

Changing Bases: Multistage Optimization for Matroids and Matchings

Anupam Gupta^{1*}, Kunal Talwar², and Udi Wieder²

¹ Computer Science Dept., Carnegie Mellon University, Pittsburgh, PA 15213.
anupamg@cs.cmu.edu

² Microsoft Research Silicon Valley, 1065 La Avenida, Mountain View, CA, USA.
{kunal, uwieder}@microsoft.com

Abstract. This paper is motivated by the fact that many systems need to be maintained continually while the underlying costs change over time. The challenge then is to continually maintain near-optimal solutions to the underlying optimization problems, without creating too much churn in the solution itself. We model this as a multistage combinatorial optimization problem where the input is a sequence of cost functions (one for each time step); while we can change the solution from step to step, we incur an additional cost for every such change.

We first study the multistage matroid maintenance problem, where we need to maintain a base of a matroid in each time step under the changing cost functions and acquisition costs for adding new elements. The online version of this problem generalizes online paging, and is a well-structured case of the metrical task systems. E.g., given a graph, we need to maintain a spanning tree T_t at each step: we pay $c_t(T_t)$ for the cost of the tree at time t , and also $|T_t \setminus T_{t-1}|$ for the number of edges changed at this step. Our main result is a polynomial time $O(\log m \log r)$ -approximation to the online multistage matroid maintenance problem, where m is the number of elements/edges and r is the rank of the matroid. This improves on results of Buchbinder et al. [7] who addressed the *fractional* version of this problem under uniform acquisition costs, and Buchbinder, Chen and Naor [8] who studied the fractional version of a more general problem. We also give an $O(\log m)$ approximation for the offline version of the problem. These bounds hold when the acquisition costs are non-uniform, in which case both these results are the best possible unless $P=NP$.

We also study the perfect matching version of the problem, where we must maintain a perfect matching at each step under changing cost functions and costs for adding new elements. Surprisingly, the hardness drastically increases: for any constant $\varepsilon > 0$, there is no $O(n^{1-\varepsilon})$ -approximation to the multistage matching maintenance problem, even in the offline case.

1 Introduction

In a typical instance of a combinatorial optimization problem the underlying constraints model a static application frozen in one time step. In many applications however, one needs to solve instances of the combinatorial optimization

* Research performed while the author was at Microsoft Research Silicon Valley. Research was partly supported by NSF awards CCF-0964474 and CCF-1016799, and by a grant from the CMU-Microsoft Center for Computational Thinking.

problem that changes over time. While this is naturally handled by re-solving the optimization problem in each time step separately, changing the solution one holds from one time step to the next often incurs a transition cost. Consider, for example, the problem faced by a vendor who needs to get supply of an item from k different producers to meet her demand. On any given day, she could get prices from each of the producers and pick the k cheapest ones to buy from. As prices change, this set of the k cheapest producers may change. However, there is a fixed cost to starting and/or ending a relationship with any new producer. The goal of the vendor is to minimize the sum total of these two costs: an "acquisition cost" $a(e)$ to be incurred each time she starts a new business relationship with a producer, and a per period cost $c_t(e)$ of buying in period t from the each of the k producers that she picks in this period, summed over T time periods. In this work we consider a generalization of this problem, where the constraint "pick k producers" may be replaced by a more general combinatorial constraint. It is natural to ask whether simple combinatorial problems for which the one-shot problem is easy to solve, as the example above is, also admit good algorithms for the multistage version.

The first problem we study is the *Multistage Matroid Maintenance* problem (MMM), where the underlying combinatorial constraint is that of maintaining a base of a given matroid in each period. In the example above, the requirement the vendor buys from k different producers could be expressed as optimizing over the k -uniform matroid. In a more interesting case one may want to maintain a spanning tree of a given graph at each step, where the edge costs $c_t(e)$ change over time, and an acquisition cost of $a(e)$ has to be paid every time a new edge enters the spanning tree. (A formal definition of the MMM problem appears in Section 2.) While our emphasis is on the online problem, we will mention results for the offline version as well, where the whole input is given in advance.

A first observation we make is that if the matroid in question is allowed to be different in each time period, then the problem is hard to approximate to any non-trivial factor (see Section A.1) even in the offline case. We therefore focus on the case where the same matroid is given at each time period. Thus we restrict ourselves to the case when the matroid is the same for all time steps.

To set the baseline, we first study the offline version of the problem (in Section 3), where all the input parameters are known in advance. We show an LP-rounding algorithm which approximates the total cost up to a logarithmic factor. This approximation factor is no better than that using a simple greedy algorithm, but it will be useful to see the rounding algorithm, since we will use its extension in the online setting. We also show a matching hardness reduction, proving that the problem is hard to approximate to better than a logarithmic factor; this hardness holds even for the special case of spanning trees in graphs.

We then turn to the online version of the problem, where in each time period, we learn the costs $c_t(e)$ of each element that is available at time t , and we need to pick a base S_t of the matroid for this period. We analyze the performance of our online algorithm in the competitive analysis framework: i.e., we compare the cost of the online algorithm to that of the optimum solution to the offline instance thus generated. In Section 4, we give an efficient randomized $O(\log |E| \log(rT))$ -competitive algorithm for this problem against any oblivious adversary (here E is the universe for the matroid and r is the rank of the matroid), and show

that no polynomial-time online algorithm can do better. We also show that the requirement that the algorithm be randomized is necessary: any deterministic algorithm must incur an overhead of $\Omega(\min(|E|, T))$, even for the simplest of matroids.

Our results above crucially relied on the properties of matroids, and it is natural to ask if we can handle more general set systems, e.g., p -systems. In Section 5, we consider the case where the combinatorial object we need to find each time step is a perfect matching in a graph. Somewhat surprisingly, the problem here is significantly harder than the matroid case, even in the offline case. In particular, we show that even when the number of periods is a constant, no polynomial time algorithm can achieve an approximation ratio better than $\Omega(|E|^{1-\epsilon})$ for any constant $\epsilon > 0$.

1.1 Techniques

We first show that the **MMM** problem, which is a packing-covering problem, can be reduced to the analogous problem of maintaining a spanning set of a matroid. We call the latter the *Multistage Spanning set Maintenance* (**MSM**) problem. While the reduction itself is fairly clean, it is surprisingly powerful and is what enables us to improve on previous works. The **MSM** problem is a covering problem, so it admits better approximation ratios and allows for a much larger toolbox of techniques at our disposal. We note that this is the only place where we need the matroid to not change over time: our algorithms for **MSM** work when the matroids change over time, and even when considering matroid intersections. The **MSM** problem is then further reduced to the case where the holding cost of an element is in $\{0, \infty\}$, this reduction simplifies the analysis.

In the offline case, we present two algorithms. We first observe that a greedy algorithm easily gives an $O(\log T)$ -approximation. We then present a simple randomized rounding algorithm for the linear program. This is analyzed using recent results on contention resolution schemes [13], and gives an approximation of $O(\log rT)$, which can be improved to $O(\log r)$ when the acquisition costs are uniform. This LP-rounding algorithm will be an important constituent of our algorithm for the online case.

For the online case we again use that the problem can be written as a covering problem, even though the natural LP formulation has both covering and packing constraints. Phrasing it as a covering problem (with box constraints) enables us to use, as a black-box, results on online algorithms for the fractional problem [9]. This formulation however has exponentially many constraints. We handle that by showing a method of adaptively picking violated constraints such that only a small number of constraints are ever picked. The crucial insight here is that if x is such that $2x$ is not feasible, then x is at least $\frac{1}{2}$ away in ℓ_1 distance from any feasible solution; in fact there is a single constraint that is violated to an extent half. This insight allows us to make non-trivial progress (using a natural potential function) every time we bring in a constraint, and lets us bound the number of constraints we need to add until constraints are satisfied by $2x$.

1.2 Related Work

Our work is related to several lines of research, and extends some of them. The paging problem is a special case of **MMM** where the underlying matroid is a uniform one. Our online algorithm generalizes the $O(\log k)$ -competitive algorithm

for weighted caching [5], using existing online LP solvers in a black-box fashion. Going from uniform to general matroids loses a logarithmic factor (after rounding), we show such a loss is unavoidable unless we use exponential time.

The MMM problem is also a special case of classical Metrical Task Systems [6]; see [1, 4] for more recent work. The best approximations for metrical task systems are poly-logarithmic in the size of the metric space. In our case the metric space is specified by the total number of bases of the matroid which is often exponential, so these algorithms only give a trivial approximation.

In trying to unify online learning and competitive analysis, Buchbinder et al. [7] consider a problem on matroids very similar to ours. The salient differences are: (a) in their model all acquisition costs are the same, and (b) they work with fractional bases instead of integral ones. They give an $O(\log n)$ -competitive algorithm to solve the fractional online LP with uniform acquisition costs (among other unrelated results). Our online LP solving generalizes their result to arbitrary acquisition costs. They leave open the question of getting integer solutions online (Seffi Naor, private communication), which we present in this work. In a more recent work, Buchbinder, Chen and Naor [8] use a regularization approach to solving a broader set of fractional problems, but once again can do not get integer solutions in a setting such as ours.

