a permutation $\sigma \in \mathcal{S}_E$ with weighted MLOP cost at most K , i.e.,

$$\min_{\sigma \in \mathcal{S}_E} \sum_i r(E_{i,\sigma}) c(\sigma^{-1}(i)) \leq K.$$

We now argue that weighted matroid MLOP for a matroid $M = (E, r)$ reduces to matroid MLOP as long as the total integer costs are bounded by a polynomial in $|E|$. This simply follows by duplicating an element $e \in E$ a $c(e)$ number of times, and solving unweighted matroid MLOP on the modified instance. For each duplication of e , $r(E_{i,\sigma})$ is counted $c(e)$ times for each duplicate for any permutation $\sigma \in \mathcal{S}_E$.

Lemma 4 Weighted matroid MLOP with cost function c can be reduced to the matroid MLOP in time $\text{poly}(|E|, c(E))$ where $c(E) := \sum_{e \in E} c(e)$.

Proof Given a matroid $M = (E, r)$ with cost function c , let $N = (E', r')$ be the corresponding matroid where for each $e \in E(M)$ we add $c(e) - 1$ parallel elements to get E' . Let $m' = |E'| = c(E)$ and let $\sigma' \in \mathcal{S}_{E'}$ be an optimal ordering for the matroid MLOP on N . Since we only add parallel elements, the rank function of M induces a natural rank function on N .

Let $r'(N) = k$. By Lemma 3, there exists a partition $\{X_1, \dots, X_k\}$ of E' such that $\bigcup_{i=1}^j X_i$ is a flat for all $1 \leq j \leq k$ and if $e \in X_i$ and $e' \in X_\ell$ for $i < \ell$, then $\sigma'(e) < \sigma'(e')$. Suppose e' is parallel with e . As $\{e', e\}$ is a dependent set, we have that $e \in X_i$ if and only if $e' \in X_i$ for all $1 \leq i \leq k$. We now define a new ordering σ'' by rearranging the elements in σ' such that parallel elements are grouped together by consecutive indices. As parallel elements appear in the same X_i , we have $\sum_{i=1}^{m'} r(E'_{i,\sigma'}) = \sum_{i=1}^{m'} r(E'_{i,\sigma''})$.

Note that σ'' induces a permutation $\sigma \in \mathcal{S}_E$ on the original weighted matroid M such that for any distinct $e, e' \in E$ we have $\sigma(e) < \sigma(e')$ if and only if $\sigma'(e) < \sigma'(e')$. Note then as parallel elements appear in the same partition set X_i we have,

$$\sum_{i=1}^{m'} r'(E'_{i,\sigma'}) = \sum_{i=1}^{m'} r'(E'_{i,\sigma''}) = \sum_{i=1}^m r(E_{i,\sigma})c(\sigma^{-1}(i)).$$

Thus, the optimal weighted matroid MLOP value for M with cost function c is at most the optimal matroid MLOP value on N . Furthermore, one can easily verify that if $\sigma \in \mathcal{S}_E$ obtains the optimal weighted matroid MLOP value for the matroid M with cost function c , there is a corresponding permutation $\sigma' \in \mathcal{S}_{E'}$ that obtains the same matroid MLOP value for N . Thus, the optimal values for both problems are equal. As we only added $c(E) - |E|$ additional elements to N , this is a $\text{poly}(|E|, c(E))$ time reduction. $\square$

6.2 Reducing MLVC to graphic matroid MLOP

We now show that graphic matroid MLOP is as hard as minimum latency vertex cover (MLVC). In MLVC we are given a graph $G = (V(G), E(G))$, and seek to find a permutation of vertices that minimizes the total edge cost, where the cost of each edge is the maximum label of its vertices, i.e.,

$$\min_{\pi \in \mathcal{S}_{V(G)}} \sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\}.$$

Theorem 8 *Minimum latency vertex cover (MLVC) problem can be reduced in polynomial time to the graphic matroid MLOP.*

Proof We will consider an instance of MLVC for a graph G , and construct an auxiliary graph H from G . We will then show that MLVC is equivalent to solving weighted graphic matroid MLOP on H with a specific cost function c . By showing a bound on

the cost of edges $c(E(H)) := \sum_{e \in E(H)} c(e)$ in terms of a polynomial of $|E(G)|$, by applying Lemma 4, we will complete the reduction.

(a) Construction of the graphic matroid MLOP instance: Let G be the given graph with n vertices and m edges. We may assume without loss of generality, G has no isolated vertices, as otherwise an optimal MLVC solution assigns isolated vertices last which play no role in the MLVC cost. Therefore, we have that $n \leq 2m$, by counting the endpoints of the edges which upper bounds the number of vertices.

Let H be a copy of G with an additional vertex z connected to each vertex of G , i.e., $V(H) = V(G) \cup \{z\}$ and $E(H) = E(G) \cup \{(z, v) : v \in V(G)\}$. Let T be the spanning tree of H with $E(T) = \{(z, v) : v \in V(G)\}$. Therefore, H has $n+1$ vertices, and $m+n$ edges. Let $\eta := 9m^2 + 2$ and define $c(\cdot)$ to be a cost function defined on $E(H)$, such that $c(e) = \eta$ if $e \in E(T)$, and $c(e) = 1$ otherwise. Therefore, the total cost of edges in H is polynomially bounded by size of the input graph G :

$$\begin{aligned} c(E(H)) &= \sum_{e \in E(H)} c(e) = \sum_{e \in E(G)} 1 + \sum_{e \in E(T)} \eta \\ &= m + (9m^2 + 2)n \leq m + (9m^2 + 2)2m. \end{aligned}$$

Now, for the sake of brevity, let $E := E(H)$, and $m' = |E(H)|$. Let $\sigma \in \mathcal{S}_E$ be an optimal ordering for weighted graphic matroid MLOP over H with costs $c(\cdot)$. Note,

$$\begin{aligned} \text{MLOP}(H, c, \sigma) &:= \sum_{i=1}^{m'} r(E_{i,\sigma}) c(\sigma^{-1}(i)) \\ &= \sum_{i=1}^{m'} r(E_{i,\sigma}) + \sum_{e \in E(T)} r(E_{\sigma(e),\sigma}) (\eta - 1). \end{aligned}$$

(b) Optimal solutions of graphic matroid MLOP are “good”: We now argue that optimal solutions to graphic matroid MLOP on H have a particular structure. We will argue that if the weights for edges of T are large enough, then analogous to Lemma 3, each edge of T must belong to a different flat induced by σ . This will be useful for relating the solutions of of graphic matroid MLOP on H to MLVC on G .

Let a permutation $\pi \in \mathcal{S}_E$ be *good* if its prefix sets in the ordering has no two edges of T induce the same rank, i.e., $r(E_{i,\pi}) \neq r(E_{j,\pi})$ for all distinct edges $\pi^{-1}(i), \pi^{-1}(j) \in E(T)$. We now argue that σ is an optimal permutation for weighted graphic matroid MLOP over H only if σ is also a good permutation. To show this, we will argue a stronger claim: any good permutation must achieve a lower MLOP objective on H compared to any non-good permutation.

Let $\sigma' \in \mathcal{S}_E$ be an arbitrary good permutation, and for the sake of contradiction, assume that there exists an optimal permutation σ which is not-good. We will argue σ' must have lower MLOP value than σ , giving a contradiction.

Note $\text{MLOP}(H, c, \sigma') \geq \text{MLOP}(H, c, \sigma)$. Furthermore, we have that

$$\sum_{i=1}^{m'} r(E_{i,\sigma'}) \leq \sum_{i=1}^{m'} i \leq (m')^2 = (m+n)^2 \leq 9m^2,$$

as $m+n \leq 3m$. The difference in MLOP objective values of σ and σ' is as follows,

$$\begin{aligned} 0 &\geq \text{MLOP}(H, c, \sigma) - \text{MLOP}(H, c, \sigma') \\ &= \sum_{i=1}^{m'} r(E_{i,\sigma}) - \sum_{i=1}^{m'} r(E_{i,\sigma'}) + (\eta-1) \left(\sum_{e \in E(T)} r(E_{\sigma(e),\sigma}) - r(E_{\sigma'(e),\sigma'}) \right) \\ &\geq - \sum_{i=1}^{m'} r(E_{i,\sigma'}) + (\eta-1) \left(\sum_{e \in E(T)} r(E_{\sigma(e),\sigma}) - r(E_{\sigma'(e),\sigma'}) \right) \\ &\geq -9m^2 + (\eta-1) \left(\sum_{e \in E(T)} r(E_{\sigma(e),\sigma}) - r(E_{\sigma'(e),\sigma'}) \right). \end{aligned}$$

As $\eta-1 > 9m^2$, we have that $\sum_{e \in E(T)} r(E_{\sigma(e),\sigma}) - r(E_{\sigma'(e),\sigma'}) \leq 0$, otherwise the right hand side will be strictly positive (contradicting optimality of σ , and we are done).

Let Y_i be the collection of prefix sets induced by σ up to the edges of T such that the rank is exactly i , i.e., $Y_i := \{E_{\sigma(e),\sigma} : e \in E(T) \text{ and } r(E_{\sigma(e),\sigma}) = i\}$. Note that more than one $E_{\sigma(e),\sigma}$ might belong to Y_i . Furthermore we have $\{Y_i : 1 \leq i \leq n\}$ partitions $\{E_{\sigma(e),\sigma} : e \in E(T)\}$. Similarly, let $Y'_i := \{E_{\sigma'(e),\sigma'} : \sigma' \in E(T) \text{ and } r(E_{\sigma'(e),\sigma'}) = i\}$. As σ' is a good permutation, we have that $|Y'_i| = 1$ for all i . Note,

$$0 \geq \sum_{e \in E(T)} r(E_{\sigma(e),\sigma}) - r(E'_{\sigma'(e),\sigma'}) = \sum_{i=1}^n i(|Y_i| - |Y'_i|) = \sum_{i=1}^n i(|Y_i| - 1).$$

We first argue that $\sum_{j=1}^i |Y_j| \leq i$ for all $1 \leq i \leq n$. Note the elements of $\bigcup_{j=1}^i Y_j$ form a chain of subsets. Define U_i to be the maximum element of $\bigcup_{j=1}^i Y_j$. In particular, as $U_i \in Y_j$ for some $j \leq i$, we have $r(U_i) \leq i$. Furthermore, as $E(T)$ is an independent set,

$$r(U_i) \geq |\{E_{\sigma(e),\sigma} : e \in E(T) \text{ and } E_{\sigma(e),\sigma} \subseteq U_i\}| = \sum_{j=1}^i |Y_j|.$$

Thus we have $\sum_{j=1}^i |Y_j| \leq i$. Note as $\{Y_i : 1 \leq i \leq n\}$ partitions $\{E_{\sigma(e),\sigma} : e \in E(T)\}$, we also have $\sum_{j=1}^n |Y_j| = E(T) = n$.

We now argue that the sum $\sum_{i=1}^n i(|Y_i| - 1)$ is minimized, i.e., $\sum_{i=1}^n i(|Y_i| - 1) = 0$, if only if $|Y_i| = 1$ for all $1 \leq i \leq n$. Suppose $|Y_j| \neq 1$ for some j . Let k be the first

index such that $|Y_k| \neq 1$. As $k \geq \sum_{i=1}^k |Y_i| = |Y_k| + (k-1)$, we have that $|Y_k| = 0$. As $\sum_{i=1}^n |Y_i| = n$, for some $l > k$, we have that $|Y_l| > 1$. Moving an element from Y_l to Y_j would decrease the sum $\sum_{i=1}^n i(|Y_i| - 1)$. Thus $\sum_{i=1}^n i(|Y_i| - 1)$ is minimized, i.e., $\sum_{i=1}^n i(|Y_i| - 1) = 0$, only if $|Y_i| = 1$ for all i . As we assumed σ to be not good, we have $\sum_{i=1}^n (|Y_i| - 1) > 0$, contradicting the optimality of σ . Thus we may conclude any optimal permutation for graphic matroid MLOP on H must also be a good permutation.

(c). Translation of optimal solutions for graphic matroid MLOP to MLVC: We have now argued all optimal solutions on graphic matroid MLOP on H are good, i.e., each edge $e \in E(T)$ has a unique labeling $r(E_{e,\sigma})$. As every vertex of $V(G)$ is incident with exactly one edge of $E(T)$, this ordering of $E(T)$ naturally induces a permutation on $V(G)$. We claim such an ordering will be an optimal MLVC ordering on $V(G)$.

Note as T is a star, all other edges of $E(G) = E \setminus E(T) = E(H) \setminus E(T)$ each form a unique triangle with the edges of T . Let $\pi : V(G) \rightarrow [n]$ such that $\pi(v) = r(E_{(v,z),\sigma})$ where $(v, z) \in E(T)$. As σ is a good permutation, we have that $\pi \in S_{V(G)}$.

We claim that for all $e = (u, v) \in E(G)$, we have $r(E_{\sigma(e),\sigma}) = \max\{\pi(v), \pi(u)\}$ for any optimal permutation σ of $V(H)$. If $r(E_{\sigma(e),\sigma}) > \max\{\pi(u), \pi(v)\}$, consider the ordering σ' in which e appears before the edge $f \in E(T)$ where $r(E_{\sigma(f),\sigma}) = \max\{\pi(u), \pi(v)\}$. This would have strictly decreasing MLOP objective value, a contradiction. Now suppose $r(E_{\sigma(e),\sigma}) < \max\{\pi(v), \pi(u)\}$, then the set

$$A = \{e\} \cup \{f' : f' \in E(T) \text{ and } r(E_{\sigma(f'),\sigma}) \leq r(E_{\sigma(e),\sigma})\}$$

is an independent set of size $r(E_{\sigma(e),\sigma}) + 1$. Now let $f \in A \setminus \{e\}$ such that $r(E_{\sigma(f),f}) = r(E_{\sigma(e),e})$. As either $E_{\sigma(f),f}$ or $E_{\sigma(e),e}$ contains A , we have a contradiction.

Thus we may conclude

$$\text{MLOP}(H, c, \sigma) = \sum_{(u,v) \in E(G)} \max\{\pi(u), \pi(v)\} + \sum_{i=1}^n i \cdot (9m^2 + 2).$$

As $c(E(H))$ is bounded by a polynomial in m , by Lemma 4, the MLVC problem can be reduced in polynomial time to the graphic matroid MLOP. $\square$

6.3 Equivalence of MLVC and MSVC in decision form

The following theorem shows that solving an MLVC instance on a simple graph $G = (V, E)$ where $|V| = n$ is equivalent to solving an MSVC instance on its complement $\bar{G} = (V, E(K_n) \setminus E)$. Since MSVC is known to be NP-hard, by Theorem 8, this will imply graphic matroid MLOP is NP-hard.

Theorem 9 *Let G be a simple graph on n vertices. For any labeling $\pi \in S_{V(G)}$, the MLVC objective on G corresponds to the MSVC objective on its complement graph*

$\overline{G}$ with a linear shift, i.e.,

$$\sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\} = (n^3 - n)/3 - (n+1)|E(\overline{G})| \\ + \sum_{(x,y) \in E(\overline{G})} \min\{\pi'(x), \pi'(y)\},$$

where $\pi' := n + 1 - \pi \in \mathcal{S}_{V(\overline{G})}$.

Proof Note that any labeling of the vertices of a complete graph K_n gives an optimal MLVC objective value of $\sum_{i=1}^n (i-1)i = (n^3 - n)/3$. It follows for all $\pi \in \mathcal{S}_{V(G)}$,

$$\sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\} + \sum_{(x,y) \in E(\overline{G})} \max\{\pi(x), \pi(y)\} = (n^3 - n)/3.$$

This key observation in turn gives the equivalence between MLVC and MSVC as following:

$$\begin{aligned} \sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\} &= (n^3 - n)/3 - \sum_{(x,y) \in E(\overline{G})} \max\{\pi(x), \pi(y)\} \\ &= (n^3 - n)/3 + \sum_{(x,y) \in E(\overline{G})} \min\{-\pi(x), -\pi(y)\} \\ &= (n^3 - n)/3 - (n+1)|E(\overline{G})| \\ &\quad + \sum_{(x,y) \in E(\overline{G})} ((n+1) + \min\{-\pi(x), -\pi(y)\}) \\ &= (n^3 - n)/3 - (n+1)|E(\overline{G})| \\ &\quad + \sum_{(x,y) \in E(\overline{G})} \min\{n+1-\pi(x), n+1-\pi(y)\}. \end{aligned}$$

As $\pi' = n + 1 - \pi \in \mathcal{S}_{V(\overline{G})}$, this completes the proof. $\square$

As MSVC is NP-hard [26], we have that

Corollary 4 *MLVC is NP-hard.*

In Theorem 8, we have reduced any instance of MLVC to graphic matroid MLOP. Combining this with Corollary 4, we have the promise.

Theorem 2 *Graphic matroid MLOP is NP-hard.*

In Corollary 3, we showed if a matroid MLOP is NP-hard for a family of matroids, we have that the corresponding dual family is NP-hard as well. It follows,

Corollary 5 *Cographic matroid MLOP is NP-hard.*

6.4 Graphic matroid MLOP for cactus graphs is in P

We are now interested in further pushing the known boundaries of NP-hardness of graphic matroid MLOP, in particular show that there is a polynomial time algorithm to solve graphic matroid MLOP for cactus graphs. To achieve this we first introduce a new formulation for matroid MLOP which we believe will be of independent interest. In matroid MLOP we optimize over permutations of the ground set. In this new formulation, we first optimize over the bases of the matroid, and then over all permutations of the selected basis. To see this, given a basis B of a matroid and permutation $\pi \in \mathcal{S}_B$, we construct an ordering σ with the following rule, for each $e \notin B$, find the minimal prefix set X of B such that $X \cup e$ is dependent. Place e anywhere after X but before the next element of B in σ . If B and π are chosen as described, this will always result in an optimal MLOP permutation. Now we present this argument in detail.

Let $M = (E, r)$ be a loopless matroid, let $r(M) = k$ and let $\sigma \in \mathcal{S}_E$ have optimal matroid MLOP value. By Lemma 3, there exists a partition of E , say $X = \{X_i : 1 \leq i \leq k\}$ such that $\bigcup_{i=1}^j X_i$ is a flat for all $1 \leq j \leq k$ and there exists a basis $B = \{b_1, \dots, b_k\}$ such that $b_i \in X_i$. Furthermore, we have that if $e \in X_i$ and $e' \in X_\ell$ for $i < \ell$, then $\sigma(e) < \sigma(e')$.

We now observe how this partition $\{X_1, \dots, X_k\}$ interacts with the values of $r(E_{\sigma(e), \sigma})$ for optimal σ . For all $e \in E \setminus B$, $B + e$ has a unique circuit, $C(B, e)$. As $C(B, e) - e$ is an independent set, we have $|\{r(E_{\sigma(e'), \sigma}) : e' \in C(B, e) - e\}| = |C(B, e) - e|$. As $C(B, e)$ is a dependent set and $\bigcup_{i=1}^j X_i$ is a flat for all $1 \leq j \leq k$, we have that

$$r(E_{\sigma(e), \sigma}) = \max\{r(E_{\sigma(e'), \sigma}) : e' \in C(B, e) - e\}.$$

Furthermore as $B = \{b_1, \dots, b_k\}$ with $b_i \in X_i$ for all $1 \leq i \leq k$, we have that there is a one-to-one correspondence between $\{b_1, \dots, b_k\}$ and $\{r(E_{\sigma(e'), \sigma}) : e' \in B\} = \{1, \dots, r(M)\}$. With this in mind, we define for all $\pi \in \mathcal{S}_B$ and fundamental circuits $C(B, e)$, the set $C(B, e)_\pi := \{\pi(e') : e' \in C(B, e) - e\}$. Note that $C(B, e)_\pi$ is the set of positions in the ordering π of the edges present in $C(B, e) - e$.

We now build a permutation σ of E as follows. First select a basis B of the matroid, and permutation $\pi \in \mathcal{S}_B$. Given this ordering of basis elements, we create a linear extension σ of this order by ensuring that:

- For all distinct $b, b' \in B$, $\sigma(b) < \sigma(b')$ if and only if $\pi(b) < \pi(b')$;
- For all $e \in E \setminus B$, if $\max C(B, e)_\pi = i$, then $\sigma(\pi^{-1}(i)) < \sigma(e) < \sigma(\pi^{-1}(i+1))$.

This process always constructs a permutation $\sigma \in \mathcal{S}_E$, and if the correct basis and π are chosen, will find the optimal matroid MLOP permutation. In particular,

Proposition 1 *Matroid MLOP is equivalent to the following problem,*

$$\min_{B \in \mathcal{B}(M)} \min_{\pi \in \mathcal{S}_B} \sum_{e \in E(M) \setminus B} \max C(B, e)_\pi.$$

The characterization Proposition 1 leads to a new class of matroids in which matroid MLOP is in P.

Theorem 10 *Let $\mathcal{X}$ be a family of matroids such that for all $M = (E, r_M) \in \mathcal{X}$ with $|E| = m$, the number of bases of M is $|\mathcal{B}(M)| \in O(g(m))$, and the rank of M is $r_M(m) \in O(h(m))$, for some $g, h : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$. Then, every matroid MLOP instance in $\mathcal{X}$ can be solved in time $O(g(m) \cdot \text{poly}(m, g(m)) \cdot (h(m))!)$. In particular, if g is polynomial in m and h is bounded by a constant, then matroid MLOP for $\mathcal{X}$ is in P .*

Proof By Proposition 1, matroid MLOP for $\mathcal{X}$ has the following formulation,

$$\min_{B \in \mathcal{B}(M)} \min_{\pi \in \mathcal{S}_B} \sum_{e \in E(M) \setminus B} \max C(B, e)_{\pi}.$$

By [46], iterating over every basis requires $\text{poly}(m, |\mathcal{B}(M)|)$ time. As $|\mathcal{B}(M)| \leq g(m)$ and $|\mathcal{S}_B| \leq (h(m))!$, simply iterating over every basis B and its corresponding permutations will solve matroid MLOP for $\mathcal{X}$ in time $O(g(m) \cdot \text{poly}(m, g(m)) \cdot (h(m))!)$. $\square$

We will now use Proposition 1 to solve graphic matroid MLOP for cactus graphs. We will first argue that the selection of spanning tree is arbitrary in finding an optimal solution for cactus graphs. Then we order greedily with respect to the size of the circuits of the graph to find an optimal solution.

Theorem 5 *Given a simple cactus graph G , there is a polynomial time algorithm that solves graphic matroid MLOP on G .*

Proof Let G be a cactus graph. We may assume G is connected, as every graphic matroid M has a connected graph H such that $M = M[H]$. Note as G is a cactus graph, each edge of G belongs to at most one cycle. Our algorithm is as follows:

1. Order the cycles by length in nondecreasing order, temporarily regarding a bridge as a cycle of infinite length.
2. Output any linear extension that respects this prior ordering. That is, first output all edges in the shortest cycle (in any order), followed by all edges in the next shortest cycle (in any order), and so on.