Shachnai et al. [28] consider “reoptimization” problems: given a starting solution and a new instance, they want to balance the transition cost and the cost on the new instance. This is a two-timestep version of our problem, and the short time horizon raises a very different set of issues (since the output solution does not need to itself hedge against possible subsequent futures). They consider a number of optimization/scheduling problems in their framework.

Cohen et al. [15] consider several problems in the framework of the stability-versus-fit tradeoff; e.g., that of finding “stable” solutions which given the previous solution, like in reoptimization, is the current solution that maximizes the quality minus the transition costs. They show maintaining stable solutions for matroids becomes a repeated two-stage reoptimization problem; their problem is poly-time solvable, whereas matroid problems in our model become NP-hard. The reason is that the solution for two time steps does not necessarily lead to a base from which it is easy to move in subsequent time steps, as our hardness reduction shows. They consider a multistage offline version of their problem (again maximizing fit minus stability) which is very similar in spirit and form to our (minimization) problem, though the minus sign in the objective function makes it difficult to approximate in cases which are not in poly-time.

In dynamic Steiner tree maintenance [21, 24, 18] where the goal is to maintain an approximately optimal Steiner tree for a varying instance (where terminals are added) while changing few edges at each time step. In dynamic load balancing [2, 16] one has to maintain a good scheduling solution while moving a small number of jobs around. The work on lazy experts in the online prediction community [11] also deals with similar concerns.

There is also work on “leasing” problems [25, 3, 26]: these are optimization problems where elements can be obtained for an interval of any length, where the cost is concave in the lengths; the instance changes at each timestep. The main differences are that the solution only needs to be feasible at each timestep

(i.e., the holding costs are $\{0, \infty\}$), and that any element can be leased for any length ℓ of time starting at any timestep for a cost that depends only on ℓ , which gives these problems a lot of uniformity. In turn, these leasing problems are related to “buy-at-bulk” problems.

2 Maintaining Bases to Maintaining Spanning Sets

Given reals $c(e)$ for elements $e \in E$, we will use $c(S)$ for $S \subseteq E$ to denote $\sum_{e \in S} c(e)$. We denote $\{1, 2, \dots, T\}$ by $[T]$.

We assume basic familiarity with matroids: see, e.g., [27] for a detailed treatment. Given a matroid $\mathcal{M} = (E, \mathcal{I})$, a *base* is a maximum cardinality independent set, and a *spanning set* is a set S such that $\text{rank}(S) = \text{rank}(E)$; equivalently, this set contains a base within it. The *span* of a set $S \subseteq E$ is $\text{span}(S) = \{e \in E \mid \text{rank}(S + e) = \text{rank}(S)\}$. The *matroid polytope* $\mathcal{P}_I(\mathcal{M})$ is defined as $\{x \in \mathbb{R}_{\geq 0}^{|E|} \mid x(S) \leq \text{rank}(S) \forall S \subseteq E\}$. The *base polytope* $\mathcal{P}_B(\mathcal{M}) = \mathcal{P}_I(\mathcal{M}) \cap \{x \mid x(E) = \text{rank}(E)\}$. We will sometimes use m to denote $|E|$ and r to denote the rank of the matroid.

Formal Definition of Problems An instance of the *Multistage Matroid Maintenance* (MMM) problem consists of a matroid $\mathcal{M} = (E, \mathcal{I})$, an *acquisition cost* $a(e) \geq 0$ for each $e \in E$, and for every timestep $t \in [T]$ and element $e \in E$, a *holding cost* cost $c_t(e)$. The goal is to find bases $\{B_t \in \mathcal{I}\}_{t \in [T]}$ to minimize

$$\sum_t (c_t(B_t) + a(B_t \setminus B_{t-1})), \quad (2.1)$$

where we define $B_0 := \emptyset$. A related problem is the *Multistage Spanning set Maintenance* (MSM) problem, where we want to maintain a spanning set $S_t \subseteq E$ at each time, and cost of the solution $\{S_t\}_{t \in [T]}$ (once again with $S_0 := \emptyset$) is

$$\sum_t (c_t(S_t) + a(S_t \setminus S_{t-1})). \quad (2.2)$$

Maintaining Bases versus Maintaining Spanning Sets The following lemma shows the equivalence of maintaining bases and spanning sets. This enables us to significantly simplify the problem and avoid the difficulties faced by previous works on this problem.

Lemma 1. *For matroids, the optimal solutions to MMM and MSM have the same costs.*

Proof. Clearly, any solution to MMM is also a solution to MSM, since a base is also a spanning set. Conversely, consider a solution $\{S_t\}$ to MSM. Set B_1 to any base in S_1 . Given $B_{t-1} \subseteq S_{t-1}$, start with $B_{t-1} \cap S_t$, and extend it to any base B_t of S_t . This is the only step where we use the matroid properties—indeed, since the matroid is the same at each time, the set $B_{t-1} \cap S_t$ remains independent at time t , and by the matroid property this independent set can be extended to a base. Observe that this process just requires us to know the base B_{t-1} and the set S_t , and hence can be performed in an online fashion.

We claim that the cost of $\{B_t\}$ is no more than that of $\{S_t\}$. Indeed, $c_t(B_t) \leq c_t(S_t)$, because $B_t \subseteq S_t$. Moreover, let $D := B_t \setminus B_{t-1}$, we pay $\sum_{e \in D} a_e$ for these elements we just added. To charge this, consider any such element $e \in D$, let $t^* \leq t$ be the time it was most recently added to the cover—i.e., $e \in S_{t'}^*$ for all $t' \in [t^*, t]$, but $e \notin S_{t^*-1}$. The MSM solution paid for including e at time

t^* , and we charge our acquisition of e into B_t to this pair (e, t^*) . It suffices to now observe that we will not charge to this pair again, since the procedure to create $\{B_t\}$ ensures we do not drop e from the base until it is dropped from S_t itself—the next time we pay an addition cost for element e , it would have been dropped and added in $\{S_t\}$ as well. ■

Hence it suffices to give a good solution to the **MSM** problem. We observe that the proof above uses the matroid property crucially and would not hold, e.g., for matchings. It also requires that the *same* matroid be given at all time steps. Also, as noted above, the reduction is online: the instance is the same, and given an **MSM** solution it can be transformed online to a solution to **MMM**.

Elements and Intervals We will find it convenient to think of an instance of **MSM** as being a matroid \mathcal{M} , where each element only has an acquisition cost $a(e) \geq 0$, and it has a lifetime $I_e = [l_e, r_e]$. There are no holding costs, but the element e can be used in spanning sets only for timesteps $t \in I_e$. Or one can equivalently think of holding costs being zero for $t \in I_e$ and ∞ otherwise.

An Offline Exact Reduction. The translation is the natural one: given instance (E, \mathcal{I}) of **MSM**, create elements e_{lr} for each $e \in E$ and $1 \leq l \leq r \leq T$, with acquisition cost $a(e_{lr}) := a(e) + \sum_{t=l}^r c_t(e)$, and interval $I_{e_{lr}} := [l, r]$. (The matroid is extended in the natural way, where all the elements e_{lr} associated with e are parallel to each other.) The equivalence of the original definition of **MSM** and this interval view is easy to verify.

An Online Approximate Reduction. Observe that the above reduction created at most $\binom{T}{2}$ copies of each element, and required knowledge of all the costs. If we are willing to lose a constant factor in the approximation, we can perform a reduction to the interval model in an *online* fashion as follows. For element $e \in E$, define $t_0 = 0$, and create many parallel copies $\{e_i\}_{i \in \mathbb{Z}_+}$ of this element (modifying the matroid appropriately). Now the i^{th} interval for e is $I_{e_i} := [t_{i-1} + 1, t_i]$, where t_i is set to $t_{i-1} + 1$ in case $c_{t_{i-1}+1}(e) \geq a(e)$, else it is set to the *largest* time such that the total holding costs $\sum_{t=t_{i-1}+1}^{t_i} c_t(e)$ for this interval $[t_{i-1} + 1, t_i]$ is at most $a(e)$. This interval I_{e_i} is associated with element e_i , which is only available for this interval, at cost $a(e_i) = a(e) + c_{t_{i-1}+1}(e)$.

A few salient points about this reduction: the intervals for an original element e now partition the entire time horizon $[T]$. The number of elements in the modified matroid whose intervals contain any time t is now only $|E| = n$, the same as the original matroid; each element of the modified matroid is only available for a single interval. Moreover, the reduction can be done online: given the past history and the holding cost for the current time step t , we can ascertain whether t is the beginning of a new interval (in which case the previous interval ended at $t-1$) and if so, we know the cost of acquiring a copy of e for the new interval is $a(e) + c_t(e)$. It is easy to check that the optimal cost in this interval model is within a constant factor of the optimal cost in the original acquisition/holding costs model.

3 Offline Algorithms

Given the reductions of the previous section, we can focus on the **MSM** problem. Being a covering problem, **MSM** is conceptually easier to solve: e.g., we could

use algorithms for submodular set cover [29] with the submodular function being the sum of ranks at each of the timesteps, to get an $O(\log T)$ approximation.