We now show its correctness. As bridges are coloops, a straightforward consequence of Lemma 2 implies bridges must come last in an optimal order. Thus, without loss of generality we may assume G is bridgeless as well. By Proposition 1, graphic matroid MLOP can be formulated as follows,

$$\min_{T \in \mathcal{B}(M)} \min_{\pi \in \mathcal{S}_T} \sum_{e \in E(M) \setminus T} \max C(T, e)_{\pi}.$$

where T is a spanning tree of G , which again for convenience, we regard as a set of edges. Note as G is a cactus graph, the set of fundamental circuits corresponds to the set of cycles of G , i.e., does not depend on the choice of T . The algorithm to solve MLOP for G is clear, first select an arbitrary spanning tree T , and then order the cycles of G non-decreasing with respect to their lengths. Finally, choose a $\pi \in \mathcal{S}_T$ that respects this ordering of the circuits. It is straight forward to verify that this ordering minimizes MLOP. $\square$

7 Approximations for minimum latency set cover (MLSC)

In Sect. 6, we introduced the MLVC problem in a series of reductions to show the graphic MLOP is NP-hard. Here we study its more general version MLSC, introduced by Hassin and Levin in 2005 [27]. In Sect. 7.1 we present a randomized factor $(2 - \frac{2}{1+\ell})$ -approximation algorithm for MLSC, based on techniques from scheduling theory, where ℓ is the size of largest input subset. Our result is better than previously best-known factor of 2 for generic instances [25]. In particular, our result implies a randomized factor $\frac{4}{3}$ -approximation algorithm for MLVC, improving upon Azar et al.'s result [25].

We also show that for ℓ -uniform hypergraphs, the natural linear programming (LP) relaxation (see eq. (MLSC-LP)) has an integrality gap of at least $2 - \frac{2}{1+\ell}$. As a special case, we show that the integrality gap for MLVC is $\frac{4}{3}$. This implies that any approximation algorithm for MLVC based on the rounding of the LP relaxation (without additional inequalities) cannot improve upon our result.

In Sect. 7.2, we explore families of instances where MLVC admits polynomial time algorithms. We show an equivalence between MLA and MLVC for regular graphs in decision form. As many classes of regular graphs have previously been studied, this yields exact polynomial time algorithms for MLVC on these families of instances, and by Theorem 9, for MSVC problem for the graph complement of these families as well.

7.1 A randomized approximation algorithm for MLSC based on scheduling

Recall that minimum latency vertex cover is a special case of minimum latency set cover (MLSC). MLSC can be similarly defined, as in our notation for MLVC. Instead of a graph, we are given a hypergraph $H = (V, E)$ with the objective

$$\min_{\pi \in \mathcal{S}_{V(H)}} \sum_{e \in E(H)} \max_{v \in e} \pi(v).$$

The state-of-the-art approximation for MLSC is a factor 2, using a reduction to a well studied problem in scheduling theory, known as $1|\text{prec}|\sum w_j C_j$, or (single machine) minimum sum scheduling with precedence constraints; that is defined as follows. The input includes a set of jobs J , with corresponding processing times and weights $\{p_j\}_{j \in J}$, $\{w_j\}_{j \in J}$, along with a partially ordered set (poset) P over the jobs. We have a single machine that takes p_j amount of time to process the job j . A feasible schedule is one that processes job j earlier than job j' whenever $j <_P j'$ in the poset. The objective is to minimize (weighted) sum of all completion times, $\sum_j w_j C_j$, where each C_j is the completion time of job j , and is uniquely determined by the schedule and processing times.

MLSC has been known to be reducible to single machine minimum sum scheduling with precedence constraints since 2005 [27], using a simple construction as follows. For every vertex $v \in V$, consider a job $v \in J$ with processing time $p_v = 1$ and weight $w_v = 0$. For every hyperedge $e \in E$, consider a job $e \in J$ with processing time

$p_e = 0$ and weight $w_e = 1$. The poset P over the set of jobs $J = V \cup E$ is defined by all pairs $v <_P e$ such that $v \in V, e \in E$, and $v \in e$. For convenience, we also have for any distinct hyperedges $e, e' \in E$ if $e \subsetneq e'$ then $e <_P e'$. Furthermore, for all multiples of the same edge in E , we order them as a chain in P in some arbitrary manner. It is easy to verify the objective of this scheduling problem is equal to that of the original MLSC. Moreover, the reduction is approximation preserving, i.e., an α -approximate solution to the scheduling instance gives an α -approximate solution to MLSC [27].

Note that the 2-approximability of MLSC is immediate, using various 2-approximations for scheduling [5–7]. Furthermore, by Proposition 6, MLSC is an instance of monotone submodular MLOP. Thus the $(2 - 2/(|E| + 1))$ approximation of [1] is applicable in this case as well. A better constant than 2-approximation for all instances seems unlikely, considering hardness results for the scheduling problem [47], or the vertex cover problem that it reduces to [48–50]. We instead show an instance-dependent improvement parameterized by the maximum size of the subsets. We achieve this result by studying the dimension of the poset and its fractional dimension (e.g., studied by [8, 51] in the context of scheduling). In the rest of this section, we prove Theorem 3 using the scheduling algorithm by [8].

Theorem 3 *There is a randomized polynomial time algorithm that approximates MLSC within factor $2 - \frac{2}{1+\ell}$, where ℓ is the maximum cardinality among all hyperedges of H .*

We now define the fractional dimension of a poset, that was introduced by [52]. A poset $P'(<_{P'})$ is an *extension* of a poset $P(<_P)$, if $x <_P y$ then $x <_{P'} y$, and P' is *linear* if $x \neq y$ then we have $x <_{P'} y$ or $y <_{P'} x$. It is easy to see that the set of feasible solutions for the single machine scheduling problem are all linear extensions of the corresponding poset. Let $\mathcal{F} = \{\mathcal{L}_1, \dots, \mathcal{L}_t\}$ be a multiset of linear extensions of P . $\mathcal{F}$ is a *k-fold realizer* of P , if for every incomparable pair (x, y) of P , there are at least k linear extensions in $\mathcal{F}$ in which $y < x$. The *fractional dimension* of P is defined as $\lim_{k \rightarrow \infty} \frac{t}{k}$, where t is the size of a minimum k -fold realizer (note that the fractional dimension of a poset ≥ 2 if it is not a linear order). Ambühl et al. [8] showed $1/\text{prec} \sum w_j C_j$ can be $(2 - \frac{2}{f})$ -approximated, where f upper bounds the fractional dimension of the corresponding poset. Specifically, they proved the following.

Theorem 11 [8] *Given an efficient sampling algorithm for a k-fold realizer of P , of size t (that is, to output each of the $\mathcal{L}_i$'s with probability at least $1/t$), the problem $1/\text{prec} \sum w_j C_j$ has a randomized approximation algorithm of factor $2 - \frac{2}{t/k}$.*

Given an oracle that outputs a random linear extension P' of P such that $\Pr[j <_{P'} i] \geq b$, for every pair of incomparable jobs (i, j) in P , Theorem 11 gives a $2 - 2b$ approximate solution to the corresponding $1/\text{prec} \sum w_j C_j$. Let us call the sampling algorithm provided to the above theorem, a $\frac{k}{t}$ -balanced linear ordering oracle for P . We show that it is easy to construct an $\frac{1}{1+\ell}$ -balanced linear ordering oracle for posets corresponding to the MLSC's reformulation to scheduling. This will result in a $(2 - \frac{2}{1+\ell})$ -approximation algorithm for MLSC, using the result of Ambühl et. al Theorem 11.

Lemma 5 Consider an arbitrary MLSC problem defined over a hypergraph $H = (V, E)$. Let P be the poset obtained from the reformulation of the MLSC instance as a scheduling problem. Then, P admits a $\frac{1}{1+\ell}$ -balanced linear ordering oracle, where ℓ is the maximum size of any hyperedge in MLSC.

Proof Consider the following linear extensions to the poset P constructed randomly: pick any random ordering $\{v_{l_1}, v_{l_2}, \dots, v_{l_n}\}$ of the vertices V and let them appear in the schedule in this order. To schedule any hyperedge $e \in E$, insert e in the ordering as soon as all its incident vertices have been scheduled. If edges are scheduled concurrently, we break ties at random. It is easy to see that this random scheduling order leads to a valid linear extension, satisfying all precedence constraints of P . Let's call this linear extension P' .

Now, we claim that any random order obtained above satisfies that the probability of $j <_{P'} i$ for two incomparable jobs i, j of P is at least $\frac{1}{1+\ell}$. To see this, note that for a pair of vertices, this trivially holds as $\Pr[u <_{P'} v] = 0.5 \geq \frac{1}{1+\ell}$ for all distinct vertices u and v . Let us show the inequality holds for a pair of incomparable hyperedges. For an incomparable pair consisting of a vertex and a hyperedge, we overload the notation to treat any vertex as a hyperedge of size 1. We can now consider any two distinct incomparable hyperedges e, e' .

Let $a = |e \setminus e'|$, let $b = |e' \setminus e|$, and let $c = |e \cap e'|$. Note that $a, b > 0$, otherwise one edge is a subset of another, i.e., they are not incomparable. We compute $\Pr[e <_{P'} e']$ conditioning on the last vertex of $e \cup e'$ with respect to the random permutation. Call this last vertex $v_{e,e'}$.

$$\begin{aligned} \Pr[e <_{P'} e'] &= \Pr[e <_{P'} e' | v_{e,e'} \in e \setminus e'] \cdot \Pr[v_{e,e'} \in e \setminus e'] \\ &\quad + \Pr[e <_{P'} e' | v_{e,e'} \in e \cap e'] \cdot \Pr[v_{e,e'} \in e \cap e'] \\ &\quad + \Pr[e <_{P'} e' | v_{e,e'} \in e' \setminus e] \cdot \Pr[v_{e,e'} \in e' \setminus e] \\ &= 0 \cdot \frac{a}{a+b+c} + \frac{1}{2} \cdot \frac{c}{a+b+c} + 1 \cdot \frac{b}{a+b+c} \\ &= \frac{b+c/2}{a+b+c}. \end{aligned}$$

We will now use the following well-known inequality: for positive numbers $\alpha, \beta, \gamma, \delta$ such that $\alpha/\beta < \gamma/\delta$, we have $\frac{\alpha}{\beta} < \frac{\alpha+\gamma}{\beta+\delta} < \frac{\gamma}{\delta}$. If $c = 0$, we have $\Pr[e <_{P'} e'] = \frac{b}{a+b} \geq \frac{1}{1+\ell}$. Suppose $c > 0$, then we can write $\Pr[e <_{P'} e'] = \frac{b+c/2}{a+b+c} \geq \min\{\frac{b}{a+b}, \frac{c/2}{c}\}$. Considering that $\frac{b}{a+b}$ is minimized at $\frac{1}{1+\ell}$ subject to the constraints $1 \leq a, b \leq \ell$, we have the desired lower bound on $\Pr[e <_{P'} e']$ in both cases. $\square$

Therefore, we get a $\frac{1}{1+\ell}$ -balanced linear ordering oracle for the MLSC's scheduling reformulation, which ultimately gives us a $(2 - \frac{2}{1+\ell})$ -approximation algorithm for MLSC.

Integrality Gap for ℓ -uniform MLSC:

Next, we consider the relaxed linear program for MLSC on ℓ -uniform hypergraphs on n vertices, i.e., where each hyperedge has size ℓ . Here, variables $u_{e,t}$ represent

whether a hyperedge e is still uncovered (from MLSC perspective) until time t , and $x_{v,t}$ indicates a vertex v to be scheduled at time step t , when these are constrained to be integral.

$$\begin{aligned} (\text{MLSC-LP}) \quad & \text{minimize} \quad \sum_{e,t} u_{e,t} \\ & \text{subject to} \quad \sum_v x_{v,t} \leq 1, \quad \forall t \in \{1, \dots, n\}, \end{aligned} \quad (1)$$

$$u_{e,t} + \sum_{t' < t} x_{v,t'} \geq 1, \quad \forall v, e, t \text{ s.t. } v \in e, \quad (2)$$

$$u_{e,t}, x_{v,t} \geq 0, \quad \forall e, v, t. \quad (3)$$

The constraints (1) and (2), respectively, ensure that at most one vertex is scheduled during each time step, and every hyperedge remains unscheduled until all incident vertices are scheduled, i.e., $u_{e,t}$ is 0 only if all $v \in e$ are scheduled strictly before t .

First we show a lower bound of $2 - \frac{2}{1+\ell}$ on the integrality gap, matching the approximation factor of Theorem 3.

Proposition 2 *The integrality gap of the LP relaxation for MLSC on ℓ -uniform hypergraphs is at least $2 - \frac{2}{1+\ell}$.*

Proof Consider the complete ℓ -uniform hypergraph on n vertices. By a well-known binomial coefficient identity³, any ordering on the vertices gives the optimal objective to the combinatorial problem, which can be shown to be

$$\sum_{k=\ell}^n k \binom{k-1}{\ell-1} = \sum_{k=\ell}^n \ell \binom{k}{\ell} = \ell \binom{n+1}{\ell+1}.$$

For ℓ -uniform instances, the MLSC-LP objective can be upper bounded with a uniform fractional solution, i.e., $x_{v,t} = \frac{1}{n}$ and $u_{e,t} = 1 - \frac{t-1}{n}$ for all v, e , and t . It follows,

$$\sum_{e,t} u_{e,t} = |E| \cdot \left(\sum_{t=1}^n \left(1 - \frac{t-1}{n} \right) \right) = \binom{n}{\ell} \cdot \frac{n+1}{2} = \frac{\ell+1}{2} \binom{n+1}{\ell+1}.$$

Thus, this family of examples provides a lower bound of $\frac{2\ell}{\ell+1} = 2 - \frac{2}{1+\ell}$ for the integrality gap. $\square$

The integrality gap of MLSC-LP is therefore at least $2 - \frac{2}{1+\ell}$, but it can be more for certain families of graphs. We end this section by showing that the integrality gap of the MLSC-LP is exactly $2 - \frac{2}{1+\ell}$, for ℓ -uniform hypergraphs where the degree of each vertex is exactly d . We call these hypergraphs d -regular ℓ -uniform hypergraphs. We do not know if the integrality gap for non-regular uniform hypergraphs is strictly larger than $2 - \frac{2}{1+\ell}$.

³ The hockey-stick identity states for positive integers $\ell \leq n$, $\sum_{k=\ell}^n \binom{k}{\ell} = \binom{n+1}{\ell+1}$.

Proposition 3 Let H be any d -regular ℓ -uniform hypergraph with n vertices. Then the integrality gap for H is at most $2 - \frac{2}{1+\ell}$.

Proof We first show that the MLSC-LP has an optimal objective value $\frac{dn(n+1)}{2\ell}$ for any d -regular ℓ -uniform hypergraph with n vertices. For all fixed $1 \leq t \leq n$, summing over constraints 2 for all $e \in E$ and all $v \in e$, and we have:

$$\ell \sum_e u_{e,t} = \sum_e \sum_{v \in e} u_{e,t} \quad (4)$$

$$\stackrel{(2)}{\geq} \sum_e \sum_{v \in e} \left(1 - \sum_{t' < t} x_{v,t'} \right) \quad (5)$$

$$= dn - d \sum_{t' < t} \sum_v x_{v,t'} \quad (6)$$

$$\stackrel{(1)}{\geq} dn - d(t-1), \text{ for all } 1 \leq t \leq n. \quad (7)$$

Now summing over (7) for t from 1 to n we have:

$$\sum_{e,t} u_{e,t} \geq \frac{1}{\ell} \sum_{t=1}^n (dn - d(t-1)) = \frac{dn(n+1)}{2\ell}.$$

It is easy to see that this objective value is achieved by letting $x_{v,t} = \frac{1}{n}$ and $u_{e,t} = 1 - \frac{t-1}{n}$ for all e, v, t , as this makes all inequalities satisfied with equality.

Now consider the MLSC problem. Using randomized rounding (e.g., [53]), we will show there exists a permutation with objective value at most $2 - \frac{2}{1+\ell}$ of the LP optimal value. Let π be a uniformly random permutation of vertices, i.e. $\pi(v) = k$ with probability $1/n$ for all $1 \leq k \leq n$. Then, for any hyperedge e we have

$$\mathbb{E} [\max\{\pi(v), v \in e\}] = \frac{1}{\binom{n}{\ell}} \sum_{k=\ell}^n k \binom{k-1}{\ell-1} = \frac{\ell \binom{n+1}{\ell+1}}{\binom{n}{\ell}} = \frac{\ell(n+1)}{\ell+1}.$$

Thus, by linearity of expectation, the expectation of the objective value for MLSC is

$$\frac{dn}{\ell} \mathbb{E} [\max\{\pi(v), v \in e\}] = \frac{dn(n+1)}{\ell+1}.$$

Therefore, there exists a permutation with objective value at most $\frac{dn(n+1)}{\ell+1}$, which is $2 - \frac{2}{1+\ell}$ of the LP optimal value. $\square$

7.2 Polynomial solvable instances for MLVC and MSVC

We next discuss classes of instances of MLVC and MSVC that can be solved in polynomial time. The following theorem relates the objective value of MLA with MLVC for the family of regular graphs.

Theorem 12 *Let G be a d -regular graph on n vertices. For any labeling $\sigma \in S_n$, we have*

$$2 \cdot \sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\} = d \binom{n+1}{2} + \sum_{(x,y) \in E(G)} |\pi(x) - \pi(y)|.$$

Proof We have that,

$$\begin{aligned} \sum_{(x,y) \in E(G)} |\pi(x) - \pi(y)| &= \sum_{(x,y) \in E(G)} [2 \cdot \max\{\pi(x), \pi(y)\} - \pi(x) - \pi(y)] \\ &= - \sum_{v \in V(G)} \pi(v)d + 2 \cdot \sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\} \\ &= -d \sum_{i=1}^n i + 2 \cdot \sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\} \\ &= -d \binom{n+1}{2} + 2 \cdot \sum_{(x,y) \in E(G)} \max\{\pi(x), \pi(y)\}. \end{aligned}$$

□

By Theorem 12, we have that MLA and MLVC for regular graphs are equivalent in decision form. As the family of regular graphs is closed under graph complements, we also have by Theorem 9 that MSVC and MLVC for the family of regular graphs are equivalent in decision form as well. Thus we have the following,

Corollary 6 *For the family of regular graphs, MLA, MLVC, and MSVC are equivalent in decision form.*

As an illustration of the utility of Theorem 12, we introduce Hamming graphs $H(d, c)$, which are obtained from d Cartesian graph products of the complete graph K_c . Motivated by designing error correcting codes, Harper [17] solved the MLA problem for hypercubes, i.e. $H(d, 2)$ where d is any positive integer. Later, Nakano [54] generalized this result to all Hamming graphs $H(d, c)$ where d and c are positive integers. As Hamming graphs are regular, we have the following corollary of Theorem 12.

Corollary 7 *MLVC is polynomial time solvable for Hamming graphs.*

The literature for the MLA problem is vast and many other instances of regular graphs have been previously solved. Thus Theorem 12, while simple, provides a powerful tool for providing polynomial time algorithms for many families of regular graphs. Some of these families of graphs include toroidal grids [55], complete p -partite graphs [56], and de Bruijn graphs of order 4 [57]. This list is by no means exhaustive, and we refer the reader to the following surveys for further reading [9–12]. Furthermore by Theorem 9, the complements of these families also have polynomial time algorithms for the MSVC problem.

8 Improved approximation for monotone submodular MLOP

Monotone submodular MLOP was introduced by Iwata et al. [1], where the authors also provided a factor $(2 - \frac{2}{1+|E|})$ -approximation algorithm using the Lovász extension of submodular functions. Fokkink et al. [16] studied the submodular search problem, which generalizes monotone submodular MLOP, and gave an approximation factor based on the total curvature of the submodular function. It was not known if a tighter approximation was possible. They considered the greedy contraction of the principal partition induced by the submodular function, an idea that has been used as early as 1992 by Pinar [15]. In this section, we give a different analysis to the same algorithm and improve the approximation factor to

$$2 - \frac{1 + \ell_f}{1 + |E|} \text{ where } \ell_f = \frac{f(E)}{\max_{x \in E} f(\{x\})}.$$

Our result can be applied to special cases including matroid and graphic matroid MLOP. For general matroid MLOP, our approximation factor is $2 - \frac{1+r(E)}{1+|E|}$, which is strictly smaller than 2 when $r(E) = \Omega(|E|)$ (e.g., graphic matroid on sparse graphs). Note that both approximation factors given by [16] based on total curvature and [1] based on Lovász extension are asymptotically 2 for all non-trivial instances of matroid MLOP.

Throughout this section, let E be a nonempty set of size m and $f : 2^E \rightarrow \mathbb{R}$ be a normalized ($f(\emptyset) = 0$) monotone submodular set function. Without loss of generality, we can also assume that the maximum minimizer of the submodular function is the empty set,⁴ i.e., $f(S) > 0$ for all $S \neq \emptyset$. Recall from Sect. 4 that the steepness of a set function f is defined as $\kappa_f = \max_{x \in E} f(\{x\})$, and linearity of f is $\ell_f = \frac{f(E)}{\kappa_f}$. By submodularity and monotonicity of f , for all $S \subseteq T$ we have $f(T) \leq f(S) + \kappa_f |T \setminus S|$.

Note for any non-trivial (i.e., $f(E) > 0$) normalized monotone submodular function $f : 2^E \rightarrow \mathbb{R}$, we have $1 \leq \ell_f \leq |E|$. Both of the bounds are tight, as the lower bound $\ell_f = 1$ is attained when f is the rank function on a graphic matroid with 2 vertices and $|E|$ parallel edges between them, while the upper bound $\ell_f = |E|$ is attained when $f(S) = |S|$ for all $S \subseteq E$. Thus, the linearity ℓ_f is a measure of how uniform and linear a submodular function is. The function will have high linearity if each singleton has approximately same function value, and the function is approximately linear, i.e., all submodular relations $f(S) + f(T) \geq f(S \cap T) + f(S \cup T)$ are close to being tight. In the special case where $f(S)$ is the rank function of some matroid, we have $\kappa_f = 1$ and $\ell_f = f(E)$ (the rank of the matroid).

In this section, we show a $(2 - \frac{1+\ell_f}{1+|E|})$ -approximation factor to monotone submodular MLOP using any linear extension of the principal partition with respect to the submodular function. Recall that a principal partition is a set of nested sets $\emptyset = \Pi_0 \subsetneq \dots \subsetneq \Pi_s = E$ ($s \geq 1$) and a set of critical values $\lambda_0 < \lambda_1 < \dots < \lambda_{s+1}$,

⁴ We can simply contract the maximal minimizer U . The elements in U must appear (in any order) before the other elements $E \setminus U$ in any optimal solution for MLOP on monotone submodular functions (see Appendix A.1).




Fig. 2 Diagram of our lower and upper bounds in grey as well as the optimal solution in red. The black circles represent the principal partitions

such that for all $0 \leq i \leq s$, Π_i is the unique maximal optimal solution to $\min_{X \subseteq E} f(X) - \lambda|X|$, for all $\lambda \in (\lambda_i, \lambda_{i+1})$ (Sect. 4).