In Section B, we give a dual-fitting proof of the performance of the greedy algorithm. Here we give an LP-rounding algorithm which gives an $O(\log rT)$ approximation; this can be improved to $O(\log r)$ in the common case where all acquisition costs are unit. (While the approximation guarantee is no better than that from submodular set cover, this LP-rounding algorithm will prove useful in the online case in Section 4). Finally, the hardness results of Section C.2 show that we cannot hope to do much better than these logarithmic approximations.

3.1 The LP Rounding Algorithm

We now consider an LP-rounding algorithm for the **MMM** problem; this will generalize to the online setting, whereas it is unclear how to extend the greedy algorithm to that case. For the LP rounding, we use the standard definition of the **MMM** problem to write the following LP relaxation.

$$\begin{aligned} \min \sum_{t,e} a(e) \cdot y_t(e) + \sum_{t,e} c_t(e) \cdot z_t(e) \\ \text{s.t. } z_t \in \mathcal{P}_B(\mathcal{M}) \quad \forall t \\ y_t(e) \geq z_t(e) - z_{t-1}(e) \quad \forall t, e \\ y_t(e), z_t(e) \geq 0 \end{aligned} \tag{LP2}$$

It remains to round the solution to get a feasible solution to **MSM** (i.e., a spanning set S_t for each time) with expected cost at most $O(\log n)$ times the LP value, since we can use Lemma 1 to convert this to a solution for **MMM** at no extra cost. The following lemma is well-known (see, e.g. [10]), we give a proof in the appendix for completeness.

Lemma 2. *For a fractional base $z \in \mathcal{P}_B(M)$, let $R(z)$ be the set obtained by picking each element $e \in E$ independently with probability z_e . Then $E[\text{rank}(R(z))] \geq r(1 - 1/e)$.*

Theorem 1. *Any fractional solution can be randomly rounded to get solution to **MSM** with cost $O(\log rT)$ times the fractional value, where r is the rank of the matroid and T the number of timesteps.*

Proof. Set $L = 32 \log(rT)$. For each element $e \in E$, choose a random threshold τ_e independently and uniformly from the interval $[0, 1/L]$. For each $t \in T$, define the set $\widehat{S}_t := \{e \in E \mid z_t(e) \geq \tau_e\}$; if \widehat{S}_t does not have full rank, augment its rank using the cheapest elements according to $(c_t(e) + a(e))$ to obtain a full rank set S_t . Since $\Pr[e \in \widehat{S}_t] = \min\{L \cdot z_t(e), 1\}$, the cost $c_t(\widehat{S}_t) \leq L \times (c_t \cdot z_t)$. Moreover, $e \in \widehat{S}_t \setminus \widehat{S}_{t-1}$ exactly when τ_e satisfies $z_{t-1}(e) < \tau_e \leq z_t(e)$, which happens with probability at most

$$\frac{\max(z_t(e) - z_{t-1}(e), 0)}{1/L} \leq L \cdot y_t(e).$$

Hence the expected acquisition cost for the elements newly added to \widehat{S}_t is at most $L \times \sum_e (a(e) \cdot y_t(e))$. Finally, we have to account for any elements added to extend \widehat{S}_t to a full-rank set S_t .

Lemma 3. *For any fixed $t \in [T]$, the set \hat{S}_t contains a basis of \mathcal{M} with probability at least $1 - 1/(rT)^8$.*

The proof of the lemma is a Chernoff bound, and appears in the Appendix. Now if the set \hat{S}_t does not have full rank, the elements we add have cost at most that of the min-cost base under the cost function $(a_e + c_t(e))$, which is at most the optimum value for (LP2). (We use the fact that the LP is exact for a single matroid, and the global LP has cost at least the single timestep cost.) This happens with probability at most $1/(rT)^8$, and hence the total expected cost of augmenting \hat{S}_t over all T timesteps is at most $O(1)$ times the LP value. This proves the main theorem. ■

Again, this algorithm for MSM works with different matroids at each timestep, and also for intersections of matroids. To see this observe that the only requirements from the algorithm are that there is a separation oracle for the polytope and that the contention resolution scheme works. In the case of k -matroid intersection, if we pay an extra $O(\log k)$ penalty in the approximation ratio we have that the probability a rounded solution does not contain a base is $< 1/k$ so we can take a union bound over the multiple matroids.

An Improvement: Avoiding the Dependence on T . When the ratio of the maximum to the minimum acquisition cost is small, we can improve the approximation factor above. More specifically, we show that essentially the same randomized rounding algorithm (with a different choice of L) gives an approximation ratio of $\log \frac{ra_{\max}}{a_{\min}}$. We defer the argument to Section D.1, as it needs some additional definitions and results that we present in the online section.

Hardness for Offline MSM. We defer the hardness proof to Appendix C, which shows that the MSM and MMM problems are NP-hard to approximate better than $\Omega(\min\{\log r, \log T\})$ even for graphical matroids. An integrality gap of $\Omega(\log T)$ appears in Appendix A.3.

4 Online MSM

We now turn to solving MMM in the online setting. In this setting, the acquisition costs $a(e)$ are known up-front, but the holding costs $c_t(e)$ for day t are not known before day t . Since the equivalence given in Lemma 1 between MMM and MSM holds even in the online setting, we can just work on the MSM problem. We show that the online MSM problem admits an $O(\log |E| \log rT)$ -competitive (oblivious) randomized algorithm. To do this, we show that one can find an $O(\log |E|)$ -competitive fractional solution to the linear programming relaxation in Section 3, and then we round this LP relaxation online, losing another logarithmic factor.

4.1 Solving the LP Relaxations Online

Again, we work in the interval model outlined in Section 2. Recall that in this model, for each element e there is a unique interval $I_e \subseteq [T]$ during which it is alive. The element e has an acquisition cost $a(e)$, no holding costs. Once an element has been acquired (which can be done at any time during its interval), it can be used at all times in that interval, but not after that. In the online

setting, at each time step t we are told which intervals have ended (and which have not); also, which new elements e are available starting at time t , along with their acquisition costs $a(e)$. Of course, we do not know when its interval I_e will end; this information is known only once the interval ends.

We will work with the same LP as in Section 3.1, albeit now we have to solve it online. The variable x_e is the indicator for whether we acquire element e .

$$\begin{aligned} P := \min \sum_e a(e) \cdot x_e \\ \text{s.t.} \quad z_{et} &\in \mathcal{P}_B(\mathcal{M}) & \forall t \\ z_{et} &\leq x_e & \forall e, t \in I_e \\ x_e, z_{et} &\in [0, 1] \end{aligned} \tag{LP3}$$

Note that this is not a packing or covering LP, which makes it more annoying to solve online. Hence we consider a slight reformulation. Let $\mathcal{P}_{ss}(\mathcal{M})$ denote the *spanning set polytope* defined as the convex hull of the full-rank (a.k.a. spanning) sets $\{\chi_S \mid S \subseteq E, \text{rank}(S) = r\}$. Since each spanning set contains a base, we can write the constraints of (LP3) as:

$$\mathbf{x}_{E_t} \in \mathcal{P}_{ss}(\mathcal{M}) \quad \forall t, \text{ where } E_t = \{e : t \in I_e\}. \tag{4.3}$$

Here we define \mathbf{x}_S to be the vector derived from \mathbf{x} by zeroing out the x_e values for $e \notin S$. It is known that the polytope $\mathcal{P}_{ss}(\mathcal{M})$ can be written as a (rather large) set of covering constraints. Indeed, $\mathbf{x} \in \mathcal{P}_{ss}(\mathcal{M}) \iff (\mathbf{1} - \mathbf{x}) \in \mathcal{P}_I(\mathcal{M}^*)$, where \mathcal{M}^* is the dual matroid for \mathcal{M} . Since the rank function of \mathcal{M}^* is given by $r^*(S) = r(E \setminus S) + |S| - r(E)$, it follows that (4.3) can be written as

$$\begin{aligned} \sum_{e \in S} x_e &\geq r(E) - r(E \setminus S) & \forall t, \forall S \subseteq E_t \\ x_e &\geq 0 & \forall e \in E \\ x_e &\leq 1 & \forall e \in E. \end{aligned} \tag{LP4}$$

Thus we get a covering LP with “box” constraints over E . The constraints can be presented one at a time: in timestep t , we present all the covering constraints corresponding to E_t . We remark that the newer machinery of [8] may be applicable to LP4. We next show that a simpler approach suffices³. The general results of Buchbinder and Naor [9] (and its extension to row-sparse covering problems by [19]) imply a deterministic algorithm for fractionally solving this linear program online, with a competitive ratio of $O(\log |E|) = O(\log m)$. However, this is not yet a polynomial-time algorithm, the number of constraints for each timestep being exponential. We next give an adaptive algorithm to generate a small yet sufficient set of constraints.

Solving the LP Online in Polynomial Time. Given a vector $\mathbf{x} \in [0, 1]^E$, define $\tilde{\mathbf{x}}$ as follows:

$$\tilde{x}_e = \min(2x_e, 1) \quad \forall e \in E. \tag{4.4}$$

Clearly, $\tilde{\mathbf{x}} \leq 2\mathbf{x}$ and $\tilde{\mathbf{x}} \in [0, 1]^E$. We next describe the algorithm for generating covering constraints in timestep t . Recall that [9] give us an online algorithm

³ Additionally, Lemma 4 will be useful in improving the rounding algorithm.