Theorem 13 *Let $\{\Pi_i\}_{0 \leq i \leq s}$ be the principal partition of a non-trivial monotone submodular function $f : 2^E \rightarrow \mathbb{R}$ satisfying $f(\emptyset) = 0$. Let $\kappa_f = \max_{x \in E} f(\{x\})$ and $\ell_f = \frac{f(E)}{\kappa_f}$. Let $\sigma \in \mathcal{S}_E$ be any linear extension of the principal partition, i.e., $E|_{|\Pi_i|, \sigma} = \Pi_i$ for all $1 \leq i \leq s$. Then, the MLOP objective value of σ is at most factor $2 - \frac{1+\ell_f}{1+|E|}$ of the optimal solution.*

Since $1 \leq \ell_f \leq |E|$, our result is a refinement on the $2 - \frac{2}{1+|E|}$ factor approximation of monotone submodular MLOP in [1]. For the lower bound, our key lemma (Lemma 7) is a more general version of the well-known fact (see [14, 16, 58]) that any member of the principal partition Π_i is the “sparsest” subset⁵ in the Π_{i-1} -contracted submodular function $f|_{\Pi_{i-1}}$, i.e., $\frac{f(S) - f(\Pi_{i-1})}{|S| - |\Pi_{i-1}|} \geq \frac{f(\Pi_i) - f(\Pi_{i-1})}{|\Pi_i| - |\Pi_{i-1}|}$ for all $S \supsetneq \Pi_{i-1}$. We show, in Lemma 7, that this algebraic statement holds for all subsets S where $|S| \neq |\Pi_{i-1}|$, allowing us to lower bound the MLOP value of an arbitrary chain. For the upper bound, we consider any MLOP solution that is a linear extension of the principal partitions. The increase of the function value can be upper bounded using κ_f , the linearity parameter of submodular function f , as well as the function value at the principal partitions.

See Fig. 2 for an illustration. The horizontal axis denotes the sizes of subsets appearing in an MLOP solution, and the vertical axis denotes the cost that these subsets incur

⁵ In our notation, “sparsest” subsets are the ones minimizing $\frac{f(S)}{|S|}$. Such subsets are referred to as being “densest” by Fokkink et al. [16], as they maximize $\frac{|S|}{f(S)}$.

in the MLOP objective. The coordinates of the black circles are the sizes and costs of the principal partitions. Between two adjacent black circles in the figure, the lower bound is the linear segment joining them, and the upper bound is formed using two linear segments, the first with positive slope κ_f and the second with slope 0. The red points represent subsets in an optimal MLOP solution, and we show that they always lie inside the triangular shaded regions formed by the lower and upper bounds. In particular, the principal partitions must appear in any optimal MLOP solution, which is also a consequence of Theorem 1 in [16].

The proofs for the lower and upper bounds are highly algebraic, and a lot of calculations are deferred to the appendix. One of the challenges is that for the upper bound, the difference between function values of two adjacent subsets in the principal partition may not be an integer multiple of κ_f , thus additional steps are needed to deal with rounding as the upper bound approaches each horizontal segment.

8.1 Lower and upper bound on MLOP objective value

Consider a monotone submodular function $f : 2^E \rightarrow \mathbb{R}$ satisfying $f(S) = 0$ if and only if $S = \emptyset$, its principal partition $\{\Pi_i\}_{0 \leq i \leq s}$ and the corresponding critical values $\{\lambda_i\}_{1 \leq i \leq s}$ (Sect. 4). The following lemma gives the relationship between the critical values and the principal partition [14].

Lemma 6 *The principal partition $\{\Pi_i\}_{0 \leq i \leq s}$ and corresponding critical values $\{\lambda_i\}_{1 \leq i \leq s}$ satisfy the following relation:*

$$\lambda_i = \frac{f(\Pi_i) - f(\Pi_{i-1})}{|\Pi_i| - |\Pi_{i-1}|}, \text{ for all } 1 \leq i \leq s.$$

Furthermore, Π_{i-1} and Π_i are the unique minimal and maximal minimizers of $\min_{X \subseteq E} f(X) - \lambda_i |X|$.

We include a proof of Lemma 6 in Appendix A.2 for completeness. It simply uses the definition of the principal partition and submodularity of the set function.

The following lemma lower gives an lower bound on the function value of any subset. As mentioned before, this lemma is more general than stating that $\Pi_i \setminus \Pi_{i-1}$ is the unique maximal sparsest subset with respect to $f_{|\Pi_{i-1}|}$, the Π_{i-1} -contracted submodular function.

Lemma 7 *Let $f : 2^E \rightarrow \mathbb{R}$ be a normalized monotone submodular function with $f(S) > 0$ if $S \neq \emptyset$, and principal partition $\{\Pi_i\}_{0 \leq i \leq s}$. Let $S \subseteq E$, then*

$$\frac{f(S) - f(\Pi_{i-1})}{|S| - |\Pi_{i-1}|} \geq \frac{f(\Pi_i) - f(\Pi_{i-1})}{|\Pi_i| - |\Pi_{i-1}|},$$

for all i such that $|S| \neq |\Pi_{i-1}|$.

Proof One can simply fix an arbitrary critical value λ_i , and use the fact that $f(\Pi_i) - \lambda_i |\Pi_i| \leq f(S) - \lambda_i |S|$ for any $S \subseteq E$. Rearranging terms we get $f(S) - f(\Pi_{i-1}) \geq$

$\lambda_i(|S| - |\Pi_{i-1}|)$. Substituting the value of $\lambda_i = \frac{f(\Pi_i) - f(\Pi_{i-1})}{|\Pi_i| - |\Pi_{i-1}|}$ (Lemma 6) gives us the desired result. $\square$

Using the above lemma, we can sum up appropriate bounds for each subset $E_{i,\sigma}$ for any ordering σ , and obtain the following lower bound for monotone submodular MLOP. The proof after summation is purely algebraic manipulation, which is deferred to appendix.

Proposition 4 *Let $f : 2^E \rightarrow \mathbb{R}$ be a normalized monotone submodular function with $f(S) > 0$ if $S \neq \emptyset$, and principal partition $\{\Pi_i\}_{0 \leq i \leq s}$. Let $\sigma \in \mathcal{S}_E$, then*

$$\sum_{k=1}^m f(E_{k,\sigma}) \geq \frac{1}{2}(|E| + 1)f(E) - \frac{1}{2} \sum_{i=1}^s (f(\Pi_i)|\Pi_{i-1}| - f(\Pi_{i-1})|\Pi_i|) > 0.$$

Proof The proof is deferred to Appendix A.3. $\square$

For the upper bound, we require that the chain must contain all sets in principal partition, i.e. $E_{|\Pi_i|,\sigma} = \Pi_i$ for all i . We use the fact that each added element into the subset can increase the function value by at most κ_f to upper bound the function value of remaining sets. Pictorially, if we start from Π_{i-1} , the upper bound starts at $f(\Pi_{i-1})$ and has slope κ_f , until it reaches $f(\Pi_i)$ where it remains flat until Π_i (refer to Fig. 2). Also note that the increase of function value is integer multiple of κ_f without additional analysis, and the rounding as function value approaches $f(\Pi_i)$ has to be taken care of.

Proposition 5 *Let $f : 2^E \rightarrow \mathbb{R}$ be a normalized monotone submodular function with $f(S) > 0$ if $S \neq \emptyset$, and principal partition $\{\Pi_i\}_{0 \leq i \leq s}$. Let $\sigma \in \mathcal{S}_E$ be such that $E_{|\Pi_i|,\sigma} = \Pi_i$ for all $1 \leq i \leq s$. Then the MLOP objective value for f with permutation σ is at most*

$$f(E)|E| - \frac{f(E)^2}{2\kappa_f} + \frac{f(E)}{2} - \sum_{i=1}^s (f(E) - f(\Pi_i))(|\Pi_i| - |\Pi_{i-1}|) + \sum_{i=1}^s \frac{f(\Pi_{i-1})(f(\Pi_i) - f(\Pi_{i-1}))}{\kappa_f}.$$

Proof The proof is deferred to Appendix A.4. $\square$

8.2 Proof of Improved Approximation for MLOP in Theorem 13

Both Proposition 4 and 5 together allow us to prove Theorem 13. Our goal is to show the upper bound obtained from Proposition 5 is at most $2 - \frac{1+\ell_f}{1+|E|}$ the lower bound obtained from Proposition 4, thus showing for a $\sigma \in \mathcal{S}_E$ such that $E_{|\Pi_i|,\sigma} = \Pi_i$ for all $1 \leq i \leq s$, is our desired approximation for monotone submodular MLOP. First comparing the non-summation terms in Proposition 4 and 5 we have

$$\frac{f(E)|E| - \frac{f(E)^2}{2\kappa_f} + \frac{f(E)}{2}}{\frac{1}{2}(|E| + 1)f(E)} = \frac{2|E| - \frac{f(E)}{\kappa_f} + 1}{|E| + 1} = 2 - \frac{1 + \ell_f}{1 + |E|}.$$

To deal with the remaining summation terms, it suffices to prove that

$$\begin{aligned} \sum_{i=1}^s (f(\Pi_i)|\Pi_{i-1}| - f(\Pi_{i-1})|\Pi_i|) &\leq \sum_{i=1}^s (f(E) - f(\Pi_i))(|\Pi_i| - |\Pi_{i-1}|) \\ &\quad - \sum_{i=1}^s \frac{f(\Pi_{i-1})(f(\Pi_i) - f(\Pi_{i-1}))}{\kappa_f}, \end{aligned}$$

i.e., the decrease of upper bound from non-summation terms is at least twice the decrease of lower bound from non-summation terms. To make computation easier we rewrite the terms using differential notation. For all $1 \leq i \leq s$, let $\delta_i = f(\Pi_i) - f(\Pi_{i-1})$ and $\Delta_i = |\Pi_i| - |\Pi_{i-1}|$. By definition of κ_f , we have $0 \leq \delta_i \leq \kappa_f \Delta_i$. Note that $f(\Pi_i) = \sum_{j=1}^i \delta_j$ and $|\Pi_i| = \sum_{j=1}^i \Delta_j$. Furthermore, we have $f(\Pi_i)|\Pi_{i-1}| - f(\Pi_{i-1})|\Pi_i| = |\Pi_i|\delta_i - f(\Pi_i)\Delta_i$. Thus, the statement to be proved can be rewritten as

$$\sum_{i=1}^s \left(\delta_i \sum_{j=1}^i \Delta_j - \Delta_i \sum_{j=1}^i \delta_j \right) \leq \sum_{i=1}^s \left(\Delta_i \sum_{j=i+1}^s \delta_j - \frac{\delta_i}{\kappa_f} \sum_{j=1}^{i-1} \delta_j \right).$$

Suppose $s = 1$, then both sides of this inequality are equal to zero.

Thus, we may assume $s \geq 2$. Rearranging terms, for the left hand side we have

$$\sum_{i=1}^s \left(\delta_i \sum_{j=1}^i \Delta_j - \Delta_i \sum_{j=1}^i \delta_j \right) = \sum_{i=1}^{s-1} \sum_{j>i} \delta_j \Delta_i - \delta_i \Delta_j,$$

and the second part of right hand side can be rewritten as

$$\sum_{i=1}^s \frac{\delta_i}{\kappa_f} \sum_{j=1}^{i-1} \delta_j = \frac{1}{\kappa_f} \sum_{j=1}^{s-1} \sum_{i>j} \delta_i \delta_j = \frac{1}{\kappa_f} \sum_{i=1}^{s-1} \sum_{j>i} \delta_i \delta_j,$$

after exchanging summation order and changing variable names. As $\delta_j \leq \kappa_f \Delta_j$ and hence $\sum_{i=1}^{s-1} \sum_{j>i} \delta_j \Delta_i - \delta_i \Delta_j \leq \sum_{i=1}^{s-1} \sum_{j>i} \delta_j \Delta_i - \delta_i \frac{\delta_j}{\kappa_f}$, we have the inequality holds and thus, the proof is finished.

Recall from Theorem 7 that principal partitions $\{\Pi_i\}_{0 \leq i \leq s}$ can be found in polynomial time. Thus, we have the following:

Theorem 6 *Let $f : 2^E \rightarrow \mathbb{R}$ be a non-trivial, normalized and monotone submodular function. There exists a factor $(2 - \frac{1+\ell_f}{1+|E|})$ -approximation algorithm to MLOP with $f(\cdot)$ in polynomial time, where $\ell_f = \frac{f(E)}{\max_{x \in E} f(\{x\})}$.*

Note that our analysis works for any linear extension to the partial order on subsets induced by the principal partition. It is unclear how this analysis can be extended to more structured linear extensions. We now discuss a special case of Theorem 13, when f is the rank function of a matroid M . Since in this case, $\ell_f = f(E)$, we get:

Corollary 8 *Let $M = (E, r)$ be a matroid on ground set E with rank function r . There exists a factor $(2 - \frac{1+r(E)}{1+|E|})$ -approximation algorithm to matroid MLOP on M in polynomial time.*

For graphic matroids, this improves upon the 2-factor approximation when graph is connected and has a linear number of edges. For instance, for connected d -regular graphs with vertex set V , the approximation factor is $2 - \frac{2|V|}{2+d|V|}$, which is asymptotically $2 - \frac{2}{d}$.

8.3 Application to minimum latency set cover (MLSC)

Recall that in Sect. 7.1 we present a randomized factor $(2 - \frac{2}{1+\ell})$ -approximation algorithm for MLSC, where ℓ is the size of largest hyperedge. For the special case of MLVC the factor is $\frac{4}{3}$. In this section we make the observation that MLSC is an instance of monotone submodular MLOP, and use Theorem 13 to show that there exists a deterministic factor $(2 - \frac{\Delta+|E|}{\Delta(1+|V|)})$ -approximation algorithm for MLSC, where Δ is the maximum degree of the hypergraph $H = (V, E)$. Note that for ℓ -uniform hypergraphs this bound is never better than the one obtained in Sect. 7.1.

Recall that in MLSC, we are given a hypergraph $H = (V, E)$ with the objective

$$\min_{\pi \in \mathcal{S}_V(H)} \sum_{e \in E} \max_{v \in e} \pi(v).$$

In other words, we minimize over all permutations of the vertices, where the cost of each hyperedge is the maximum label of all vertices in it. Throughout this section we let $n = |V|$ denote the number of vertices.

For a fixed $\pi \in \mathcal{S}_V$, its reverse permutation is defined as $\pi'(v) = n + 1 - \pi(v)$ for all $v \in V$. We now prove that the MLSC value with π is the same as the MLOP value with π' on a particular monotone submodular function, which shows that MLSC is an instance of monotone submodular MLOP.

Proposition 6 *For a fixed hypergraph $H = (V, E)$, let f be the set function on V such that for all $S \subseteq V$, $f(S) = |\{e \in E : S \cap e \neq \emptyset\}|$. Then f is a monotone submodular function satisfying $f(\emptyset) = 0$. Furthermore, for all $\pi \in \mathcal{S}_V$ we have*

$$\sum_{e \in E} \max_{v \in e} \pi(v) = \sum_{i=0}^n f(V_i, \pi'),$$

where $\pi' \in \mathcal{S}_V$ is given by $\pi'(v) = n + 1 - \pi(v)$.

Proof It is straightforward to verify that f is monotone and $f(\emptyset) = 0$. For submodularity, for all $S, T \subseteq V$, observe that

$$\begin{aligned} f(S) + f(T) - f(S \cup T) - f(S \cap T) \\ = |\{e \in E : e \cap S \neq \emptyset, e \cap T \neq \emptyset, e \cap S \cap T = \emptyset\}| \geq 0. \end{aligned}$$

Now for all $0 \leq k \leq n$ let $T_k = \{e \in E : \exists v \in e, \pi'(v) \leq k\}$. Then it is straightforward to verify that for all $0 \leq k \leq n$, $f(V_{k,\sigma}) = |T_k|$ and furthermore for all $e \in E$, $|\{k : e \in T_k\}| = \max_{v \in e} \pi(v)$. Therefore we have

$$\sum_{e \in E} \max_{v \in e} \pi(v) = \sum_{i=0}^n |T_i| = \sum_{i=0}^n f(V_{i,\pi'}).$$

□

Using Theorem 13, we obtain the following approximation algorithm for MLSC where the factor is based on maximum degree of the hypergraph. Note in this case $\kappa_f = \Delta(H) = \max_{e \in E} |e|$ is the maximum degree of hypergraph H .

Corollary 2 *There is a deterministic factor $(2 - \frac{\Delta+|E|}{\Delta(1+|V|)})$ -approximation algorithm for MLSC, where Δ is the maximum degree of hypergraph $H = (V, E)$.*

For comparison, in Sect. 7 we presented a randomized scheduling-based approximation algorithm for MLSC within factor $2 - \frac{2}{1+\ell}$, where $\ell = \max_{e \in E} |e|$ is the size of largest hyperedge. The algorithm presented in this section is deterministic, but for uniform hypergraphs this bound is never better than the randomized algorithm based on scheduling.

9 Future directions

We conclude this work by presenting a list of open questions that stem from this work.

In Sects. 5 and 6 we investigated the hardness of restrictions of MLOP. In particular, we showed that graphic matroid MLOP is NP-hard. In Sect. 6.4, we saw how matroid MLOP can be viewed as an optimization problem over the bases of the matroid. However, even when a basis is fixed, the corresponding ordering problem on the ground set of elements can be non-trivial. In particular, in the context for graphic matroid MLOP on a connected graph G , consider the following optimization problem,

$$\min_{\sigma \in S_T} \sum_{e \in E(G) \setminus T} \max\{\sigma(e') : e' \in C(T, e) - e\},$$

where T is a given (fixed) edge set of a spanning tree of G , and $C(T, e)$ denotes the fundamental circuit with respect to T and e . We saw implicitly in the reduction of MLVC to graphic matroid MLOP (Theorem 8), that MLVC reduces to this problem when T is the star graph (thus, this is NP-hard). We prove in Proposition 1, that if we

allow the choice of the spanning tree T to vary over all spanning trees of G , then the above problem is equivalent to graphic matroid MLOP. Thus, when T is fixed, this problem can be viewed as a “fixed-basis” restriction of graphic matroid MLOP.

Open question 1. Given a graph G , a spanning tree T and integer k , consider the problem of whether there exists a permutation σ of $E(T)$ such that

$$\sum_{e \in E(G) \setminus E(T)} \max\{\sigma(e') : e' \in C(T, e) - e\} \leq k.$$

For what families of trees is this problem NP-hard?

This problem is known to be NP-hard only when T is a star graph, and remains open for other simple families of trees, such as in the case where T is a path.

In Sect. 7, we showed that MLSC can be $(2 - \frac{2}{1+\ell})$ -approximated using randomized scheduling techniques. Furthermore, we showed for ℓ -uniform regular hypergraphs, MLSC can be $(2 - \frac{2}{1+\ell})$ -approximated using an LP relaxation. This question for general ℓ -uniform hypergraphs remains open.

Open question 2. Does solving the LP relaxation provide an approximation guarantee for MLSC on ℓ -uniform hypergraphs by a factor of $2 - \frac{2}{1+\ell}$?

In Sect. 7.2, we show that the MLVC, MSVC, and MLA are all equivalent problems in decision form for regular graphs. Using techniques similar to Theorem 9, we can show that the optimal value for all three problems are related by linear shifts. It is known that MSVC on regular graphs can be $4/3$ -approximated (see [26]), but we have not found a formal proof that this problem is NP-hard. Thus, the following question remains open, to the best of our knowledge.

Open question 3. Are MLA, MLVC, and MSVC NP-hard for the family of simple regular graphs?

In Sect. 8, we show that monotone submodular MLOP can be approximated within factor $2 - \frac{1+\ell_f}{1+|E|}$, using principal partitions. A related open question is to develop algorithms when the principal partitions are trivial, i.e., $f(S)|E| \geq f(E)|S|$ for all $S \subseteq E$. In this case, the principal partition-based algorithm studied by Fokkink et al. [16] (and by us) will simply output an arbitrary solution.

Open question 4. Do there exist better polynomial time approximation algorithms for monotone submodular MLOP in the case where the function f satisfies $f(\emptyset) = 0$ and $f(S)|E| \geq f(E)|S|$ for all $S \subseteq E$?

In the scope of symmetric submodular MLOP, the current best known approximation factor for the special case MLA is polylogarithmic in the size of the graph, i.e., $O(\sqrt{\log n \log \log n})$, given by Feige and Lee [21]; see also Charikar et al. [22]. For the more general problem of symmetric submodular MLOP, there is currently no known efficient approximation algorithm better than $O(|E|)$.

Open question 5. Can symmetric submodular MLOP over a ground set E be approximated to a factor better than $O(|E|)$?

Acknowledgements The authors would like to thank Nikhil Bansal for insightful discussions on Theorem 3, i.e., approximation of MLVC, László Végh for remarks on principal partitions, and Jai Moondra for comments on a preliminary version of this paper. The authors would also like to thank the anonymous referees for numerous feedback and suggestions that are of great value, in particular for pointing out that MLSC is an instance of monotone submodular MLOP. The second author would like to acknowledge support from

NSF CRII-1850182. The last author would like to acknowledge support by an NSF Graduate Research Fellowship under Grant No. DGE-165004. The first and third author were supported by ARC-TRIAD Student Fellowships. The fourth author acknowledges support by NSF grant DMS-2151283.

Funding Open Access funding provided by the MIT Libraries.

Declarations

Competing Interest The authors declare that they have no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Proofs for Sect. 8

A.1 Reduction to the case where $f(S) > 0$ for all $S \neq \emptyset$

Here we formally prove the statement that in monotone submodular MLOP where $f(\emptyset) = 0$, all elements with weight zero must appear (in any order) in the beginning of any optimal MLOP solution.

Lemma 8 *Let $f : 2^E \rightarrow \mathbb{R}$ be a monotone submodular function with $f(\emptyset) = 0$. Then there exists a unique maximal set U satisfying $f(U) = 0$. Furthermore, any optimal MLOP solution $\sigma \in \mathcal{S}_E$ on f must have U as prefix, i.e., $E_{|U|, \sigma} = U$.*

Proof Let U be a maximal set with $f(U) = 0$ and let U' be any subset of E such that $f(U') = 0$. We claim that $U' \subseteq U$ which would show that U is the unique maximal set. From submodularity we have $f(U \cup U') \leq f(U) + f(U') - f(U \cap U') = -f(U \cap U')$. From monotonicity, $f(U \cap U') \geq f(\emptyset) = 0$, therefore $f(U \cup U') = 0$. Since U is maximal, we have $U \cup U' = U$, which implies $U' \subseteq U$.