\mathcal{A}_{onLP} for solving a fractional covering LP with box constraints; we use this as a black-box. (This LP solver only raises variables, a fact we will use.) In timestep t , we adaptively select a small subset of the covering constraints from (LP4), and present it to \mathcal{A}_{onLP} . Moreover, given a fractional solution returned by \mathcal{A}_{onLP} , we will need to massage it at the end of timestep t to get a solution satisfying all the constraints from (LP4) corresponding to t .

Let \mathbf{x} be the fractional solution to (LP4) at the end of timestep $t - 1$. Now given information about timestep t , in particular the elements in E_t and their acquisition costs, we do the following. Given \mathbf{x} , we construct $\tilde{\mathbf{x}}$ and check if $\tilde{\mathbf{x}}_{E_t} \in \mathcal{P}_{ss}(M)$, as one can separate for $\mathcal{P}_{ss}(M)$. If $\tilde{\mathbf{x}}_{E_t} \in \mathcal{P}_{ss}(M)$, then $\tilde{\mathbf{x}}$ is feasible and we do not need to present any new constraints to \mathcal{A}_{onLP} , and we return $\tilde{\mathbf{x}}$. If not, our separation oracle presents an S such that the constraint $\sum_{e \in S} \tilde{x}_e \geq r(E) - r(E \setminus S)$ is violated. We present the constraint corresponding to S to \mathcal{A}_{onLP} to get an updated \mathbf{x} , and repeat until $\tilde{\mathbf{x}}$ is feasible for time t . (Since \mathcal{A}_{onLP} only raises variables and we have a covering LP, the solution remains feasible for past timesteps.) We next argue that we do not need to repeat this loop more than $2n$ times.

Lemma 4. *If for some \mathbf{x} and the corresponding $\tilde{\mathbf{x}}$, the constraint $\sum_{e \in S} \tilde{x}_e \geq r(E) - r(E \setminus S)$ is violated. Then*

$$\sum_{e \in S} x_e \leq r(E) - r(E \setminus S) - \frac{1}{2}$$

Proof. Let $S_1 = \{e \in S : \tilde{x}_e = 1\}$ and let $S_2 = S \setminus S_1$. Let γ denote $\sum_{e \in S_2} \tilde{x}_e$. Thus

$$|S_1| = \sum_{e \in S} \tilde{x}_e - \sum_{e \in S_2} \tilde{x}_e < r(E) - r(E \setminus S) - \gamma$$

Since both $|S_1|$ and $r(E) - r(E \setminus S)$ are integers, it follows that $|S_1| \leq r(E) - r(E \setminus S) - \lceil \gamma \rceil$. On the other hand, for every $e \in S_2$, $x_e = \frac{1}{2} \cdot \tilde{x}_e$, and thus $\sum_{e \in S_2} x_e = \frac{\gamma}{2}$. Consequently

$$\begin{aligned} \sum_{e \in S} x_e &= \sum_{e \in S_1} x_e + \sum_{e \in S_2} x_e = |S_1| + \frac{\gamma}{2} \\ &\leq r(E) - r(E \setminus S) - \lceil \gamma \rceil + \frac{\gamma}{2}. \end{aligned}$$

Finally, for any $\gamma > 0$, $\lceil \gamma \rceil - \frac{\gamma}{2} \geq \frac{1}{2}$, so the claim follows. \blacksquare

The algorithm \mathcal{A}_{onLP} updates \mathbf{x} to satisfy the constraint given to it, and Lemma 4 implies that each constraint we give to it must increase $\sum_{e \in E_t} x_e$ by at least $\frac{1}{2}$. The translation to the interval model ensures that the number of elements whose intervals contain t is at most $|E_t| \leq |E| = m$, and hence the total number of constraints presented at any time t is at most $2m$. We summarize the discussion of this section in the following theorem.

Theorem 2. *There is a polynomial-time online algorithm to compute an $O(\log |E|)$ -approximate solution to (LP3).*

We observe that the solution to this linear program can be trivially transformed to one for the LP in Section 3.1. Finally, the randomized rounding algorithm of Section 3.1 can be implemented online by selecting a threshold $t_e \in [0, 1/L]$

the beginning of the algorithm, where $L = \Theta(\log rT)$ and selecting element e whenever \tilde{x}_e exceeds t_e : here we use the fact that the online algorithm only ever raises x_e values, and this rounding algorithm is monotone. Rerandomizing in case of failure gives us an expected cost of $O(\log rT)$ times the LP solution, and hence we get an $O(\log m \log rT)$ -competitive algorithm.

An $O(\log r \frac{a_{\max}}{a_{\min}})$ -Approximate Rounding The dependence on the time horizon T is unsatisfactory in some settings, but we can do better using Lemma 4. Recall that the $\log(rT)$ -factor loss in the rounding follows from the naive union bound over the T time steps. We can argue that when $\frac{a_{\max}}{a_{\min}}$ is small, we can afford for the rounding to fail occasionally, and charge it to the acquisition cost incurred by the linear program. The details appear in Appendix D.1.

Hardness of the online MMM and online MSM. We can show that any polynomial-time algorithm cannot achieve better than an $\Omega(\log m \log T)$ competitive ratio, via a reduction from online set cover. Details appear in Appendix D.2.

5 Perfect Matching Maintenance

We next consider the *Perfect Matching Maintenance* (PMM) problem where E is the set of edges of a graph $G = (V, E)$, and the at each step, we need to maintain a perfect matchings in G .

Integrality Gap. Somewhat surprisingly, we show that the natural LP relaxation has an $\Omega(n)$ integrality gap, even for a constant number of timesteps. The LP and the (very simple) example appears in Appendix E.1.

Hardness. Moreover, in Appendix E.2 we show that the Perfect Matching Maintenance problem is very hard to approximate:

Theorem 3. *For any $\varepsilon > 0$ it is NP-hard to distinguish PMM instances with cost N^ε from those with cost $N^{1-\varepsilon}$, where N is the number of vertices in the graph. This holds even when the holding costs are in $\{0, \infty\}$, acquisition costs are 1 for all edges, and the number of time steps is a constant.*

6 Conclusions

In this paper we studied multistage optimization problems: an optimization problem (think about finding a minimum-cost spanning tree in a graph) needs to be solved repeatedly, each day a different set of element costs are presented, and there is a penalty for changing the elements picked as part of the solution. Hence one has to hedge between sticking to a suboptimal solution and changing solutions too rapidly. We present online and offline algorithms when the optimization problem is maintaining a base in a matroid. We show that our results are optimal under standard complexity-theoretic assumptions. We also show that the problem of maintaining a perfect matching becomes impossibly hard.

Our work suggests several directions for future research. It is natural to study other combinatorial optimization problems, both polynomial time solvable ones such shortest path and min-cut, as well NP-hard ones such as min-max load balancing and bin-packing in this multistage framework with acquisition costs. Moreover, the approximability of the *bipartite* matching maintenance, as well as matroid intersection maintenance remains open. Our hardness results for the

matroid problem hold when edges have $\{0, 1\}$ acquisition costs. The unweighted version where all acquisition costs are equal may be easier; we currently know no hardness results, or sub-logarithmic approximations for this useful special case.