For the sake of contradiction, let $\sigma \in \mathcal{S}_E$ be any optimal MLOP solution where $E_{|U|, \sigma} \neq U$. Let the elements in $E = \{e_1, \dots, e_n\}$ be ordered such that $\sigma(e_j) = j$. Let $i < |U|$ be the smallest index such that $e_i \notin U$ and let j be the smallest index such that $j > i$ and $e_j \in U$. Consider a new permutation $\sigma' \in \mathcal{S}_E$ where we move e_j just before e_i and keep everything else unchanged. That is,

$$\sigma'(e) = \begin{cases} \sigma(e) & \text{if } \sigma(e) < \sigma(e_i), \\ \sigma(e_i) & \text{if } e = e_j, \\ \sigma(e) + 1 & \text{if } \sigma(e) \geq \sigma(e_i) \text{ and } \sigma(e) < \sigma(e_j), \\ \sigma(e) & \text{if } \sigma(e) \geq \sigma(e_j). \end{cases}$$

We have $E_{k,\sigma} = E_{k,\sigma'}$ for all $k < i$ and $k \geq j$. For all $i \leq k < j$, we have that $E_{k,\sigma'} = E_{k,\sigma} - e_k + e_j$. As $f(\{e_j\}) = 0$, we have by submodularity, $f(E_{k,\sigma'}) \leq f(E_{k,\sigma} - e_k) + f(\{e_j\}) = f(E_{k,\sigma} - e_k)$. Hence by monotonicity, $f(E_{k,\sigma'}) \leq f(E_{k,\sigma})$ for any $i \leq k < j$. Furthermore, we have $f(E_{i,\sigma}) > 0$ since $e_i \notin U$. As $f(E_{i,\sigma'}) = 0$, this shows that changing σ to σ' strictly decreases the MLOP value, contradicting the optimality of σ . Thus any optimal MLOP solution must contain U as a prefix. $\square$

A.2 Proof for Lemma 6

Fix $1 \leq i \leq s$, consider $\min_{X \subseteq E} f(X) - \lambda_i |X|$. From the definition of principal partitions, Π_i is a minimizer of $f(X) - \lambda |X|$ for all $\lambda \in (\lambda_i, \lambda_{i+1})$, and Π_{i-1} is a minimizer of $f(X) - \lambda |X|$ for all $\lambda \in (\lambda_{i-1}, \lambda_i)$.

Thus for all $X \subseteq E$ and sufficiently small $\varepsilon > 0$, $f(\Pi_i) - (\lambda_i + \varepsilon)|\Pi_i| \leq f(X) - (\lambda_i + \varepsilon)|X|$ and $f(\Pi_i) - \lambda_i |\Pi_i| \leq f(X) - \lambda_i |X| + \varepsilon |\Pi_i|$. Letting $\varepsilon \rightarrow 0$, we have $f(\Pi_i) - \lambda_i |\Pi_i| \leq f(X) - \lambda_i |X|$ for all $X \subseteq E$.

Furthermore, for all $X \subseteq E$ and sufficiently small $\varepsilon > 0$, $f(\Pi_{i-1}) - (\lambda_i - \varepsilon)|\Pi_{i-1}| \leq f(X) - (\lambda_i - \varepsilon)|X|$ and $f(\Pi_{i-1}) - \lambda_i |\Pi_{i-1}| \leq f(X) - \lambda_i |X| + \varepsilon |\Pi_{i-1}|$. Letting $\varepsilon \rightarrow 0$, we have $f(\Pi_{i-1}) - \lambda_i |\Pi_{i-1}| \leq f(X) - \lambda_i |X|$ for all $X \subseteq E$.

Combining these two statements we get $f(\Pi_{i-1}) - \lambda_i |\Pi_{i-1}| = f(\Pi_i) - \lambda_i |\Pi_i|$. Solving for λ_i we obtain the desired result. The proof that Π_{i-1} and Π_i are the unique minimal and maximal minimizers of $\min_{X \subseteq E} f(X - \lambda_i |X|)$ can be found in [14].

A.3 Proof for Proposition 4

We first claim that for all $S \subseteq E$ and all i , we have

$$f(S) \geq f(\Pi_{i-1}) + \frac{f(\Pi_i) - f(\Pi_{i-1})}{|\Pi_i| - |\Pi_{i-1}|} (|S| - |\Pi_{i-1}|).$$

If $|S| = |\Pi_{i-1}|$ then $f(S) \geq f(\Pi_{i-1})$ from the definition of principal partitions. If $|S| \neq |\Pi_{i-1}|$, this follows from rearranging terms from Lemma 7.

Note $\sum_{k=1}^m f(E_{k,\sigma}) = \sum_{i=1}^s \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma})$. For each $E_{j,\sigma} \subseteq E$, we apply the bound above using the unique i where $|\Pi_{i-1}| < j \leq |\Pi_i|$. Since $f(\Pi_i) > f(\Pi_{i-1})$, this lower bound is strictly positive, and therefore the summation in the end is also strictly positive.

The summation of $f(E_{j,\sigma})$ for all $|\Pi_{i-1}| < j \leq |\Pi_i|$ is then given by

$$\begin{aligned} \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) &\geq f(\Pi_{i-1})(|\Pi_i| - |\Pi_{i-1}|) + \frac{f(\Pi_i) - f(\Pi_{i-1})}{|\Pi_i| - |\Pi_{i-1}|} \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} (j - |\Pi_{i-1}|) \\ &= f(\Pi_{i-1})(|\Pi_i| - |\Pi_{i-1}|) + \frac{1}{2} (f(\Pi_i) - f(\Pi_{i-1})) (|\Pi_i| - |\Pi_{i-1}| + 1) \\ &= \frac{1}{2} (f(\Pi_i)|\Pi_i| - f(\Pi_{i-1})|\Pi_{i-1}| + f(\Pi_i) - f(\Pi_{i-1})) \\ &\quad + \frac{1}{2} (f(\Pi_{i-1})|\Pi_i| - f(\Pi_i)|\Pi_{i-1}|) \end{aligned}$$

The terms are grouped in this way so that the first part telescopes and second part does not. It follows that

$$\begin{aligned}
 \sum_{k=1}^m f(E_{k,\sigma}) &= \sum_{i=1}^s \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) \\
 &\geq \sum_{i=1}^s \left[\frac{1}{2} (f(\Pi_i)|\Pi_i| - f(\Pi_{i-1})|\Pi_{i-1}| + f(\Pi_i) - f(\Pi_{i-1})) \right. \\
 &\quad \left. + \frac{1}{2} (f(\Pi_{i-1})|\Pi_i| - f(\Pi_i)|\Pi_{i-1}|) \right] \\
 &= \frac{1}{2} (f(\Pi_s)|\Pi_s| - f(\Pi_0)|\Pi_0| + f(\Pi_s) - f(\Pi_0)) \\
 &\quad - \frac{1}{2} \sum_{i=1}^s (f(\Pi_i)|\Pi_{i-1}| - f(\Pi_{i-1})|\Pi_i|) \\
 &= \frac{1}{2} (|E| + 1)f(E) - \frac{1}{2} \sum_{i=1}^s (f(\Pi_i)|\Pi_{i-1}| - f(\Pi_{i-1})|\Pi_i|).
 \end{aligned}$$

A.4 Proof for Proposition 5

For indices i and j such that $1 \leq i \leq s$ and $|\Pi_{i-1}| + 1 \leq j \leq |\Pi_i|$ we have $f(E_{j,\sigma}) \leq \min\{f(\Pi_{i-1}) + \kappa_f(j - |\Pi_{i-1}|), f(\Pi_i)\}$. Let $a_i = \lfloor \frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} \rfloor$, which is the integer multiple of κ_f before the upper bound becomes flat ($f(\Pi_i)$). Note that we always have $0 \leq a_i \leq |\Pi_i| - |\Pi_{i-1}|$ by definition of κ_f .

Our first goal is to sum over terms between two adjacent principal partitions. We will show that

$$\begin{aligned}
 \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) &\leq f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) \\
 &\quad - \frac{(f(\Pi_i) - f(\Pi_{i-1}))^2}{2\kappa_f} + \frac{f(\Pi_i) - f(\Pi_{i-1})}{2}.
 \end{aligned} \tag{8}$$

There are three cases:

1. First suppose $a_i = 0$. We have $f(\Pi_i) - f(\Pi_{i-1}) < \kappa_f$ and (8) holds as the bound $\sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) \leq f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|)$ is true for all a_i .
2. Suppose $a_i = |\Pi_i| - |\Pi_{i-1}|$. Since $|\Pi_i| - |\Pi_{i-1}| = a_i = \lfloor \frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} \rfloor$, we have $f(\Pi_i) - f(\Pi_{i-1}) \geq \kappa_f(|\Pi_i| - |\Pi_{i-1}|)$. By definition of κ_f we have $f(\Pi_i) - f(\Pi_{i-1}) \leq \kappa_f(|\Pi_i| - |\Pi_{i-1}|)$, and therefore $f(\Pi_i) - f(\Pi_{i-1}) =$

$\kappa_f(|\Pi_i| - |\Pi_{i-1}|)$. Thus we have

$$\begin{aligned}
 \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) &\leq \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} (f(\Pi_{i-1}) + \kappa_f(j - |\Pi_{i-1}|)) \\
 &= f(\Pi_{i-1})(|\Pi_i| - |\Pi_{i-1}|) \\
 &\quad + \frac{\kappa_f(|\Pi_i| - |\Pi_{i-1}|)(|\Pi_i| - |\Pi_{i-1}| + 1)}{2} \\
 &= f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) - \kappa_f(|\Pi_i| - |\Pi_{i-1}|)^2 \\
 &\quad + \frac{\kappa_f(|\Pi_i| - |\Pi_{i-1}|)(|\Pi_i| - |\Pi_{i-1}| + 1)}{2} \\
 &= f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) - \frac{\kappa_f(|\Pi_i| - |\Pi_{i-1}|)^2}{2} + \frac{\kappa_f(|\Pi_i| - |\Pi_{i-1}|)}{2} \\
 &= f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) - \frac{(f(\Pi_i) - f(\Pi_{i-1}))^2}{2\kappa_f} + \frac{f(\Pi_i) - f(\Pi_{i-1})}{2},
 \end{aligned}$$

since in this case we have $f(\Pi_i) - f(\Pi_{i-1}) = \kappa_f(|\Pi_i| - |\Pi_{i-1}|)$. Thus (8) also holds when $a_i = |\Pi_i| - |\Pi_{i-1}|$.

3. Now suppose $0 < a_i < |\Pi_i| - |\Pi_{i-1}|$, we have:

$$\begin{aligned}
 \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) &= \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_{i-1}|+a_i} f(E_{j,\sigma}) + \sum_{j=|\Pi_{i-1}|+a_i+1}^{|\Pi_i|} f(E_{j,\sigma}) \\
 &\leq \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_{i-1}|+a_i} (f(\Pi_{i-1}) + \kappa_f(j - |\Pi_{i-1}|)) + \sum_{j=|\Pi_{i-1}|+a_i+1}^{|\Pi_i|} f(\Pi_i) \\
 &= a_i f(\Pi_{i-1}) + \frac{\kappa_f a_i (a_i + 1)}{2} + f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}| - a_i) \\
 &= f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) + \frac{\kappa_f a_i (a_i + 1)}{2} - a_i (f(\Pi_i) - f(\Pi_{i-1})).
 \end{aligned}$$

Let $\eta_i := \frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} - a_i$ be the decimal part of $\frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f}$. Substituting a_i with $\frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} - \eta_i$ with the above inequality we have

$$\begin{aligned}
 \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) &\leq f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) \\
 &\quad + \frac{\kappa_f}{2} \left(\frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} - \eta_i \right) \left(\frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} - \eta_i + 1 \right) \\
 &\quad - \left(\frac{f(\Pi_i) - f(\Pi_{i-1})}{\kappa_f} - \eta_i \right) (f(\Pi_i) - f(\Pi_{i-1})) \\
 &\leq f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) - \frac{(f(\Pi_i) - f(\Pi_{i-1}))^2}{2\kappa_f}
 \end{aligned}$$

$$+ \frac{f(\Pi_i) - f(\Pi_{i-1})}{2} + \frac{\kappa_f}{2}(\eta_i^2 - \eta_i).$$

As $0 \leq \eta_i < 1$ and we have $\eta_i^2 - \eta_i \leq 0$. It follows,

$$\begin{aligned} \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) &\leq f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) \\ &\quad - \frac{(f(\Pi_i) - f(\Pi_{i-1}))^2}{2\kappa_f} + \frac{f(\Pi_i) - f(\Pi_{i-1})}{2}. \end{aligned}$$

Thus inequality (8) always holds.