References

- [1] Abernethy, J., Bartlett, P.L., Buchbinder, N., Stanton, I.: A regularization approach to metrical task systems. In: ALT. pp. 270–284 (2010)
- [2] Andrews, M., Goemans, M.X., Zhang, L.: Improved bounds for on-line load balancing. Algorithmica 23(4), 278–301 (1999)
- [3] Anthony, B.M., Gupta, A.: Infrastructure leasing problems. In: IPCO. pp. 424–438 (2007)
- [4] Bansal, N., Buchbinder, N., Naor, J.: Metrical task systems and the k -server problem on hsts. In: ICALP (1). pp. 287–298 (2010)
- [5] Bansal, N., Buchbinder, N., Naor, J.S.: Randomized competitive algorithms for generalized caching. In: STOC’08, pp. 235–244 (2008)
- [6] Borodin, A., Linial, N., Saks, M.E.: An optimal on-line algorithm for metrical task system. J. ACM 39(4), 745–763 (Oct 1992)
- [7] Buchbinder, N., Chen, S., Naor, J., Shamir, O.: Unified algorithms for online learning and competitive analysis. JMLR 23, 5.1–5.18 (2012)
- [8] Buchbinder, N., Chen, S., Naor, J.S.: Competitive analysis via regularization. In: ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 436–444 (2014)
- [9] Buchbinder, N., Naor, J.S.: Online primal-dual algorithms for covering and packing. Math. Oper. Res. 34(2), 270–286 (2009)
- [10] Calinescu, G., Chekuri, C., Pál, M., Vondrák, J.: Maximizing a submodular set function subject to a matroid constraint (extended abstract). pp. 182–196 (2007)
- [11] Cesa-Bianchi, N., Lugosi, G.: Prediction, learning, and games. Cambridge University Press (2006)
- [12] Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: STOC. pp. 311–320 (2010)
- [13] Chekuri, C., Vondrák, J., Zenklusen, R.: Submodular function maximization via the multilinear relaxation and contention resolution schemes. In: STOC. pp. 783–792 (2011)
- [14] Chvátal, V.: A greedy heuristic for the set-covering problem. Math. Oper. Res. 4(3), 233–235 (1979)
- [15] Cohen, E., Cormode, G., Duffield, N.G., Lund, C.: On the tradeoff between stability and fit. CoRR abs/1302.2137 (2013)
- [16] Epstein, L., Levin, A.: Robust algorithms for preemptive scheduling. In: ESA, Lecture Notes in Comput. Sci., vol. 6942, pp. 567–578. Springer, Heidelberg (2011)
- [17] Fisher, M.L., Nemhauser, G.L., Wolsey, L.A.: An analysis of approximations for maximizing submodular set functions. II. Math. Prog. Stud. (8), 73–87 (1978)
- [18] Gu, A., Gupta, A., Kumar, A.: The power of deferral: maintaining a constant-competitive steiner tree online. In: STOC. pp. 525–534 (2013)
- [19] Gupta, A., Nagarajan, V.: Approximating sparse covering integer programs online. In: ICALP (1) (Jul 2012)
- [20] Guruswami, V., Khanna, S.: On the hardness of 4-coloring a 3-colorable graph. SIAM J. Discrete Math. 18(1), 30–40 (electronic) (2004)
- [21] Imase, M., Waxman, B.M.: Dynamic Steiner tree problem. SIAM J. Discrete Math. 4(3), 369–384 (1991)
- [22] Kann, V.: Maximum bounded 3-dimensional matching is max SNP-complete. Inf. Process. Lett. 37(1), 27–35 (Jan 1991)
- [23] Korman, S.: On the use of randomness in the online set cover problem. M.Sc. thesis, Weizmann Institute of Science (2005)

- [24] Megow, N., Skutella, M., Verschae, J., Wiese, A.: The power of recourse for online MST and TSP. In: ICALP (1). pp. 689–700 (2012)
- [25] Meyerson, A.: The parking permit problem. In: FOCS. pp. 274–284 (2005)
- [26] Nagarajan, C., Williamson, D.P.: Offline and online facility leasing. In: IPCO. pp. 303–315 (2008)
- [27] Schrijver, A.: Combinatorial Optimization. Springer (2003)
- [28] Shachnai, H., Tamir, G., Tamir, T.: A theory and algorithms for combinatorial reoptimization. In: LATIN. pp. 618–630 (2012)
- [29] Wolsey, L.A.: An analysis of the greedy algorithm for the submodular set covering problem. Combinatorica 2(4), 385–393 (1982)

A Lower Bounds: Hardness and Gap Results

A.1 Hardness for Time-Varying Matroids

An extension of **MMM/MSM** problems is to the case when the set of elements remain the same, but the matroids change over time. Again the goal in **MMM** is to maintain a matroid base at each time.

Theorem 4. *The **MMM** problem with different matroids is NP-hard to approximate better than a factor of $\Omega(T)$, even for partition matroids, as long as $T \geq 3$.*

Proof. The reduction is from 3D-Matching (3DM). An instance of 3DM has three sets X, Y, Z of equal size $|X| = |Y| = |Z| = k$, and a set of hyperedges $E \subseteq X \times Y \times Z$. The goal is to choose a set of disjoint edges $M \subseteq E$ such that $|M| = k$.

First, consider the instance of **MMM** with three timesteps $T = 3$. The universe elements correspond to the edges. For $t = 1$, create a partition with k parts, with edges sharing a vertex in X falling in the same part. The matroid \mathcal{M}_1 is now to choose a set of elements with at most one element in each part. For $t = 2$, the partition now corresponds to edges that share a vertex in Y , and for $t = 3$, edges that share a vertex in Z . Set the movement weights $w(e) = 1$ for all edges.

If there exists a feasible solution to 3DM with k edges, choosing the corresponding elements form a solution with total weight k . If the largest matching is of size $(1 - \varepsilon)k$, then we must pay $\Omega(\varepsilon k)$ extra over these three timesteps. This gives a k -vs- $(1 + \Omega(\varepsilon))k$ gap for three timesteps.

To get a result for T timesteps, we give the same matroids repeatedly, giving matroids $\mathcal{M}_{t \pmod 3}$ at all times $t \in [T]$. In the “yes” case we would buy the edges corresponding to the 3D matching and pay nothing more than the initial k , whereas in the “no” case we would pay $\Omega(\varepsilon k)$ every three timesteps. Finally, the APX-hardness for 3DM [22] gives the claim. ■

The time-varying **MSM** problem does admit an $O(\log rT)$ approximation, as the randomized rounding (or the greedy algorithm) shows. However, the equivalence of **MMM** and **MSM** does not go through when the matroids change over time.

The restriction that the matroids vary over time is essential for the NP-hardness, since if the partition matroid is the same for all times, the complexity of the problem drops radically.

Theorem 5. *The MMM problem with partition matroids can be solved in polynomial time.*

Proof. The problem can be solved using min-cost flow. Indeed, consider the following reduction. Create a node v_{et} for each element e and timestep t . Let the partition be $E = E_1 \cup E_2 \cup \dots \cup E_r$. Then for each $i \in [r]$ and each $e, e' \in E_i$, add an arc $(v_{et}, v_{e',t+1})$, with cost $w(e') \cdot \mathbf{1}_{e \neq e'}$. Add a cost of $c_t(e)$ per unit flow through vertex v_{et} . (We could simulate this using edge-costs if needed.) Finally, add vertices s_1, s_2, \dots, s_r and source s . For each i , add arcs from s_i to all vertices $\{v_{e1}\}_{e \in E_i}$ with costs $w(e)$. All these arcs have infinite capacity. Now add unit capacity edges from s to each s_i , and infinite capacity edges from all nodes v_{eT} to t .

Since the flow polytope is integral for integral capacities, a flow of r units will trace out r paths from s to t , with the elements chosen at each time t being independent in the partition matroid, and the cost being exactly the per-time costs and movement costs of the elements. Observe that we could even have time-varying movement costs. Whereas, for graphical matroids the problem is $\Omega(\log n)$ hard even when the movement costs for each element do not change over time, and even just lie in the set $\{0, 1\}$. ■

Moreover, the restriction in Theorem 4 that $T \geq 3$ is also necessary, as the following result shows.

Theorem 6. *For the case of two rounds (i.e., $T = 2$) the MSM problem can be solved in polynomial time, even when the two matroids in the two rounds are different.*

Proof. The solution is simple, via matroid intersection. Suppose the matroids in the two timesteps are $\mathcal{M}_1 = (E, \mathcal{I}_1)$ and $\mathcal{M}_2 = (E, \mathcal{I}_2)$. Create elements (e, e') which corresponds to picking element e and e' in the two time steps, with cost $c_1(e) + c_2(e') + w_e + w_{e'} \mathbf{1}_{e \neq e'}$. Lift the matroids \mathcal{M}_1 and \mathcal{M}_2 to these tuples in the natural way, and look for a common basis. ■

A.2 Lower Bound for Deterministic Online Algorithms

We note that deterministic online algorithms cannot get any non-trivial guarantee for the MMM problem, even in the simple case of a 1-uniform matroid. This is related to the lower bound for deterministic algorithms for paging. Formally, we have the 1-uniform matroid on m elements, and $T = m$. All acquisition costs $a(e)$ are 1. In the first period, all holding costs are zero and the online algorithm picks an element, say e_1 . Since we are in the non-oblivious model, the algorithm knows e_1 and can in the second time step, set $c_2(e_1) = \infty$, while leaving the other ones at zero. Now the algorithm is forced to move to another edge, say e_2 , allowing the adversary to set $c_3(e_2) = \infty$ and so on. At the end of $T = m$ rounds, the online algorithm is forced to incur a cost of 1 in each round, giving a total cost of T . However, there is still an edge whose holding cost was zero throughout, so that the offline OPT is 1. Thus against a non-oblivious adversary, any online algorithm must incur a $\Omega(\min(m, T))$ overhead.

A.3 An $\Omega(\min(\log T, \log \frac{a_{\max}}{a_{\min}}))$ LP Integrality Gap

In this section, we show that if the aspect ratio of the movement costs is not bounded, the linear program has a $\log T$ gap, even when T is exponentially larger than m . We present an instance where $\log T$ and $\log \frac{a_{\max}}{a_{\min}}$ are about r with $m = r^2$, and the linear program has a gap of $\Omega(\min(\log T, \log \frac{a_{\max}}{a_{\min}}))$. This shows that the $O(\min(\log T, \log \frac{a_{\max}}{a_{\min}}))$ term in our rounding algorithm is unavoidable.

The instance is a graphical matroid, on a graph G on $\{v_0, v_1, \dots, v_n\}$, and $T = \binom{n}{\frac{n}{2}} = 2^{O(n)}$. The edges (v_0, v_i) for $i \in [n]$ have acquisition cost $a(v_0, v_i) = 1$ and holding cost $c_t(v_0, v_i) = 0$ for all t . The edges (v_i, v_j) for $i, j \in [n]$ have acquisition cost $\frac{1}{nT}$ and have holding cost determined as follows: we find a bijection between the set $[T]$ and the set of partitions (U_t, V_t) of $\{v_1, \dots, v_n\}$ with each of U_t and V_t having size $\frac{n}{2}$ (by choice of T such a bijection exists, and can be found e.g. by arranging the U_t 's in lexicographical order.) . In time step t , we set $c_t(e) = 0$ for $e \in (U_t \times U_t) \cup (V_t \times V_t)$, and $c_t(e) = \infty$ for all $e \in U_t \times V_t$.

First observe that no feasible integral solution to this instance can pay acquisition cost less than $\frac{n}{2}$ on the (v_0, v_i) edges. Suppose that the solution picks edges $\{(v_0, v_i) : v_i \in U_{sol}\}$ for some set U_{sol} of size at most $\frac{n}{2}$. Then any time step t such that $U_{sol} \subseteq U_t$, the solution has picked no edges connecting v_0 to V_t , and all edges connecting U_t to V_t have infinite holding cost in this time step. This contradicts the feasibility of the solution. Thus any integral solution has cost $\Omega(n)$.

Finally, we show that on this instance, (LP2) from Section 3.1, has a feasible solution of cost $O(1)$. We set $y_t(v_0, v_i) = \frac{2}{n}$ for all $i \in [n]$, and set $y_t(v_i, v_j) = \frac{2}{n}$ for $(v_i, v_j) \in (U_t \times U_t) \cup (V_t \times V_t)$. It is easy to check that $z_t = y_t$ is in the spanning tree polytope for all time steps t . Finally, the total acquisition cost is at most $n \cdot 1 \cdot \frac{2}{n}$ for the edges incident on v_0 and at most $T \cdot n^2 \cdot \frac{1}{nT} \cdot \frac{2}{n}$ for the other edges, both of which are $O(1)$. The holding costs paid by this solution is zero. Thus the LP has a solution of cost $O(1)$

The claim follows.

B The Greedy Algorithm

The greedy algorithm for MSM is the natural one. We consider the interval view of the problem (as in Section 2) where each element only has acquisition costs $a(e)$, and can be used only in some interval I_e . Given a current subset $X \subseteq E$, define $X_t := \{e' \in X \mid I_{e'} \ni t\}$. The benefit of adding an element e to X is

$$\text{ben}_X(e) = \sum_{t \in I_e} (\text{rank}(X_t \cup \{e\}) - \text{rank}(X_t))$$

and the greedy algorithm repeatedly picks an element e maximizing $\text{ben}_X(e)/a(e)$ and adds e to X . This is done until $\text{rank}(X_t) = r$ for all $t \in [T]$.

Phrased this way, an $O(\log T)$ bound on the approximation ration follows from Wolsey [29]. We next give an alternate dual fitting proof. We do not know of an

instance with uniform acquisition costs where greedy does not give a constant factor approximation. The dual fitting approach may be useful in proving a better approximation bound for this special case.

The natural LP is:

$$\begin{aligned}
P := \min \sum_e a(e) \cdot x_e & & & \text{(LP1)} \\
\text{s.t. } \{z_{et}\}_e \in \mathcal{P}_B(\mathcal{M}) & & \forall t & \\
z_{et} \leq x(e) & & \forall e, \forall t \in I_e & \\
x_e \geq 0 & & \forall e & \\
z_{et} \geq 0 & & \forall e, \forall t \in I_e &
\end{aligned}$$

where the polytope $\mathcal{P}_B(\mathcal{M})$ is the base polytope of the matroid \mathcal{M} .

Using Lagrangian variables $\beta_{et} \geq 0$ for each e and $t \in I_e$, we write a lower bound for P by

$$\begin{aligned}
D(\beta) := \min \sum_e a(e) \cdot x_e + \sum_{e,t \in I_e} \beta_{et} (z_{et} - x_e) & \\
\text{s.t. } z_{et} \in \mathcal{P}_B(\mathcal{M}) & \forall t \\
x_e, z_{et} \geq 0 &
\end{aligned}$$

which using the integrality of the matroid polytope can be rewritten as:

$$\min_{x \geq 0} \sum_e x_e (a(e) - \sum_{e,t \in I_e} \beta_{et}) + \sum_t \text{mst}(\beta_{et}).$$

Here, $\text{mst}(\beta_{et})$ denotes the cost of the minimum weight base at time t according to the element weights $\{\beta_{et}\}_{e \in E}$, where the available elements at time t is $E_t = \{x \in E \mid t \in I_e\}$. The best lower bound is:

$$\begin{aligned}
D := \max \sum_t \text{mst}(\beta_{et}) & \\
\text{s.t. } \sum_{t \in I_e} \beta_{et} \leq a(e) & \\
\beta_{et} \geq 0. &
\end{aligned}$$

The analysis of greedy follows the dual-fitting proofs of [14, 17].

Theorem 7. *The greedy algorithm outputs an $O(\log |I_{\max}|)$ -approximation to MSM, where $|I_{\max}|$ is the length of the longest interval that an element is alive for. Hence, it gives an $O(\log T)$ -approximation.*

Proof. For the proof, consider some point in the run of the greedy algorithm where set X of elements has been picked. We show a setting of duals β_{et} such that

- (a) the dual value equals the current primal cost $\sum_{e \in X} a(e)$, and
- (b) the constraints are nearly satisfied, namely $\sum_{t \in I_e} \beta_{et} \leq a(e) \log |I_e|$ for every $e \in E$.

It is useful to maintain, for each time t , a *minimum weight base* B_t of the subset $\text{span}(X_t)$ according to weights $\{\beta_{et}\}$. Hence the current dual value equals

$\sum_t \sum_{e \in B_t} \beta_{et}$. We start with $\beta_{et} = 0$ and $X_t = B_t = \emptyset$ for all t , which satisfies the above properties.

Suppose we now pick e maximizing $\text{ben}_X(e)/a(e)$ and get new set $X' := X \cup \{e\}$. We use $X'_t := \{e' \in X' \mid I_{e'} \ni t\}$ akin to our definition of X_t . Call a timestep t “interesting” if $\text{rank}(X'_t) = \text{rank}(X_t) + 1$; there are $\text{ben}_X(e)$ interesting timesteps. How do we update the duals? For $e' \in \text{span}(X'_t) \setminus \text{span}(X_t)$, we set $\beta_{e't} \leftarrow a(e)/\text{ben}_X(e)$. Note the element e itself satisfies the condition of being in $\text{span}(X'_t) \setminus \text{span}(X_t)$ for precisely the interesting timesteps, and hence $\sum_{t \text{ interesting}} \beta_{et} = (a(e)/\text{ben}_X(e)) \cdot \text{ben}_X(e) = a(e)$. For each interesting $t \in I_e$, define the base $B'_t \leftarrow B_t + e$; for all other times set $B'_t \leftarrow B_t$. It is easy to verify that B'_t is a base in $\text{span}(X'_t)$. But is it a min-weight base? Inductively assume that B_t was a min-weight base of $\text{span}(X_t)$; if t is not interesting there is nothing to prove, so consider an interesting t . All the elements in $\text{span}(X'_t) \setminus \text{span}(X_t)$ have just been assigned weight $\beta_{e't} = a(e)/\text{ben}_X(e)$, which by the monotonicity properties of the greedy algorithm is at least as large as the weight of any element in $\text{span}(X_t)$. Since e lies in $\text{span}(X'_t) \setminus \text{span}(X_t)$ and is assigned value $\beta_{et} = a(e)/\text{ben}_X(e)$, it cannot be swapped with any other element in $\text{span}(X'_t)$ to improve the weight of the base, and hence $B'_t = B_t + e$ is a min-weight base of $\text{span}(X'_t)$.

It remains to show that the dual constraints are approximately satisfied. Consider any element f , and let $\lambda = |I_f|$. The first step where we update β_{ft} for some $t \in I_f$ is when f is in the span of X_t for some time t . We claim that $\beta_{ft} \leq a(f)/\lambda$. Indeed, at this time f is a potential element to be added to the solution and it would cause a rank increase for λ time steps. The greedy rule ensures that we must have picked an element e with weight-to-coverage ratio at most as high. Similarly, the next t for which β_{ft} is updated will have $a(f)/(\lambda-1)$, etc. Hence we get the sum

$$\sum_t \beta_{ft} \leq a(f) \left(\frac{1}{|I_f|} + \frac{1}{|I_f|-1} + \dots + 1 \right) \leq a(f) \times O(\log |I_f|).$$

Since each element can only be alive for all T timesteps, we get the claimed $O(\log T)$ -approximation. ■

Note that the greedy algorithm would solve **MSM** even if we had a different matroid \mathcal{M}_t at each time t . However, the equivalence of **MMM** and **MSM** no longer holds in this setting, which is not surprising given the hardness of Theorem 4.