For easier manipulation and telescoping later in the computation, we rewrite the terms as $f(\Pi_i)(|\Pi_i| - |\Pi_{i-1}|) = f(E)(|\Pi_i| - |\Pi_{i-1}|) - (f(E) - f(\Pi_i))(|\Pi_i| - |\Pi_{i-1}|)$ and $(f(\Pi_i) - f(\Pi_{i-1}))^2 = f(\Pi_i)^2 - f(\Pi_{i-1})^2 - 2f(\Pi_{i-1})(f(\Pi_i) - f(\Pi_{i-1}))$. Now summing over i from 1 to s we have that (8) implies,

$$\begin{aligned} \sum_{k=1}^m f(E_{k,\sigma}) &= \sum_{i=1}^s \sum_{j=|\Pi_{i-1}|+1}^{|\Pi_i|} f(E_{j,\sigma}) \\ &\leq \sum_{i=1}^s f(E)(|\Pi_i| - |\Pi_{i-1}|) - \sum_{i=1}^s (f(E) - f(\Pi_i))(|\Pi_i| - |\Pi_{i-1}|) \\ &\quad - \sum_{i=1}^s \frac{f(\Pi_i)^2 - f(\Pi_{i-1})^2}{2\kappa_f} + \sum_{i=1}^s \frac{f(\Pi_{i-1})(f(\Pi_i) - f(\Pi_{i-1}))}{\kappa_f} \\ &\quad + \sum_{i=1}^s \frac{f(\Pi_i) - f(\Pi_{i-1})}{2} \\ &= f(E)|E| - \frac{f(E)^2}{2\kappa_f} + \frac{f(E)}{2} - \sum_{i=1}^s (f(E) - f(\Pi_i))(|\Pi_i| - |\Pi_{i-1}|) \\ &\quad + \sum_{i=1}^s \frac{f(\Pi_{i-1})(f(\Pi_i) - f(\Pi_{i-1}))}{\kappa_f}. \end{aligned}$$

□

References

1. Iwata, S., Tetali, P., Tripathi, P.: Approximating minimum linear ordering problems. In: APPROX, pp. 206–217 (2012)
2. Oxley, J., Welsh, D.: Chromatic, flow and reliability polynomials: the complexity of their coefficients. *Comb. Probab. Comput.* **11**(4), 403–426 (2002)
3. Tutte, W.T.: Matroids and graphs. *Trans. Am. Math. Soc.* **90**(3), 527–552 (1959)
4. Oxley, J.G.: *Matroid Theory*, vol. 3. Oxford University Press, Oxford (2006)
5. Chekuri, C., Motwani, R.: Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Discret. Appl. Math.* **98**(1–2), 29–38 (1999)

6. Hall, L.A., Schulz, A.S., Shmoys, D.B., Wein, J.: Scheduling to minimize average completion time: off-line and on-line approximation algorithms. *Math. Oper. Res.* **22**(3), 513–544 (1997)
7. Margot, F., Queyranne, M., Wang, Y.: Decompositions, network flows, and a precedence constrained single-machine scheduling problem. *Oper. Res.* **51**(6), 981–992 (2003)
8. Ambühl, C., Mastrolilli, M., Mutsanas, N., Svensson, O.: On the approximability of single-machine scheduling with precedence constraints. *Math. Oper. Res.* **36**(4), 653–669 (2011)
9. Díaz, J., Petit, J., Serna, M.: A survey of graph layout problems. *ACM Comput. Surv. (CSUR)* **34**(3), 313–356 (2002)
10. Petit, J.: Addenda to the survey of layout problems. *Bull. EATCS* **3**(105), 177–201 (2013)
11. Bezrukov, S.L.: Edge isoperimetric problems on graphs. *Graph Theory Comb. Biol.* **7**, 157–197 (1999)
12. Lai, Y.-L., Williams, K.: A survey of solved problems and applications on bandwidth, edgesum, and profile of graphs. *J. Graph Theory* **31**(2), 75–94 (1999)
13. Kishi, G., Kajitani, Y.: Maximally distant trees and principal partition of a linear graph. *IEEE Trans. Circuit Theory* **16**(3), 323–330 (1969)
14. Fujishige, S.: *Theory of Principal Partitions Revisited*, pp. 127–162. Springer, Berlin (2009). https://doi.org/10.1007/978-3-540-76796-1_7
15. Pizaruk, N.: The boundaries of submodular functions. *Comput. Math. Math. Phys.* **32**(12), 1769–1783 (1992)
16. Fokkink, R., Lidbetter, T., Végh, L.A.: On submodular search and machine scheduling. *Math. Oper. Res.* **44**(4), 1431–1449 (2019)
17. Harper, L.H.: Optimal assignments of numbers to vertices. *J. Soc. Ind. Appl. Math.* **12**(1), 131–135 (1964)
18. Shiloach, Y.: A minimum linear arrangement algorithm for undirected trees. *SIAM J. Comput.* **8**(1), 15–18 (1979). (02)
19. Chung, F.-R.K.: On optimal linear arrangements of trees. *Comput. Math. Appl.* **10**(1), 43–60 (1984)
20. Garey, M.R., Johnson, D.S., Stockmeyer, L.: Some simplified NP-complete problems. In: *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing*, pp. 47–63 (1974)
21. Feige, U., Lee, J.R.: An improved approximation ratio for the minimum linear arrangement problem. *Inf. Process. Lett.* **101**(1), 26–29 (2007)
22. Charikar, M., Hajiaghayi, M.T., Karloff, H., Rao, S.: ℓ_2^2 spreading metrics for vertex ordering problems. *Algorithmica* **56**(4), 577–604 (2010)
23. Impagliazzo, R., Paturi, R.: On the complexity of k-SAT. *J. Comput. Syst. Sci.* **62**(2), 367–375 (2001)
24. Ambühl, C., Mastrolilli, M., Svensson, O.: Inapproximability results for maximum edge biclique, minimum linear arrangement, and sparsest cut. *SIAM J. Comput.* **40**(2), 567–596 (2011)
25. Azar, Y., Gamzu, I., Yin, X.: Multiple intents re-ranking. In: *Symposium on Theory of Computing, STOC*, pp. 669–678 (2009)
26. Feige, U., Lovász, L., Tetali, P.: Approximating min sum set cover. *Algorithmica* **40**(4), 219–234 (2004)
27. Hassin, R., Levin, A.: An approximation algorithm for the minimum latency set cover problem. In: *European Symposium on Algorithms, ESA*, pp. 726–733 (2005)
28. Díaz, J., Gibbons, A., Paterson, M., Toran, J.: The minsumcut problem. In: *Workshop on Algorithms and Data Structures*. Springer, Berlin, pp. 65–79 (1991)
29. Lin, Y., Yuan, J.: Profile minimization problem for matrices and graphs. *Acta Math. Appl. Sin.* **10**(1), 107–112 (1994)
30. Even, S., Shiloach, Y.: NP-completeness of several arrangement problems. Technical Report; Israel Institute of Technology, Department of Computer Science, vol. 43 (1975)
31. Bansal, N., Batra, J., Farhadi, M., Tetali, P.: Improved approximations for min sum vertex cover and generalized min sum set cover. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 998–1005. SIAM, Philadelphia (2021)
32. Rao, S., Richa, A.W.: New approximation techniques for some linear ordering problems. *SIAM J. Comput.* **34**(2), 388–404 (2005)
33. Burer, S., Monteiro, R.D.: A projected gradient algorithm for solving the maxcut SDP relaxation. *Optim. Methods Softw.* **15**(3–4), 175–200 (2001)
34. Barenholz, U., Feige, U., Peleg, D., et al.: Improved approximation for min-sum vertex cover, MCS06-07. Technical Report, Computer Science and Applied Mathematics (2006)
35. Stanković, A.: Some results on approximability of minimum sum vertex cover. arXiv preprint [arXiv:2212.11882](https://arxiv.org/abs/2212.11882) (2022)

36. Bansal, N., Gupta, A., Krishnaswamy, R.: A constant factor approximation algorithm for generalized min-sum set cover. In: Symposium on Discrete Algorithms, SODA, pp. 1539–1545 (2010)
37. Skutella, M., Williamson, D.P.: A note on the generalized min-sum set cover problem. *Oper. Res. Lett.* **39**(6), 433–436 (2011)
38. Im, S., Sviridenko, M., van der Zwaan, R.: Preemptive and non-preemptive generalized min sum set cover. *Math. Program.* **145**(1–2), 377–401 (2014)
39. Tsaparas, P., Ntoulas, A., Terzi, E.: Selecting a comprehensive set of reviews. In: ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 168–176 (2011)
40. Luo, W., Chakraborty, N., Sycara, K.: Distributed dynamic priority assignment and motion planning for multiple mobile robots with kinodynamic constraints. In: American Control Conference (ACC), pp. 148–154. IEEE (2016)
41. Fokkink, R., Lidbetter, T., Végh, L.A.: On submodular search and machine scheduling. *Math. Oper. Res.* **44**(4), 1431–1449 (2019)
42. Happach, F.: Min-sum set cover, or-scheduling, and related problems. Ph.D. dissertation, Technische Universität München (2020)
43. Happach, F., Hellerstein, L., Lidbetter, T.: A general framework for approximating min sum ordering problems. *arXiv preprint arXiv:2004.05954* (2020)
44. Schrijver, A.: Combinatorial Optimization. Springer, Berlin (2003)
45. Narayanan, H.: The principal lattice of partitions of a submodular function. *Linear Algebra Appl.* **144**, 179–216 (1991)
46. Khachiyan, L., Boros, E., Elbassioni, K., Gurvich, V., Makino, K.: On the complexity of some enumeration problems for matroids. *SIAM J. Discrete Math.* **19**(4), 966–984 (2005)
47. Bansal, N., Khot, S.: Optimal long code test with one free bit. In: Foundations of Computer Science, FOCS, pp. 453–462 (2009)
48. Correa, J.R., Schulz, A.S.: Single-machine scheduling with precedence constraints. *Math. Oper. Res.* **30**(4), 1005–1021 (2005)
49. Ambühl, C., Mastrolilli, M.: Single machine precedence constrained scheduling is a vertex cover problem. *Algorithmica* **53**(4), 488–503 (2009)
50. Khot, S., Regev, O.: Vertex cover might be hard to approximate to within $2 - \varepsilon$. *J. Comput. Syst. Sci.* **74**(3), 335–349 (2008)
51. Hochbaum, D.S.: Efficient bounds for the stable set, vertex cover and set packing problems. *Discrete Appl. Math.* **6**(3), 243–254 (1983)
52. Brightwell, G.R., Scheinerman, E.R.: Fractional dimension of partial orders. *Order* **9**(2), 139–158 (1992)
53. Lau, L.C., Ravi, R., Singh, M.: Iterative Methods in Combinatorial Optimization, vol. 46. Cambridge University Press, Cambridge (2011)
54. Nakano, K.: Linear layout of generalized hypercubes. *Int. J. Found. Comput. Sci.* **14**(01), 137–156 (2003)
55. Muradyan, D., Piliposjan, T.: Minimal numberings of vertices of a rectangular lattice. *Akad. Nauk. Armjan. SRR* **1**(70), 21–27 (1980). (in Russian)
56. Muradyan, D., Piliposyan, T.: The problem of finding the length and width of the complete p-partite graph. *Fluchen. Zapiski Erevan. Gosunivers* **2**, 18–26 (1980). (in Russian)
57. Harper, L.: Chassis layout and isoperimetric problems. *Jet Propuls. Lab. SPS* **11**, 37–66 (1970)
58. Nagano, K., Kawahara, Y., Aihara, K.: Size-constrained submodular minimization through minimum norm base. In: Proceedings of the 28th International Conference on Machine Learning (ICML-11), pp. 977–984 (2011)