C Missing Proofs for the Offline Case (Section 3)

C.1 Missing Proofs

Proof of Lemma 2: We use the results of Chekuri et al. [13] (extending those of Chawla et al. [12]) on so-called contention resolution schemes. Formally, for a matroid \mathcal{M} , they give a randomized procedure π_z that takes the random set $R(z)$ and outputs an independent set $\pi_z(R(z))$ in \mathcal{M} , such that $\pi_z(R(z)) \subseteq R(z)$, and

for each element e in the support of z , $\Pr[e \in \pi_z(R(z)) \mid e \in R(z)] \geq (1 - 1/e)$. (They call this a $(1, 1 - 1/e)$ -balanced CR scheme.) Now, we get

$$\begin{aligned}\mathbf{E}[\text{rank}(R(z))] &\geq \mathbf{E}[\text{rank}(\pi_z(R(z)))] = \sum_{e \in \text{supp}(z)} \Pr[e \in \pi_z(R(z))] \\ &= \sum_{e \in \text{supp}(z)} \Pr[e \in \pi_z(R(z)) \mid e \in R(z)] \cdot \Pr[e \in R(z)] \\ &\geq \sum_{e \in \text{supp}(z)} (1 - 1/e) \cdot z_e = r(1 - 1/e).\end{aligned}$$

The first inequality used the fact that $\pi_z(R(z))$ is a subset of $R(z)$, the following equality used that $\pi_z(R(z))$ is independent with probability 1, the second inequality used the property of the CR scheme, and the final equality used the fact that z was a fractional base. ■

Proof of Lemma 3: The set \widehat{S}_t is obtained by threshold rounding of the fractional base $z_t \in \mathcal{P}_B(\mathcal{M})$ as above. Instead, consider taking L different samples $T^{(1)}, T^{(2)}, \dots, T^{(L)}$, where each sample is obtained by including each element $e \in E$ independently with probability $z_t(e)$; let $T := \cup_{i=1}^L T^{(i)}$. It is easy to check that $\Pr[\text{rank}(T) = r] \leq \Pr[\text{rank}(\widehat{S}_t) = r]$, so it suffices to give a lower bound on the former expression. For this, we use Lemma 2: the sample $T^{(1)}$ has expected rank $r(1 - 1/e)$, and using reverse Markov, it has rank at least $r/2$ with probability at least $1 - 2/e \geq 1/4$. Now focusing on the matroid \mathcal{M}' obtained by contracting elements in $\text{span}(T^{(1)})$ (which, say, has rank r'), the same argument says the set $T^{(2)}$ has rank $r'/2$ with probability at least $1/4$, etc. Proceeding in this way, the probability that the rank of T is less than r is at most the probability that we see fewer than $\log_2 r$ heads in $L = 32 \log r T$ flips of a coin of bias $1/4$. By a Chernoff bound, this is at most $\exp\{-(7/8)^2 \cdot (L/4)/3\} = 1/(rT)^8$. ■

C.2 Hardness for Offline MSM

Theorem 8. *The MSM and MMM problems are NP-hard to approximate better than $\Omega(\min\{\log r, \log T\})$ even for graphical matroids.*

Proof of Theorem 8: We give a reduction from Set Cover to the MSM problem for graphical matroids. Given an instance (U, \mathcal{F}) of set cover, with $m = |\mathcal{F}|$ sets and $n = |U|$ elements, we construct a graph as follows. There is a special vertex r , and m set vertices (with vertices s_i for each set $S_i \in \mathcal{F}$). There are m edges $e_i := (r, s_i)$ which all have inclusion weight $a(e_i) = 1$ and per-time cost $c_t(e) = 0$ for all t . All other edges will be zero cost short-term edges as given below. In particular, there are $T = n$ timesteps. In timestep $j \in [n]$, define subset $F_j := \{s_i \mid S_i \ni u_j\}$ to be vertices corresponding to sets containing element u_j . We have a set of edges $(e_i, e_{i'})$ for all $i, i' \in F_j$, and all edges (x, y) for $x, y \in \{r\} \cup \overline{F_j}$. All these edges have zero inclusion weight $a(e)$, and are only alive at time j . (Note this creates a graph with parallel edges, but this can be easily fixed by subdividing edges.)

In any solution to this problem, to connect the vertices in F_j to r , we must buy some edge (r, s_i) for some $s_i \in F_j$. This is true for all j , hence the root-set

edges we buy correspond to a set cover. Moreover, one can easily check that if we acquire edges (r, s_i) such that the sets $\{S_i : (r, s_i) \text{ acquired}\}$ form a set cover, then we can always augment using zero cost edges to get a spanning tree. Since the only edges we pay for are the (r, s_i) edges, we should buy edges corresponding to a min-cardinality set cover, which is hard to approximate better than $\Omega(\log n)$. Finally, that the number of time periods is $T = n$, and the rank of the matroid is $m = \text{poly}(n)$ for these hard instances. This gives us the claimed hardness. ■

D Missing Proofs for the Online Case (Section 4)

D.1 An $O(\log r \frac{a_{\max}}{a_{\min}})$ -Approximate Rounding

The dependence on the time horizon T is unsatisfactory in some settings, but we can do better using Lemma 4. Recall that the $\log(rT)$ -factor loss in the rounding follows from the naive union bound over the T time steps. We now argue that when $\frac{a_{\max}}{a_{\min}}$ is small, we can afford for the rounding to fail occasionally, and charge it to the acquisition cost incurred by the linear program.

Let us divide the period $[1 \dots T]$ into disjoint ‘‘epochs’’, where an epoch (except for the last) is an interval $[p, q)$ for $p \leq q$ such that the total fractional acquisition cost $\sum_{t=p}^{q-1} \sum_e a(e) \cdot y_t(e) \geq r \cdot a_{\max} > \sum_{t=p}^{q-2} \sum_e a(e) \cdot y_t(e)$. Thus an epoch is a minimal interval where the linear program spends acquisition cost $\in [r \cdot a_{\max}, 2r \cdot a_{\max}]$, so that we can afford to build a brand new tree once in each epoch and can charge it to the LP’s fractional acquisition cost in the epoch. Naively applying Theorem 1 to each epoch independently gives us a guarantee of $O(\log rT')$, where T' is the maximum length of an epoch.

However, an epoch can be fairly long if the LP solution changes very slowly. We break up each epoch into phases, where each phase is a maximal subsequence such that the LP incurs acquisition cost at most $\frac{a_{\min}}{4}$; clearly the epoch can be divided into at most $R := \frac{8ra_{\max}}{a_{\min}}$ disjoint phases. For a phase $[t_1, t_2]$, let $Z_{[t_1, t_2]}$ denote the solution defined as $Z_{[t_1, t_2]}(e) = \min_{t \in [t_1, t_2]} z_t(e)$. The definition of the phase implies that for any $t \in [t_1, t_2]$, the L_1 difference $\|Z_{[t_1, t_2]} - z_t\|_1 \leq \frac{1}{4}$. Now Lemma 4 implies that $\tilde{Z}_{[t_1, t_2]}$ is in $\mathcal{P}_{ss}(M)$, where \tilde{Z} is defined as in (4.4).

Suppose that in the randomized rounding algorithm, we pick the threshold $t_e \in [0, 1/L']$ for $L' = 64 \log R$. Let $\mathcal{G}_{[t_1, t_2]}$ be the event that the rounding algorithm applied to $Z_{[t_1, t_2]}$ gives a spanning set. Since $\tilde{Z}_{[t_1, t_2]} \leq 2Z_{[t_1, t_2]}$ is in $\mathcal{P}_D(M)$ for a phase $[t_1, t_2]$, Lemma 3 implies that the event $\mathcal{G}_{[t_1, t_2]}$ occurs with probability $1 - 1/R^8$. Moreover, if $\mathcal{G}_{[t_1, t_2]}$ occurs, it is easy to see that the randomized rounding solution is feasible for all $t \in [t_1, t_2]$. Since there are R phases within an epoch, the expected number of times that the randomized rounding fails any time during an epoch is $R \cdot 1/R^8 = R^{-7}$.

Suppose that we rerandomize all thresholds whenever the randomized rounding fails. Each rerandomization will cost us at most ra_{\max} in expected acquisition cost. Since the expected number of times we do this is less than once per epoch, we can charge this additional cost to the ra_{\max} acquisition cost incurred by the LP during the epoch. Thus we get an $O(\log R) = O(\log \frac{r a_{\max}}{a_{\min}})$ -approximation. This argument also works for the online case; hence for the common case where

all the acquisition costs are the same, the loss due to randomized rounding is $O(\log r)$.

D.2 Hardness of the Online MMM and Online MSM

In the online set cover problem, one is given an instance (U, \mathcal{F}) of set cover, and in time step t , the algorithm is presented an element $u_t \in U$, and is required to pick a set covering it. The competitive ratio of an algorithm on a sequence $\{u_t\}_{t \in [n']}$ is the ratio of the number of sets picked by the algorithm to the optimum setcover of the instance $(\{u_t : t \in [n]\}, \mathcal{F})$. Korman [23, Theorem 2.3.4] shows the following hardness for online set cover:

Theorem 9 ([23]). *There exists a constant $d > 0$ such that if there is a (possibly randomized) polynomial time algorithm for online set cover with competitive ratio $d \log m \log n$, then $NP \subseteq BPP$.*

Recall that in the reduction in the proof of Theorem 8, the set of long term edges depends only on \mathcal{F} . The short term edges alone depend on the elements to be covered. It can then be verified that the same approach gives a reduction from online set cover to online MSM. It follows that the online MSM problem does not admit an algorithm with competitive ratio better than $d \log m \log T$ unless $NP \subseteq BPP$. In fact this hardness holds even when the end time of each edge is known as soon as it appears, and the only non-zero costs are $a(e) \in \{0, 1\}$.

E Missing Proofs for Perfect Matchings (Section 5)

Recall that in Section 5 we wanted were given a graph and wanted, at each step, to maintain a perfect matching in this graph. We claimed an integrality gap and a hardness result for this problem; here are the proofs.

E.1 Integrality Gap for PM-Maintenance

The natural LP relaxation we use is:

$$\begin{aligned} \min \sum_t c_t \cdot \mathbf{x}_t + \sum_{t,e} y_t(e) \\ \text{s.t. } \mathbf{x}_t &\in PM(G) & \forall t \\ y_t(e) &\geq x_t(e) - x_{t+1}(e) & \forall t, e \\ y_t(e) &\geq x_{t+1}(e) - x_t(e) & \forall t, e \\ x_t(e), y_t(e) &\geq 0 \end{aligned}$$

The polytope $PM(G)$ is now the perfect matching polytope for G .

Lemma 5. *There is an $\Omega(n)$ integrality gap for the PMM problem.*

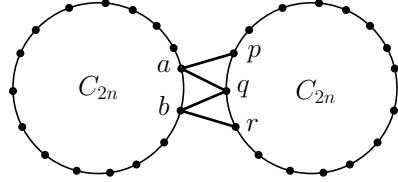


Fig. E.1. Integrality gap example

Proof. Consider the instance in the figure, and the following LP solution for 4 time steps. In x_1 , the edges of each of the two cycles has $x_e = 1/2$, and the cross-cycle edges have $x_e = 0$. In x_2 , we have $x_2(ab) = x_2(pq) = 0$ and $x_2(ap) = x_2(bq) = 1/2$, and otherwise it is the same as x_1 . x_3 and x_5 are the same as x_1 . In x_4 , we have $x_4(ab) = x_4(qr) = 0$ and $x_4(aq) = x_4(br) = 1/2$, and otherwise it is the same as x_1 . For each time t , the edges in the support of the solution x_t have zero cost, and other edges have infinite cost. The only cost incurred by the LP is the movement cost, which is $O(1)$.

Consider the perfect matching found at time $t = 1$, which must consist of matchings on both the cycles. (Moreover, the matching in time 3 must be the same, else we would change $\Omega(n)$ edges.) Suppose this matching uses exactly one edge from ab and pq . Then when we drop the edges ab, pq and add in ap, bq , we get a cycle on $4n$ vertices, but to get a perfect matching on this in time 2 we need to change $\Omega(n)$ edges. Else the matching uses exactly one edge from ab and qr , in which case going from time 3 to time 4 requires $\Omega(n)$ changes. ■

E.2 Hardness of PM-Maintenance

In this section we prove the following hardness result:

Theorem 10. *For any $\varepsilon > 0$ it is NP-hard to distinguish PMM instances with cost N^ε from those with cost $N^{1-\varepsilon}$, where N is the number of vertices in the graph. This holds even when the holding costs are in $\{0, \infty\}$, acquisition costs are 1 for all edges, and the number of time steps is a constant.*

Proof. The proof is via reduction from 3-coloring. We assume we are given an instance of 3-coloring $G = (V, E)$ where the maximum degree of G is constant. It is known that the 3-coloring problem is still hard for graphs with bounded degree [20, Theorem 2].

We construct the following gadget X_u for each vertex $u \in V$. (A figure is given in Figure E.2.)

- There are two cycles of length 3ℓ , where ℓ is odd. The first cycle (say C_u^1) has three distinguished vertices u'_R, u'_G, u'_B at distance ℓ from each other. The second (called C_u^2) has similar distinguished vertices u''_R, u''_G, u''_B at distance ℓ from each other.
- There are three more “interface” vertices u_R, u_G, u_B . Vertex u_R is connected to u'_R and u''_R , similarly for u_G and u_B .
- There is a special “switch” vertex s_u , which is connected to all three of $\{u_R, u_G, u_B\}$. Call these edges the *switch* edges.

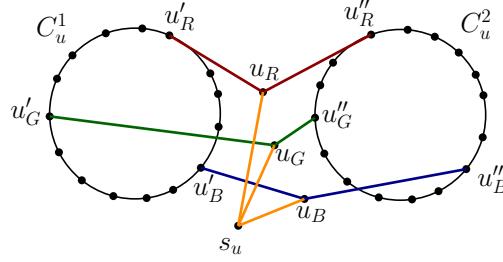


Fig. E.2. Per-vertex gadget

Due to the two odd cycles, every perfect matching in X_u has the structure that one of the interface vertices is matched to some vertex in C_u^1 , another to a vertex in C_u^2 and the third to the switch s_u . We think of the subscript of the vertex matched to s_u as the color assigned to the vertex u .

At every odd time step $t \in T$, the only allowed edges are those within the gadgets $\{X_u\}_{u \in V}$: i.e., all the holding costs for edges within the gadgets is zero, and all edges between gadgets have holding costs ∞ . This is called the “steady state”.

At every even time step t , for some matching $M_t \subseteq E$ of the graph, we move into a “test state”, which intuitively tests whether the edges of a matching M_t have been properly colored. We do this as follows. For every edge $(u, v) \in M_t$, the switch edges in X_u, X_v become unavailable (have infinite holding costs). Moreover, now we allow some edges that go between X_u and X_v , namely the edge (s_u, s_v) , and the edges (u_i, v_j) for $i, j \in \{R, G, B\}$ and $i \neq j$. Note that any perfect matching on the vertices of $X_u \cup X_v$ which only uses the available edges would have to match (s_u, s_v) , and one interface vertex of X_u must be matched to one interface vertex of X_v . Moreover, by the structure of the allowed edges, the colors of these vertices must differ. (The other two interface vertices in each gadget must still be matched to their odd cycles to get a perfect matching.) Since the graph has bounded degree, we can partition the edges of G into a constant number of matchings $M_1, M_2, \dots, M_\Delta$ for some $\Delta = O(1)$ (using Vizing’s theorem). Hence, at time step 2τ , we test the edges of the matching M_τ . The number of timesteps is $T = 2\Delta$, which is a constant.

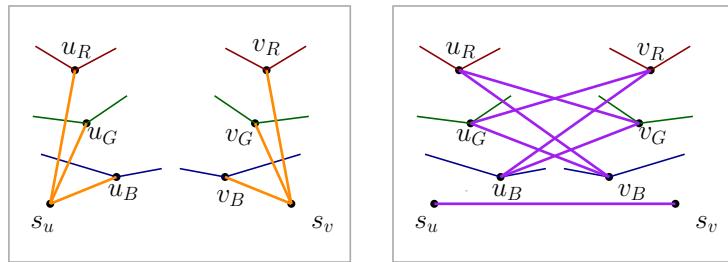


Fig. E.3. On the left, the steady-state edges incident to the interface and switch vertices of edge (u, v) . The test-state edges are on the right.

Suppose the graph G was indeed 3-colorable, say $\chi : V \rightarrow \{R, G, B\}$ is the proper coloring. In the steady states, we choose a perfect matching within each gadget X_u so that $(s_u, u_{\chi(u)})$ is matched. In the test state $2t$, if some edge (u, v) is in the matching M_t , we match (s_u, s_v) and $(u_{\chi(u)}, v_{\chi(v)})$. Since the coloring χ was a proper coloring, these edges are present and this is a valid perfect matching using only the edges allowed in this test state. Note that the only changes are that for every test edge $(u, v) \in M_t$, the matching edges $(s_u, u_{\chi(u)})$ and $(s_v, v_{\chi(v)})$ are replaced by (s_u, s_v) and $(u_{\chi(u)}, v_{\chi(v)})$. Hence the total acquisition cost incurred at time $2t$ is $2|M_t|$, and the same acquisition cost is incurred at time $2t + 1$ to revert to the steady state. Hence the total acquisition cost, summed over all the timesteps, is $4|E|$.

Suppose G is not 3-colorable. We claim that there exists vertex $u \in U$ such that the interface vertex not matched to the odd cycles is different in two different timesteps—i.e., there are times t_1, t_2 such that u_i and u_j (for $i \neq j$) are the states. Then the length of the augmenting path to get from the perfect matching at time t_1 to the perfect matching at t_2 is at least ℓ . Now if we set $\ell = n^{2/\varepsilon}$, then we get a total acquisition cost of at least $n^{2/\varepsilon}$ in this case.

The size of the graph is $N := O(n\ell) = O(n^{1+2/\varepsilon})$, so the gap is between $4|E| = O(n) = O(N^\varepsilon)$ and $\ell = N^{1-\varepsilon}$. This proves the claim. \blacksquare